# The standard issue

Thomas Basbøll

2012–2013

## Contents



## 40 Paragraphs

"The world exists to end up in a book," said Stéphane Mallarmé. For today's scholars, it might seem as though it exists to end up in a journal article. But it would be more accurate to say that their objects of inquiry are constructed precisely in order to end up in such an article. It is the standard unit of knowledge-production.

Poets produce poems. Scholars produce articles. Both convert experience into a particular kind of expression with a particular kind of form, one that is recognizable to their readers. A well-written journal article will present a single, easily identifiable claim; it will show that something is the case. And it will provide an argument not just for the truth of that claim but for its relevance for a particular line of inquiry. It will also situate both the claim and the line of inquiry in a world of shared concern that goes beyond the narrow, scholarly interests of both the writer and the reader. Within those narrow limits, however, it will respect the field's theoretical and methodological commitments. Before it is over, it will offer a simple one-paragraph statement of the argument for the central claim (thesis) of the paper, one that assumes that the very knowledgeable and highly intelligent reader has understood the rest of the paper.

The article will consist of roughly 40 paragraphs. Five of them will provide the introductory and concluding remarks. Five of them will establish a general, human background. Five of them will state the theory that informs the analysis. Five of them will state the method by which the data was gathered. The analysis (or "results" section) will make roughly three overarching claims (that support the main thesis) in three five-paragraph sections. The implications of the research will be outlined in five paragraphs. These are ball-park figures, not hard and fast rules, but "knowing" something for academic purposes means being able to articulate yourself in roughly these proportions.

The scholarly conversation depends on respecting these proportions. That is: scholars expect to be talked to in a particular way, they do not expect a unique, transformative literary experience. Jonathan Mayhew has rightly described literature as the kind of writing that "kicks your ass with its transformative power". Academic writing serenely disdains to kick your ass like this. Rather, as I'm fond of saying, it is the purpose of a journal article to artfully disappoint our expectations of a particular object of inquiry. If you want to be a scholar, it is a good idea to learn that art.

### The first three paragraphs

It is difficult to overstate the importance of a good introduction. If your reader does not have a good sense of your argument by the end of the third paragraph (before reading the 600th word), there is something seriously wrong with your paper. Or, perhaps more tellingly, *if you are unable to* outline your argument straightforwardly and clearly in three paragraphs, you will be unable to write a good paper. When I talk about what a scholarly article is, I always use the opportunity to sketch "the ideal introduction". It consists of exactly three paragraphs and no more than six-hundred words.

The first paragraph tells us about the world we are living in. This should obviously be the world that your paper helps us to better understand. It's the world that needs to be understood in precisely the way you understand it. But in this paragraph we (your readers) don't want this understanding, we just want a recognizable description of the world we share with you. Talk to us like we only need to be reminded that this is where we live. It should be familiar to us and based on widely available sources. While you should avoid the letter of a statement like "We live in a world of ..." or "Ours is an age of ...", this is very much the spirit of the first paragraph. It's a time for commonplaces; it provides a shared place for you and your readers. In an important sense, you are here describing the *practices* that your paper is about. And these practices are interesting because there is some *problem* with them.

The second paragraph tells us about the science that studies this world. It summarizes the body of scholarship that has taken an interest in the problem that is described in the first paragraph. There are two good ways and one common but bad way to structure this paragraph. It can state either a constitutive consensus in the literature or a just as constitutive controversy. Scholarship will normally be characterized either by broad agreement about some issue (which your work will then challenge) or by a standing disagreement (where your paper will provide support to one side). Many papers these days begin by identifying a "gap" in the literature (which the paper then proposes to fill), but this is a false start. The gap is only interesting because what you have found there bears upon some interesting consensus or controversy. So you should fill in the gap in advance (i.e., in this second paragraph) with the *theoretical assumptions* that shape your readers' *expectations* of your subject matter. Indeed, if the first paragraph is about practice, this paragraph is about *theory*; the problem persists despite precisely this theory.

