

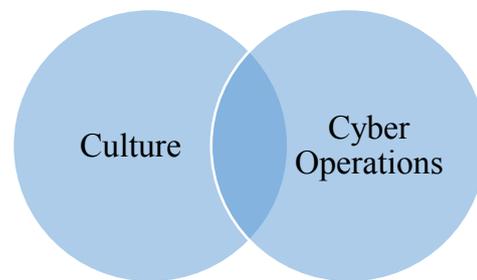
The Research Process: A Short Guide for Undergraduate Students

Joseph M. Hatfield, US Naval Academy

The research process isn't intuitive. It must be taught through the passing of transferrable skills from one person to another and perfected through the experience of trial and error. Although this document references "the" research process, there really isn't one single process that adequately characterizes every research project. This is similar to when children are taught about "the scientific method," despite the fact that there really is no such thing.¹ Scientists use many methods, these methods change over time, and scientists even disagree about which methods are the right ones. Nevertheless, there are *general* stages that *often* occur in *many* research projects; and this is enough to allow for textbooks and others to discuss general stages of scientific research. Since my own research focuses on topics within human factors in cyber operations, ethics, and other human-related issues, the examples used below will have a social science leaning. Readers performing other types of research, engineering and computer science for example, may need to abstract away from the examples used. But the general research process is agnostic with respect to subject-area.

Step 1: The Venn diagram stage

All research begins with overlapping conceptual categories. When asked what their research project is about, students typically respond by providing two or three categories between which their topic is situated. For example, a student might respond by saying "I'm doing a capstone project about how culture impacts cyber operations." Such students often *think* they have a developed research topic but in fact they are merely at what might be called the Venn diagram stage. John Venn (1834-1923) was an English logician whose circular diagrams are used to show interconnected relationships. Students should be perfectly happy to be at this stage so long as they realize that that this is just the beginning. **Real research work begins only after a research question has been carefully crafted**, and such a question has not yet been accomplished at the Venn diagram stage of research.



A research Venn diagram

Take a look at our hypothetical student's research statement again from above. They are simply saying that their yet-to-be-written research question will have "culture" and "cyber operations" as two of its topics. In fact, by putting them together they are saying that their research question rests in the intersection—the center of the Venn diagram—between these two topics. But how many questions can one ask within such an intersection? An infinite number! The key mistake that nearly every new researcher makes is to attempt to begin research at this stage. If a student goes to the library or search engine with a potential infinite number of possible research questions then almost *everything* will seem relevant. Researchers will begin filing away book after book, article after article, and quickly overwhelm themselves with mountains of information that they feel they must go through. Researchers quickly become disenchanted, leading to procrastination and eventual underperformance or failure.

The Venn diagram stage is important as an intellectual starting point. It gives researchers the opportunity to throw ideas around, brainstorm, and broadly state what concepts, topics, and categories one wants to

¹ For classic works making this point see Paul Feyerabend. 2010. *Against Method (4th Edition)*, New York: Verso Books, and Thomas S. Kuhn. 1996. *The Structure of Scientific Revolutions (3rd Edition)*, Chicago: University of Chicago Press.

work with in their project. “I’ve always been interested in human culture and in cyber operations and wonder if these are related to each other in some way.” Statements like this are indicators that one is at the Venn diagram stage. Venn circles come and go during the process of revision at this early step in the research process—and that fluidity is perfectly normal.

Step 2: Crafting a Research Question

After one has settled on the topical categories of interest and has developed a clear Venn diagram one is ready to begin the next research step—crafting a research question. By that, what is meant is a full and complete question that must begin with a capital letter and end with a question-mark. The question utilizes the conceptual categories found in one’s Venn diagram but formulates these into a complete sentence. Let’s give this a shot:

Attempt 1: How does culture impact cyber operations?

Notice here that it is written as a real question, not a statement or a fragment. It also includes the hypothetical student’s Venn categories. So at least it has the right form, and if one finds that they have written something other than a full question they must rewrite it until it possesses these features. But how do I know if this is a good research question? There is actually a step-by-step process that helps.²

First, write down the research question on a piece of paper, marker board, or chalkboard. Next, draw a box around every key term in the question, letting only completely insignificant terms (e.g. “the,” “a,” “and”) get a pass. **Note:** sometimes such terms really are significant and may change the meaning of the question. And sometimes terms can be combined if they naturally go together.