The third paragraph tells us about your paper. "In this paper, I show that..." is a nice, tight way to do this. Notice that supporting such a sentence requires you not to offer evidence but to outline your paper; it's a statement about your paper not about your object. So here you have to say something about, especially, your *method* (what have you *done* to put yourself in a position to know you are right). It should also briefly sketch the content of each section of the analysis (what have you discovered to support your conclusion) and leave us with a good sense of the implications of the paper as a whole (a paper will normally have

a section for implications, so you may just summarize that section). The implications may be either theoretical or practical: you may show that practice ought to fall in line with a perfectly good theory, thus solving the problem by making the world a more "ideal" place, or that the theory has be adjusted to better capture the "real-world" practices, thus at least acknowledging the problem. Or you may argue for some combination of such implications.

These three paragraphs, finally, should each be organized around a claim that can be expressed in a single, declarative sentence. The rest of the paragraph merely supports that claim. Notice that the *thesis* of your paper is stated only within a larger claim about its being the thesis of your paper. And that claim has been nested in a claim about the world and a claim about the research that has already been done about that world. Since the world is construed in terms of some interesting *problem*, there should be no need for an explicit "statement of the problem". But if your editor (or teacher) insists, there's no harm in providing it.

### The Introduction & The Conclusions

Just this week, during one of my undergraduate workshops, I had something of an epiphany about the rhetoric of an academic article. I've long argued that the *only* difference between the introduction and conclusion is a rhetorical one. The two sections *say* the same thing, which is to say, the conclusion does not contain information that the introduction has not already presented; but they say it differently or, more precisely, they address different audiences. And the difference is a very small and precise one: the reader of your conclusion is the same reader as the reader of your introduction except that the reader of your introduction has *not yet* read your paper and the reader of your conclusion has *just* read it. Beyond that, I used to say, you're free to write the conclusion as you choose ... oh yes, except that you only have two paragraphs to accomplish what you did with three in your introduction.

But look what I discovered while workshopping a conclusion this week. We began by sharpening the introduction to give us the ideal form. The first paragraph provided a description of the *world* in which the paper is needed. "We live in an age of..." it might (too) typically begin. The second paragraph introduces the relevant *science*: "Scholars agree that..." (or, conversely, "Current scholarship is divided about...") It is only the third paragraph that introduces your thesis: "In this paper, I show that..." Notice that until the third paragraph, no mention is made of you or your thesis; you are positioning your thesis in a world of shared concern and a body of current scholarship.

And notice also that even in that third paragraph you don't actually *state* your thesis. The key sentence of that paragraph is not a claim about your *object*; it is a claim about your *paper*. This means that the relevant kind of support is provided by a description of your paper ("After recalling the recent history of efforts to ..., I will outline my theoretical framework. The analysis uses the method of ..., which gives access to ... On this basis, I conclude that ... and emphasize the importance of ... in rethinking best practices in this area.") It is not actually an argument for the truth of your thesis. It is a description or outline of such an argument.

I have said before that the first paragraph of your conclusion and the last paragraph of your introduction should mirror each other. In the introduction there should be a paragraph about what you *will say*, while in the conclusion the same paragraph should tell us what you *have said*. But this can be quite boring ("In this paper, I have shown that..." etc.) We now arrive at the epiphany.

"In this paper, I show that..." will continue with some proposition. "In this paper, I show that organizational designs not market forces are to blame for the financial crisis." That proposition, of course, states the major thesis of the paper. Now, what your paper should be putting you in a position to do is precisely to state that thesis, plainly and straightforwardly. After reading your paper, the reader should be in a position to understand a simple, efficient, 6-sentence, 150-word argument for your thesis. Let, **the penultimate paragraph** of your paper, i.e., the first paragraph of your conclusion be that argument.

Its key sentence will be your thesis. "Organizational designs not market forces are to blame for the financial crisis." Leave out explicit mentions of method and instead state the major empirical claims that your method allowed you to discover. Also leave out any explicit mentions of theory, but use theoretical terms as though they are part of the vocabulary shared by you and your reader (if your paper is any good, by this point they are.) Three of the sentences will normally simply state the sub-theses that your analysis arrives at. However, you choose to do it, keep it simple. This is the moment when you show the reader that s/he's understood what you've been trying to say. Here's the simplest possible statement of your argument *for the most informed possible reader*.