How *does* *culture* *impact* *cyber operations?*

Now one must look at the question and ask three things: (i) what facts does the question *assume* to be true, (ii) what specifically does one *mean* by each boxed term, and (iii) is each term *necessary*? Let’s start with the analysis of assumptions. Every question assumes that certain things are true which may or may not be. This is so for a very good reason. In order to give meaning to the very terms that appear within the question, certain things must be true, and it is these things that, if untrue, cause major problems later in the research process. If, for example, one’s question asks whether computer hardware vulnerabilities, as a class, are a greater problem than software vulnerabilities, the question assumes that both hardware and software vulnerabilities exist. It also assumes that hardware and software vulnerabilities are separable in some meaningful way to allow for comparisons. Since some hardware vulnerabilities can be mitigated through software updates, perhaps this is even an unjustified assumption. At the very least one must be aware of the assumptions they are making when doing research. These assumptions must be made explicit in the writing of the research paper.

The question as represented in *Attempt 1* rests upon one major assumption, that culture does in fact impact cyber operations, otherwise it would be formulated simply as “Does culture impact cyber operations?” If a researcher decides they’d like to revise their question at this stage, by dropping the “How,” such a

² My articulation of this step in the research process is indebted to the faculty of the Politics and International Studies department at Cambridge University, where I attended graduate school. It was at Cambridge that the iterative stages of crafting a research question was first made explicit to me. I claim no originality in this and fondly refer to it as “Cambridge questioning” in discussions with my own students.

revision is perfectly fine. In doing so, one has both identified a key assumption and a term that was unnecessary to the researcher's revised goal.

Attempt 2: Does culture impact cyber operations?

Now each term in the question must be analyzed (i.e. their meaning must be further refined). The three terms that stand out as in need of semantic fine-tuning are "culture," "impact," and "cyber operations." In almost every case, the terms used in initial research questions are far too broad. Researchers typically leave the context that provides significance to their words implicit, so the task of refinement involves making these contextual features explicit by adding more words to the question (e.g. adjectives) or substituting more general words for terms that are specific. For example, suppose our hypothetical research student was actually thinking about examples such as Russia's cyber operations against the US Presidential Election in 2016 or the state-sponsored *Stuxnet* attack on Iran's nuclear power facility at Natanz. In other words, it was the national security culture of state actors that the researcher implicitly meant by "culture" and it was offensive cyber-attacks specifically that they had in mind. Finally, what about the term "impact"? An impact occurs when an object hits another object. That's its literal meaning. In this case, the use is metaphorical and one should always try to avoid the use of metaphor within a research question to avoid ambiguity. The researcher, upon reflection, may state that what was meant by "impact" was a relationship between cause and effect—that is, a *causal* relationship. Taking this into account, the question can be revised as follows.

Attempt 3: Does a nation's security culture affect a nation's offensive cyber operations?

Now we're really starting to get somewhere because the question's key terms are becoming refined enough that actual research can begin to be performed. Except what exactly do we mean by "security culture"? Perhaps someone has already written about national security cultures? A query into the academic literature reveals that a number of authors have discussed the concept of security culture in a way that is helpful for this research.³ In that case, the term "security culture" can be left in the research question because an operative meaning has been found and will be explicitly articulated within the paper. Also, since "a nation's offensive cyber operations" encompass a huge number of things, this too needs further refinement. Does this mean, for example, the manner in which such operations are ordered (e.g. the chain of command), the resources allocated (e.g. as a proportion to GDP), or the types of operations utilized (e.g. DDoS, Phishing, Sybil Attacks)? Choosing the last of these, the question can be further revised.

Attempt 4: Does a nation's security culture affect the types of offensive cyber operations it employs?

Perhaps at this point it seems clear that in order to answer this question one would need to look at every nation, or at least a sufficient sample of nations, and this is something the researcher does not think they have the time or resources to do. A simple way to solve this would be to focus on a smaller subset of nations—perhaps even just one. The nation of choice may be determined by interest, but equally by mundane factors such as one's access to research material.

Attempt 5: Does Russia's security culture affect the types of offensive cyber operations it employs?

At last! This is a properly scoped question for a research project. Further refinements can always be made, of course. If one were an historian one might vary the question slightly by specifying a time-period

³ Gustafson, K.C. 2010. "Echo of Empires: Russia's Inheritance of Byzantine Security Culture." *Journal of Slavic Military Studies* 23(4): 574-596. Welch, Stephen. 2013. "Political Culture." In *Intelligence Elsewhere: Spies and Espionage outside the Anglosphere*, Philip H. J. Davies and Kristian C. Gustafson (eds), pp. 13-26.

(e.g. “Did Russia’s security culture affect the types of offensive cyber operations it employed from 1991-2016?”). But this question is clearly refined and the researcher is far better prepared to begin research having attained a refined research question than they were at the Venn diagram stage.