There's one paragraph left. To write it, think about your "implications" section. What change (whether in theory or practice) in the mind of your reader does your paper suggest? How should the reader see or do things differently after granting the rightness of your conclusions? In a word, how have you changed the reader's world? Write the last paragraph as a description of this new world. The first and last paragraph, set side by side, should describe the "before" and "after" images of the world that your research pertains to—the world that needs your research. It is the world that needs to change in the way suggested by these images.

When I explained this to a group a faculty members recently something truly profound hit me: that last paragraph is also the first paragraph of your *next* paper.

## The Background

It is commonplace to begin a paper with a few commonplaces about the world in which we live. This world will be described in a way that emphasizes the social practices that the paper will offer a scientific analysis of. But the description will not itself be scientific. It will, of course, be "knowledgeable", but the knowledge it contains will not be dependent on either the theory or the method that supports the analysis. Instead, it will be based on sources that are also available to the reader, that is, published work.

That first paragraph of the paper, which establishes a common place where the reader and writer can share their interest in a particular corner of social life, will sometimes require elaboration. That is what the background section is for. The reader may know something about lean management practices, for example, but nothing about the manufacture and distribution of cardboard boxes in Sweden. Or the reader may need to know some general, historical facts about the company or companies that the writer has studied. The information that the writer provides here will, again, be available to the reader (and the writer does well to cite some good sources of the information along the way) but the writer is providing it anyway, for the reader's convenience. It is information that the reader *could be* aware of but *presumably isn't*. The reader will not find it presumptuous of the writer to assume that reader is ignorant about it.

The first paragraph of the introduction and the (roughly) five paragraphs of the background section really state the same claim. The introductory paragraph is the top of the iceberg (the summit), the background section is the tip of the iceberg (the part that is above water), and the dignity of the movement of this iceberg comes from everything the writer knows but does not say about the company, country or industry that the paper discusses. This is important to keep in mind as part of your growing base of knowledge. If you are doing a case study in a particular industry you will have a great deal of specialized knowledge about a particular company or, even more specifically, a particular team in a particular company. But you are not an expert on what that team does if you do not learn something about the factual world that it is embedded in. A scholar who has spent a lot of time within a company should be an interesting conversationalist about both that company and the industry it works in. That scholar should also be a reliable "go to person" for ordinary facts related to it.

The content of the background section can be easily distinguished from that of the "analysis" or "results" section by being "factual" but not, properly speaking, "empirical". The facts that are adduced here are not drawn from the *experience* of the writer, but from hearsay and reading (in the case of the former, preferably confirmed by the latter). By contrast, the facts that are presented in the results section are supported by "data" that have been gathered by way of a "method". This has an important consequence: the writer has a special authority to speak about his or her results. Because a valid method has been applied, we are entitled to trust the author's presentation. Moreover, we have no easy way of validating the results themselves. To do so, we'd have to requisition the writer's data. Even then, we'd have to trust that it wasn't just made up. So, to be really sure, we'd have repeat the study, i.e., carry out our own observations of the same phenomena, i.e., have the same experiences. It is only in extreme cases that we'd go through that kind of trouble.

But the background section can be "fact checked" in a more workaday sense. The reader can simply read the sources that the writer uses to support his or her claims. We can even find some better sources, where these are available. That is, because the background consists of commonplace claims, the writer has no special authority to make them. The text is also open to critique in a quite different way here than the results section is.

### The Theory

There is arguably nothing more standard in a social science paper than the theory section. Most journals will demand not only a theory but some "theoretical contribution". Like PhD dissertations, however, papers are sometimes written with a theory *separately* in mind. The author will announce two objectives: first, to make that contribution to theory development, normally by producing a "literature review", and, second, to make an empirical contribution, i.e., to present a set of results. While such papers do sometimes get published too, I don't recommend this approach. Your empirical results ought to have theoretical implications. You should not "develop" your theory independently or in advance of your results.