In my experience, the refinement of an adequate research question **may take weeks of hard work**. That investment pays huge dividends as the research process moves forward! This should not be seen as a static process. Later findings may cause a researcher to go back and further refine or even change their question still more. It is also beneficial to share one’s attempted research questions with others, particularly with experienced researchers and subject matter experts. Outsiders often see meanings, assumptions, and aspects of preliminary research questions which are hidden from the researchers drafting the questions.

Step 3: The Literature Review

Now that an adequate research question has been crafted researchers must find out if anyone has already answered it before. To do this one does something called a “literature review.” The word “literature” here doesn’t have to do with Shakespearian plays or the *The Adventures of Tom Sawyer*. Rather, it signifies the written works relating to a certain topic or subject. The literature on English poetry will be vast, while the literature on household plumbing that references pipes of a certain diameter will be quite narrow. **A literature review is an analysis of what has been published on the topics central to one’s research question.** It involves searching libraries, search engines, and other databases for material relevant to answering the research question. As noted above, a library’s query tool is a great place to begin (and often end) a literature review, as it provides access to a large sample of books, academic articles, and other resources. The use of Boolean searches [for example: (“security culture”) AND (“nation state”) OR (“nation-state”)] and other techniques allow one to whittle down a potentially daunting number of sources to a few dozen that are specifically relevant to one’s question.

When conducting a literature review it is rarely necessary to read an entire book or academic article to discover whether a source is germane to one’s research question. Reading through the abstract, using key word searches, inspecting the bibliography or index, or skimming the article or book chapter are time-saving analysis techniques.

Once one discovers a source of interest and relevance, one can create what researchers call an “annotated bibliography.” An annotated bibliography includes the bibliographical reference information followed by descriptive and evaluative paragraphs or sentences (i.e. the annotation).

Annotated Bibliography Example

Gustafson, K.C. 2010. “Echo of Empires: Russia’s Inheritance of Byzantine Security Culture.” *Journal of Slavic Military Studies* 23(4): 574-596.

Gustafson argues that the way the Byzantines managed their security and intelligence was a function of the state’s political culture, the same political culture that was later inherited by the Russian state, and which has served the Soviet and subsequent post-Soviet Russian state (Gustafson 2010). This analysis supports the conjecture that a deeply ingrained security culture underlies superficial changes in governmental structure, from the Byzantines to today. If that is so, the tactics utilized by the modern Russian state (e.g. offensive cyber operations) ought to bear the mark of this security culture.

An annotated bibliography can consist of as little or as much information and commentary as one deems appropriate. At a minimum, however, one should try to achieve three things:

- **Determine whether the research question has been answered before.** If so, the researcher must decide whether they agree with the answer or not. If they agree, there is no more work to be done and the research question must be changed—even if only slightly—to maintain its originality. If the researcher doesn't agree, their originality rests, partly or solely, on their ability to argue against an established point of view.
- **Motivate the reader's interest in the question.** Good researchers show their audience why the research question is important, which can be done in a number of ways. Perhaps the question has never been asked or answered. Perhaps it has been asked but the answers researchers have offered were incomplete or assumed something that is actually false. By showing any of these to be the case through one's literature review one has generated interest in the minds of one's readers.
- **Provide sufficient background to the reader so that they are able to follow one's argument.** Most readers will need a certain amount of information given to them in order that they might understand what the research attempts to show or explain. For example, in the case of the present research question a reader cannot be expected to know that other research has argued that contemporary Russia shares with the Byzantine Empire its security culture. The literature review can discuss works, such as the Gustafson example above, that establish that connection, thereby freeing the researcher from the obligation to re-establish that idea and allowing them to move forward with that idea's implications for offensive cyber operations.

Once a researcher has a sufficient annotated bibliography one has effectively established the previous scholarly work predating the researcher's current project. Thus, combining the last two steps, one has asked or refined a new research question and shown why attaining the answer to that question is valuable. And in the annotated bibliography one has generated the reference material necessary to properly inform one's reader about the background information required to understand the current research project.

Step 4: Generating hypotheses

Not all research is empirical in nature. Research may rearrange known facts in interesting new ways to argue for conclusions that nobody has thought of before. Albert Einstein's *Special Theory of Relativity* was of this kind. Fruitful research conducted on questions involving definition, theory, ontology, history, or ethical and normative topics often fall into this category. However, a great deal of research seeks to answer questions that do ask for new data and the parsing of new observations. In such cases, researchers should seek to generate hypotheses prior to performing the data collection and analysis.

Hypotheses are specially-crafted guesses at what the data will reveal. More accurately, an hypothesis is a supposition or proposed explanation made on the basis of limited evidence as a starting point for further more detailed investigation. The suppositions shouldn't come out of thin air but should be backed by the literature or be reasonable extrapolations from the research question's meaning. Multiple hypotheses are often generated for a given research question. For example:

Does Russia's security culture affect the types of offensive cyber operations it employs?

H₁ Russia's security culture strongly determines what types of offensive cyber operations it employs.

H₂ Russia's security culture partly determines what types of offensive cyber operations it employs.

H₃ Russia's security culture does not determine what types of offensive cyber operations it employs.

Good hypotheses help researchers by articulating testable answers to the research question asked. In doing so, correct investigative methods are more readily identified. An hypothesis must be specific enough to be testable, and the criterion for its truth or falsity must be made clear *prior* to the research commencing. In the examples above (H₁- H₃) the phrases “strongly determines,” “partly determines,” and “does not determine,” are terms of ordinary language that need to be directly linked to criteria for their truth or falsity. In science this is sometimes called *operationalizing* terms.⁴ While crafting the example research question, the meaning of the term “security culture” was found within the academic literature, and supposing that further research showed Russia’s security culture to consist of eight attributes (e.g. the belief that notable state affairs are done in secret; the view that the governmental regime, not the people, is the referent object for security; and so on), then the falsification criterion for H₁ might be operationalized by postulating that if six or more attributes are found to be present in at least 80% of the analyzed cases then this will mean that H₁ has been confirmed. In other words, this is the operational meaning of the term “strongly determines” in H₁. Operationalizing helps bring clarity to key terms and avoid the problem of confirmation bias—the tendency to interpret data in such a way that it confirms what one already believes.

Step 5: Identifying the Correct Methods

Most researchers, even many published ones, fail to distinguish discussions about "methods" from those about "methodology" and they tend to use the word "methodology" to incorrectly describe the methods they use. But these are different concepts. A method is a specific structured way of performing some research task while a methodology (from the Greek: *λόγος*, Logos, meaning "reasoned speech" about methods) is a philosophical view *about* methods, including their content, employment, and legitimacy.⁵

Many methodological debates center around the nature of knowledge and the human mind’s relation to the world. A methodological *positivist*, for example, argues that the mind is sufficiently separate from the world it observes through sensory experience, and by extension the data gathered thereof, that when performed correctly, research methods, be they quantitative or qualitative, lead to neutral objective conclusions.⁶ Descriptions of observational data are theory-neutral, for the positivist, and can therefore confirm or falsify empirical hypotheses in a straightforward way. By contrast, a methodological *reflectivist* argues that observation is theory-laden; that is, how we describe what our senses tell us depends upon other beliefs we already hold.⁷ Therefore, research methods do not observe neutral data and research conclusions are strongly biased by the interests and beliefs of the researcher. Positivism and reflectivism are only two positions along a spectrum of methodological views.

Most undergraduate research assumes a positivist methodology, particularly in the so-called “hard” sciences, such as physics, chemistry, and the like. This is partly due to the relatively uncontroversial nature of the observation data with which these sciences work, leading to less concern about observational

⁴ The classic statement on operationalization comes from Percy Williams Bridgman. 1927. *The Logic of Modern Physics*. New York: MacMillan.

⁵ One of the best discussions of the distinction between method and methodology appears in Patrick Thaddeus Jackson. 2016. *The Conduct of Inquiry in International Relations (2nd Edition)*, New York: Routledge. Although Jackson’s discussion takes place within the context of political science it can be applied elsewhere. See specifically pp. 26-57.

⁶ For a defense of the philosophical basis of this view see Otto Neurath’s “Protocol Sentences,” pp. 199-208, and A. J. Ayer’s “Verification and Experience,” pp.228-243, in *Logical Positivism*, edited by A.J. Ayer, 1959. New York: The Free Press.

⁷ For a defense of the philosophical basis of this view see W.V.O. Quine and J. S. Ullian. *The Web of Belief (2nd Edition)*, particularly pp. 20-49. 1978. New York: McGraw-Hill, and Wilfrid Sellars, 1956. *Empiricism and the Philosophy of Mind*. Minneapolis: University of Minnesota Press.

bias. Positivism has also become the assumed methodological position for much of the rest of society, even for those who have never heard the word “methodology.” This probably has to do with the exalted cultural position enjoyed by the hard sciences. Researchers working in the social sciences, by contrast, have been much more open to non-positivist methodological views. This again partly reflects the more controversial observation data that are the concern of anthropology, psychology, economics, political science, and the other social sciences.