As Pierre Bourdieu said, a theory is a "program of perception". In your theory section you are telling the reader how you see the world. In the social sciences, this means announcing which of the available theories of a particular, say, social practice you are letting inform your vision of that practice. It is how you are construing (or even outright constructing) your object. It is very important that, as a program of perception, you let your theory assign

a series of descriptive tasks, marked by your concepts. Your theory tells you how you have to describe the world in your analysis (or "results" section). And for this reason it is important make an inventory of the concepts (the theoretical terms) you will use in your paper. While your analysis will *use* these concepts, your theory section will *account* for them.

A theory is built out of concepts. Concepts inform our vision; they make us see a set of facts, actions or events *as* something, i.e., as being of a particular kind. When we see something as, say, a "technology of self" or a "sensemaking process" or an "abstract machine" we have subsumed some part of the social world under a concept. Indeed, there will usually be more concrete objects and therefore more particular concepts: conduct, action, affects. This is why some people also call concepts "categories of observation". They simply make it possible to see particular things. More precisely, they make us construe the flux of experience *as* made up of empirical objects of a particular kinds. And objects in turn are simply limits on the possible. An object's "properties", i.e., the specific truths you can state about an object, are ultimately limits on the way they can be combined with other objects (defined by the theory).

These possibilities are precisely what you want to remind the reader of. They are the expectations that your results will artfully disappoint. A good research paper in the social sciences shows that the objects that constitute the social world are capable of being combined in ways we did not expect. And it was our theory that conditioned us to hold those expectations in the first place.

### The Methods

The methods section is perhaps the most conventional part of a standard research paper in the social sciences. There is a certain amount of room for creative syntheses of theory, and often a great deal of freedom in how the results themselves are presented. But the methods you use, and they way you talk about them, have to make your reader trust you. After all, you are setting the reader up for an artful disappointment. It is the methods section that ensures that the reader will actually *feel* that disappointment rather than merely doubt your results.

The most important thing, however, is that you tell the reader the truth about what you did. This section must accurately describe how you gathered the data on which you base your conclusions. How many interviews did you do? Over how long a period did you do onsite observation? How many surveys did you send out? How did you select your subjects and informants? How big was the data set you drew from industry databases?

But you must also show an awareness of the standard methodologies in your field. Though the distinction is sometimes blurred, method is what you did, methodology is the account of why it is the right thing to do. The standards here are often expressed by others in classic papers or handbooks. In the fields I normally work with, for example, if you're doing a case study you are likely to cite Kathleen Eisenhardt's paper "Building Theories from Case Study Research" (AMR, 1989) and, just as likely, Robert Yin's widely read handbook, *Case Study Research* (Sage, 2009). You do not have to agree with everything they say, and both approaches themselves have a history of reception, i.e., their views have been adapted and modified in particular studies that may look more like your own. The point is merely that in order to be taken seriously as a "case study", your paper must acknowledge precisely the tradition of case-study research that is informed and guided by such "standard" statements (accounts) of method.

The good news is that once you have shown that you are an intelligent reader of the methodologies that are available in your field, your reader is likely to trust even your intelligent breaches of those methods. In one sense, it is true what Paul Feyerabend said many years ago in *Against Method*: "The only principle that does not inhibit progress is: anything goes" (p. 7). But once you've done "whatever" it takes to reach your conclusions, you must make a compelling case for what you've done. And here it is always a good idea to gesture respectfully at the conventional wisdom that constitutes your field of research.

## The Results

In a standard-issue academic paper, the analysis of your results occupies about 37.5% of the text, or fifteen out of forty paragraphs. This morning I'd like to say something about how those fifteen paragraphs should be written. Keep in mind that this is the section in which you are the epistemic authority. Unlike the theory and method sections, your reader is not presumed to understand in advance what you are saying. Unlike the background section, your reader is not presumed to have access to sources that might contradict you. In the results section, you have the data. Your readers are, on the whole, going to trust you when you tell them what your data says.