Setting methodology aside, one can split research *methods* into two major groupings—quantitative methods and qualitative methods. Quantitative methods apply to information that can be expressed in numerical categories or relationships (e.g. numbers and statistics). Qualitative methods typically involve the careful in-depth study of a much smaller number of cases. There is a natural research tradeoff between numerical breadth and qualitative depth. A study of the password habits of thousands of computer users allows researchers to make very powerful inferences about the password habits of large populations of computer users. However, simply in order to run such a large study, researchers must necessarily reduce interaction with their research subjects—for example, by using questionnaires that can be disseminated and completed *en mass*. On the other hand, by performing in-depth interviews with a few dozen computer users, each carefully selected as representatives of the population about which one wants to make inferences, researchers may gain deeper insight by gaining access to much more detailed and contextual information about why people make password-related decisions. Ideally, researchers could perform in-depth interviews with thousands of people; but in most cases the time and resources involved are prohibitive.

There are many types of *quantitative* analysis. A statistical regression analysis, for example, is a quantitative method used to determine the relationship between variables. More specifically, regression determines the degree of variation in an output variable (y) that is attributable to one or more input variables (x). Descriptive statistics allow one to calculate numbers that tell researchers about the distribution of a data-set, for example a sample mean, range, and standard deviation. Z-tests and t-tests are statistical comparisons that allow one to determine whether a difference between two sample means is statistically significant, and so on. Quantitative analysis is appropriate when one’s research question asks about quantifiable relationships and the researcher has access to a sufficiently robust amount of numerical data to answer the question. There are lots of excellent data analysis books and websites dedicated to quantitative methods, whether general statistics textbooks or books dedicated to specific types of data analysis (e.g. time-series analysis).⁸

Qualitative analysis generally trades numerical breadth for contextual depth. By focusing on a comparatively smaller number of cases, qualitative analysis allows researchers to uncover a tremendous amount of data that superficial numerical relationships cannot. This adds textured nuance to our understanding of the phenomena about which we seek to gain knowledge. The paradigmatic case study method, for example, uses the careful analysis of a small number of representative cases, of, say, election interferences by Russia, to make general inferences about Russian election interferences—or even about election vulnerabilities broadly. Structured interviews, similarly, allow researchers to understand the motives, beliefs, desires, and intentions that form the basis of human behavior—for example, in the password-related choices example above. Discourse analysis is a qualitative method that studies the use of language relating to a topic of interest. By tracking how language is used, researchers may gain deep

⁸ A great short primer in data analysis from an industry point of view is Donald J. Wheeler. 2000. *Understanding Variation* (2nd Edition). Knoxville, TN: SPC Press.

insights into what people may think, feel, or believe—and by extension how they behave. There are a number of good resources dedicated to qualitative methods.⁹

Some research questions require both the employment of quantitative and qualitative methods to find a satisfactory answer. This is often referred to as “mixed-methods” research. There is no *a priori* reason to exclude one or the other type of method. Rather, the research method chosen should be a natural extension of one’s research question.

Step 6: Conducting the Research

At this stage, one performs the data collection and analysis, writes the computer code, gives sufficient thought to ethical, normative, or nonempirical scholarship, and so on. There is less to say about this stage of the research process because this is the stage of *action*—performing the tasks necessary to generate results.

Research must be conducted in an ethical way. Researchers are expected to exhibit complete integrity in all of their activities. Academic dishonesty, including plagiarism of any type, violates this principle. Plagiarism is the act of presenting another’s words, ideas, or work as one’s own, whether accidentally or deliberately. Another issue that sometimes arises is the artificial generation of data-sets or results. This too is never justified and is one of the central reasons behind the scientific norm that one must be ready to make one’s data available to other researchers so that the study can be independently reproduced.

Research involving human subjects must be approved by an Institutional Review Board (IRB). Even seemingly innocuous research tasks, such as asking people’s permission to participate in a survey, must receive IRB approval before commencing. The IRB review and approval process can be quite lengthy, so adequate time must be allocated.

Step 7: Writing up the Findings

At this point, one has moved from the Venn diagram stage, through the generation of the research question, the literature review, and hypothesis generation. The correct methods were identified to answer the research question and the subsequent data gathering and analysis has been undertaken. But all of this is ultimately meaningless if the results aren’t shared with others in a way that is clear and understandable. The research process culminates in a research paper (or book if it is a large project) that conveys not just the results but also a good amount of information from the previous six steps.

Many undergraduates receive scholarly information strictly from textbooks or academic lectures, which nicely summarize and make intelligible to undergraduate audiences peer-reviewed scholarly research. However, this means that undergraduates are sometimes less directly familiar with the sort of academic writing they are then asked to perform for their own senior thesis, capstone project, or other culminating undergraduate research effort. It will therefore be useful to describe what makes academic writing unique and how one might improve one’s own writing.

What is academic writing?

Academic writing differs from other writing in its epistemic role, its manner of expression, and its standards for publication. Let’s take these one by one. The term “epistemic” comes from the Greek word ἐπιστήμη (pron. epistēmē) meaning “knowledge,” “science,” or “understanding.” Academic writing is

⁹ An excellent source on qualitative research is Gary King, Robert O. Keohane, and Sydney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton, NJ: Princeton University Press.

ultimately meant to **persuade an audience of some conclusion by means of valid logical inference** rather than mere rhetoric, sophistry, or appeals to tradition or authority. This is called making an “argument” and the conclusion is called a “thesis.” The term “argument” here isn’t the same as when someone says they’ve had an argument with their sibling or spouse. Rather, in academia an argument is a line of reasoning from premises or established facts to a conclusion or thesis.

This knowledge-building role means that academics must express themselves differently than novelists, speech writers, entertainers, and so on. Academics must back up any claim they make by referencing the primary study or observation that first established the referenced facts, or by referencing secondary sources that reported on these primary sources. So academic writing uses **sourcing, citation, and bibliography** to keep track of the origins of every claim made.

Academic writing is also often full of **words that may be unfamiliar** to undergraduate audiences. Why? Because undergraduates are usually relative beginners and most academic writing has as its intended audience older more experienced scholars. When anyone encounters an unfamiliar word they must either attempt to gain an understanding of what the unfamiliar word means from its context or look it up in a dictionary.¹⁰ This is true whether one is reading a newspaper, Tolstoy’s *Anna Karenina*, or a highway billboard. Reading should be an *activity* (notice the root of this word – active) not a passive enterprise. Language isn’t merely a container or vehicle within which ideas are placed for the purpose of communication. Rather, language *constitutes* these ideas, and the more words one becomes familiar with the more one experiences.

Academic writing also differs from other forms in the standards required for its publication. Academic journals and books undergo a process of **blinded peer review**. That is, the author’s name and institutional affiliation are typically removed from the submitted manuscript and then sent to multiple experts within the field, whose identities are also withheld from the author, for a blind review. Most manuscripts receive meticulous criticism and are either rejected or are sent back to the author with a request to resubmit the manuscript after dealing with any deficiencies found by the reviewers. Nonacademic writing, such as newspaper and magazine articles, blog posts, website content, and so on, usually need only satisfy the editors, who always know the author’s identity—potentially allowing all sorts of bias into the process. This isn’t to say that blinded peer-review is perfect, but the process of publishing in academia is more rigorous than elsewhere. In academic writing, nonacademic sources should *supplement* but not replace proper academic sourcing. Nonacademic sources can helpfully facilitate discussion or make vivid to an undergraduate audience conclusions established by rigorous academic work.

Furthermore, academic writing must **break new ground**, however incremental, to be published; and this is often achieved, particularly in the sciences, by leveraging the expertise of multiple authors. Novelty typically takes two forms. First, an academic may perform a new study that others haven’t yet performed, thereby introducing the world to a new set of facts. Second, an academic may rearrange old facts in a new way, thereby showing the world that new inferences are possible. Sometimes, of course, both of these are accomplished within the same piece of academic writing.

One final point to make about academic writing is that unlike other forms of expression, such as e-mails, texts, and in everyday conversation, the **use of the first-person**—the “I/we” perspective—in academic writing is kept to a minimum. In high school this is often taught in an overly strident or draconian

¹⁰ For a wonderful essay justifying the use of words that might be unfamiliar to an author’s audience see William F. Buckley’s “I Am Lapidary But Not Eristic When I Use Big Words,” *New York Times*, 30 Nov. 1986, <http://movies2.nytimes.com/books/00/07/16/specials/buckley-bigwords.html>

manner, as if the word “I” never appears in scholarly work. This is simply false; it does so appear. However, since the purpose of academic writing is to persuade others through scientific inference, the third-person—the “he/she/it/they” perspective—is the most appropriate point of view to adopt. Yet if a situation arises where the first-person is needed it is best to simply use it rather than create a tortured sentence just to avoid its use. The rule is: if you mean “I,” say “I,” but you shouldn’t mean “I” very often in academic writing.