I suggest you divide the section into two to four main parts, perhaps framed by an opening paragraph or two and a concluding summary. (It possible to do away with this frame by letting your methods section outline the structure of your results section and opening the implications section with a summary of the results.) The results section will be making an overall claim, which will have been summarized in the third paragraph of your introduction ("This paper shows that…"). But this claim will find support in a number of sub-claims. I suggest coming up with two to four claims mainly to give us that much-needed sense of finitude, each claim can then be given three to four paragraphs of support. In some cases, however, you'll have ten or more claims, each of which will have its own paragraph. The point of these proportions is not impose a set form on the presentation of your results but to get you to think about how the space of your results section will be structured.

The structure of the space will make it easier to organize your time. You goal, as always, is to get your prose into good enough shape to let you write a coherent paragraph (six sentences of prose that support a single, well-defined claim) in thirty minutes. Since you'll be writing fifteen paragraphs, you'll need seven and a half hours to do it, or three two-and-a-half-hour sessions. So when you are looking over your data, try to analyze it into claims that your are able to support in this way. This is the key to "prosing your world", to putting what you have *seen* into writing.

A good results section is an account of your *observations*. It's not an account of everything you've seen or heard in the field, nor an exhaustive presentation of your data set. It is a summary of the data in support of empirical claims, i.e., claims about "what is the case". Your background section also makes claims about what is the case, but let's call these claims merely "factual" and not "empirical". Like your empirical claims (your results), you must strive for the *truth*, but unlike them they are not based on observation. It is in your results section that your claims are based on your own first-hand research experience (that's what "empirical" means). This is where you come to *represent* the world of practice. It is where you make it available for theorizing.

## The Implications

These posts about the standard-issue social science paper (or standard issues in social science papers) are intended to get you thinking about the separate components of your knowledge base that put you in a position to write such a paper. You need not just "empirical results" but also an everyday "factual background", for example. It is not enough that you have seen something in your data; the result has to stand out in the everyday world of practice. I'll write about the theory section tomorrow to complete this series; this morning I want to talk about how important it is that your results have a specifiable set of "implications". Before you can expect to publish your paper, after all, you must be aware of how your reader's mind should be *changed* by reading. What does your work imply?

I have already said that the results section should *artfully disappoint* your reader's expectations. And the reader's expectations, in an academic paper, are shaped by the theory that the reader and writer share. But disappointments come in many shapes and sizes. Roughly speaking, a paper may highlight two kinds of implication, theoretical and practical. The results may evoke a disappointment in the theory, implying that the theory should be modified to accommodate the new facts brought to light by the study. Theories are always "underdetermined" by observation, i.e., our knowledge of particular facts, so we are getting used to having to modify our theories as new work brings new facts to light. You want to make sure that the implications of your study are precisely a *modification* of your theory, not its outright *rejection*. There are rare moments in the history of a discipline when a paper that argues for the replacement of one theory with another can be published (these moments are what Kuhn called "crises", which precede his "scientific revolutions"). But normally you want to imply only that a theory, which should remain in place (and dominant) after you have published your work, needs to rethink certain assumptions.

The other kind of implication is practical. Here we suggest that we are really disappointed in the world, the behavior of practitioners. If only they understood the theory (which is right about these matters), and lived up to our expectations of them, their practices would be much more successful. That is, we are making recommendations to practitioners in light of empirical study of their practices, framed by a particular theory. Bob Sutton's "no asshole rule" is such a recommendation. He is telling practitioners that "research shows" that hiring assholes (or keeping them around if you have accidentally already hired one) is a bad idea. He uses a variety of theories to show that assholes undermine your ability to bring about what he calls a "civilized" workplace. He is adjusting our view of who contributes to the value of an organization.

The implications section is best imagined as consisting of five paragraphs or one-eighth of the 40-paragraph paper. That will mean you are drawing between 3 and 5 specific implications (depending on how you write the section). And you are perfectly entitled to draw both kinds of implication: recommendations for practice and contributions to theory. Just keep it simple. The main critical standard here is a logical one: your recommendations have to "follow" from your background (description of practice), your theory, and your results in a logical way.