Resources for improving one’s academic writing

Have you ever heard someone describe themselves as “not a math person”? Research has shown that to be a myth. Unless we’re talking about higher-end mathematics (e.g. quantum physics), nearly everyone can be great at math so long as they practice and put in the work. The same is true of writing. **Good writing is difficult and takes practice**, but nobody should describe themselves as “a bad writer” unless they simply mean that they haven’t put in the time, effort, and practice to become a good writer. This is as true of academic essays as it is of poetry and fiction.¹¹

Great writers are great readers; **to write well one must read well**. Reading well begins by selecting great texts from which one absorbs—sometimes without realizing it—correct grammar usage, new vocabulary, effective sentence construction, appropriate paragraphing, and a host of other skills. If one seeks to write excellent essays, one ought to read accomplished general essayists such as Ralph Waldo Emerson, C.S. Lewis, or George Orwell.¹² Those seeking to improve one’s writing of essays in a particular subject—philosophy, for example—should read skillful essayists in that subject, such as Richard Rorty, James Rachels, or Michael Walzer.¹³

Reading well also requires time, enough to fully absorb an essay’s content without distraction. This means actively planning to read when one is at one’s best—unhurried, alert, and able to immerse oneself. An essay doesn’t have to be read in one sitting, and great essays should be read multiple times. Some writing is so profound, so full of ideas, that one cannot really read them for the first time until one has already read it once or twice before. Plato’s *Republic* is writing of this caliber, as is the poetry of Emily Dickinson.

Reading well is also an active process. The mantra I give to my students is: **you aren’t reading unless you’re writing**. Students should have a physical notebook next to them as they read so that they can record the author’s thesis, main points, and, even more importantly, capture the student’s own thoughts,

¹¹ Edgar Allan Poe wrote a wonderful essay entitled “The Philosophy of Composition” describing the process required to write his most famous poem, “The Raven.” <https://www.eapoe.org/works/essays/philcomp.htm>

¹² Recommended essays from each author: “Character” by Ralph Waldo Emerson, <https://emersoncentral.com/texts/essays-second-series/character/#complete-essay>, “The Inner Ring” by C.S. Lewis, <https://www.lewissociety.org/innerring/>, and “Politics and the English Language” by George Orwell, https://www.orwell.ru/library/essays/politics/english/e_polit.

¹³ Recommended essays from these authors: “Philosophy as a Kind of Writing: an essay on Derrida” by Richard Rorty, <https://www.scribd.com/doc/187334534/Richard-Rorty-Consequences-of-Pragmatism-Essays-1972-1980-1982>, “The Challenge of Cultural Relativism” by James Rachels, <https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=3&ved=2ahUKewiglvD6ldnmAhXxYN8KHT-PBZ0QFjACegQIBBAH&url=https%3A%2F%2Ffrintintin.colorado.edu%2F~vancecd%2Fphil1100%2FRachels1.pdf&usg=AOvVaw3zrJzbzr5z2hS0ztgLf0PI>, “Political Action: the problem of Dirty Hands” by Michael Walzer, https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=4&cad=rja&uact=8&ved=2ahUKewi6sPqdltmAhVDZN8KHTeMAa4QFjADegQIAxAl&url=https%3A%2F%2Fwww.sss.ias.edu%2Ffiles%2Fpdfs%2FWalzer%2FPolitical-action.pdf&usg=AOvVaw0t9me7nxdc_SBCt5FpNb-x.

observations, and potential objections.¹⁴ Very few people can read through academic writing and remember the argument's details weeks or months later without having them written down. By including both information about the text and one's own observations, thoughts, and criticisms, the notebook keeps a record not just of the text but of *one's reading of the text*. Adding the bibliographical information at the top and referencing page numbers along the side next to each thought or main point gives the reader a handy reference tool to use in one's later writing.

Great writers strive for **clarity of expression**. George Orwell's previously referenced essay "Politics and the English Language" enumerates six rules writers should follow to achieve this:

1. Never use a metaphor, simile, or other figure of speech which you are used to seeing in print.
2. Never use a long word where a short one will do.
3. If it is possible to cut a word out, always cut it out.
4. Never use the passive where you can use the active.
5. Never use a foreign phrase, a scientific word, or a jargon word if you can think of an everyday English equivalent.
6. Break any of these rules sooner than say anything outright barbarous.

Much academic writing runs afoul of Orwell's six rules, but most *great* academic writers follow his advice, even when appearances seem otherwise. For example, the reason William F. Buckley Jr. (in the essay referenced in the footnote above) uses a term like "eristic" [ERISTIC: (i ris/ tik) adj. of or provoking controversy, or given to sophistical argument and specious reasoning] to describe a judge's behavior, is because, Buckley explains, the word captures exactly the meaning he wished to convey. No other word would do. Absent any equivalent word, the writer is left crafting a longer descriptive phrase which would violate Orwell's second rule. Of course, Buckley was likely aware that by choosing "eristic" some readers would be forced to look up its meaning in a dictionary. But he accepted this as the price to pay for semantic precision. Similarly, Orwell's warning against the use of "pretentious diction" (e.g. foreign words, scientific words, or jargon) is based upon his view that such vocabulary is often "used to dress up a simple statement and give an air of scientific impartiality to biased judgements." But note that this is not violated by, for example, one doctor's note to another stating that her patient "has contracted *varicella*" (chicken pox). This is because, due to their common training, both doctors are presumed to have this term as part of their working vocabularies.

The secret to writing well rests in three tasks: (a) choosing the correct audience with which to communicate, (b) using an appropriate real person (preferably someone one knows well) to represent that audience, and (c) always keeping that **representative person** in mind as each sentence is written. The representative individual has been appropriately selected if she possesses the same background information, education level, vocabulary, and set of interests as the audience with whom the writing seeks to communicate. By writing *to her* one will ensure one is effectively choosing the right words, explaining what needs to be explained, and so on. As John R. Trimble, the author of my favorite book on writing, notes, bad writing mostly stems from an author's forgetting that they are trying to communicate to a real audience.¹⁵ Following these three steps helps mitigate this pitfall.

¹⁴ Research continues to show that longhand notes are superior to computer notetaking. Pam A. Mueller and Daniel M. Oppenheimer. 2014. "The Pen is Mightier than the Keyboard: Advantages of Longhand over Laptop Note Taking." *Psychological Science* 25(6): 1-10.

¹⁵ John R. Trimble. 2010. *Writing with Style: Conversations on the Art of Writing (3rd Edition)*. New York: Pearson.

Great academic writing has a **clear structure**. Such organization helps the reader determine what argument is being made, what evidence is being offered, what portion of the writing simply describes work by previous authors, and what portion breaks new scientific ground. The importance of structure can be seen in the way academic journal articles are organized into commonly occurring sections. Students should familiarize themselves with these:

- **Abstract:** a short summary of the article, its thesis, and main point appearing at the beginning of the article (often in smaller or italicized font).
- **Introduction:** academics often begin articles by stating their research question, by stating their thesis plainly up front, or motivating the reader to want to continue by showing a gap in the literature that the article promises to fill.
- **Literature review:** academic articles usually review past academic work on the question taken up by the present article, or, if the question hasn't been asked before, the article reviews work closely related to its own argument. Sometimes it is within the literature review that the author discusses gaps in the scientific work that the article promises to fill.
- **Data analysis or research method** (often called "Methodology," errantly as was explained above): if a new study has been conducted, this section explains how the data was gathered, processed, and analyzed. Sometimes authors attempt to justify their methods as legitimate for the task at hand, the data type, or the question asked.
 - Sometimes methods are used that may not be familiar to readers. If, for example, a statistical method is unfamiliar to a reader, who is left puzzled by tables of resultant coefficients, they must either research the method themselves or trust the author (and the peer-review system of publication) to have done the analysis correctly and made valid conclusions from the data.
- **Conclusions:** this section combines the literature review (known facts) with the data analysis (new facts) to produce new or novel claims, to refute older conclusions, or to open up the possibility of new and potentially more fruitful questions.
- **Final summary:** academic writing often ends by summarizing what the article accomplished, walking the reader through each of the aforementioned steps.

What constitutes a good academic essay?

A good undergraduate essay clearly states its thesis early (even in the first sentence) and telegraphs to the reader the argument's main points which the author will make in order to justify the thesis. It then does what it says; it follows these steps, using academic sourcing to contextualize the author's argument against what others have said. As stressed earlier, nonacademic sources, such as websites, news articles, and so on, are used only to *supplement*—not replace—proper academic sourcing. Toward the end, or throughout, good essays will discuss counterarguments that can be made against the thesis, its main points, or any assumptions employed. Good essays don't necessarily have to provide an extensive rebuttal to these, and a simple "it is beyond the scope of this essay to address this objection" will sometimes suffice. Good essays end by providing a summary of the essay's thesis and main points, as well as recommending new questions for further research, as appropriate. There is no cut and dry answer to the question of how many academic sources are required since this will depend upon how sources are used and the scope of the argument.

Tarak's Tree

Tarak Barkawi, Professor of International Relations at the London School of Economics, introduced me to a wonderful mnemonic in the form of a Christmas tree for thinking about academic essays. I call this mnemonic "Tarak's Tree" in honor of him. Thinking of Tarak's Tree may help one remember to include the most critical items necessary for an effective academic essay.

The base of the tree is the research question. Connected perpendicularly to the base is the trunk of the tree which represents the thesis or answer to the research question. Along the sides of the trunk, supporting the thesis, are branches which represent the argument's main points. Along each branch are ornaments that represent the data used to support each main point, whether that data is newly gathered or referenced from another academic's work. At the top of the tree is the star which represents the self-critical standpoint that is required of any academic essay. That is, the star represents the need to make any assumptions explicit and to be humble by being up front about both the strengths and the weaknesses of one's findings.

