

PERSUASION WITH CASE STUDIES

NICOLAJ SIGGELKOW
University of Pennsylvania

The goal of every author is to write a paper that readers (and reviewers) find convincing. Since writers of papers based on case research do not have recourse to the canonical statement “results are significant at $p < 0.05$ ” that helps assuage readers’ skepticism of empirical papers, researchers using case research often feel they are fighting an uphill battle to persuade their readers. In this short essay, I provide some thoughts guided by my experience of reading, reviewing, and writing papers based on case-based research over the last decade. These are clearly only the views of this particular writer and thus should be taken with a considerable grain of salt. I am seeking here more to provoke thought than to provide answers.

What makes a case study persuasive? The first big obstacle that many writers feel they face is the charge of having too small a sample. Yet, imagine the following scenario, adapted from Ramachandran (1998): You cart a pig into my living room and tell me that it can talk. I say, “Oh really? Show me.” You snap with your fingers and the pig starts talking. I say, “Wow, you should write a paper about this.” You write up your case report and send it to a journal. What will the reviewers say? Will the reviewers respond with “Interesting, but that’s just one pig. Show me a few more and then I might believe you”? I think we would agree that that would be a silly response. A single case can be a very powerful example.

Perhaps not surprisingly, the management field is not alone in its debate about the value of small-versus large-sample research. In neurology, where a lot of knowledge has been gleaned from case studies of individual patients with particular brain injuries (lesions), a similar debate is underway. Ramachandran, a prominent neurologist, uses the example above to make his case for case research.

So should we now rejoice and simply cite Ramachandran to motivate and justify our case-based research? Well, we had better not forget that the above scenario involved *a talking pig*. That was quite a deal. Thus, my first main point is that if you

want to write a case study that derives its excitement and justification through little more than the description of a particular phenomenon, make sure you have a talking pig. If not, a purely descriptive study will be a hard sell.

The second charge that case-based researchers often feel obliged to defend themselves against is that of nonrepresentativeness. “You have a biased sample,” reviewers might say. Let us again have a quick look at the field of neurology. One of the most celebrated case studies in that field is of a man named Phineas Gage. Living in the second half of the 19th century, Gage was the foreman of a construction crew preparing the bed for a new railroad line. Part of his job was to fill holes, first with gunpowder and then with sand, which was then packed in with a large tamping iron. Unfortunately, at one hole Gage forgot the sand, created a spark with his tamping iron, and ignited the charge. The tamping iron, weighing thirteen and a half pounds, shot through his head, landing 30 yards behind him.

Remarkably, Gage survived and continued to live for 12 more years, despite the large hole in his head and major destruction to his brain’s frontal lobes. However, both psychologically and behaviorally, he was a changed man. For example, while he had previously been considered a smart man who energetically executed his plans, he now was capricious and vacillating, devising many plans but not following through with any of them. Similarly, whereas before he had been described as having a temperate personality, he was now impatient and profane, particularly when advice given to him conflicted with his desires. These psychological and behavioral changes led observers to draw inferences about what functions might be performed by the frontal lobes.

If one were to write a paper on this case, the charge, “You haven’t picked a representative person” would be absurd. Of course Gage was not randomly picked, and for good reasons. Likewise, the organization that one has chosen to study may very well not be selected randomly. In fact, it is often desirable to choose a particular organization precisely because it is very special in the sense of allowing one to gain certain insights that other organizations would not be able to provide. Few au-

This essay is an extensively extended version of remarks delivered at the 2006 Academy of Management meeting. Many thanks to Jan Rivkin and Sara Rynes for helpful suggestions.

thors defend their case choices, though, and some even try to claim that they have a “representative sample.” To me, that is a mismatch of method and goals: to say something representative, you need to pick a different methodology.

The neurological example also suggests a response to a third possible reviewer request: “You need to collect more data by studying more cases.” In the case of Gage, it is easy to reply, “There are not that many people out there with large holes in their frontal lobes.” If the organization one studies is a Phineas Gage, it is much easier to defend the research site. Of course, there is a price to pay. In studying a “special” organization, one needs to be careful with the kinds of conclusions that one draws. The specialness pays off, however, if it permits particular insights that allow one to draw inferences about more normal firms. Otherwise, the interest of the findings is much more limited. (In the end, we are more interested in people who *have* frontal lobes than in those who do not.)

Lastly, there often is an important difference between Phineas Gage and the organizations one studies. At least as a reasonable first cut, one can take the flying rod as an exogenous, random event that caused the hole in Gage’s frontal lobes and made him so special. In contrast, with many organizations, the features that make them so special are endogenous. In those cases, more care needs to be taken with respect to the conclusions one can draw.

So how does one increase the likelihood of publishing case research if one doesn’t have a talking pig or a Phineas Gage? In this situation, a paper cannot just stand on its descriptive feet, but also has to provide a conceptual insight. My rule of thumb is that the grander the theoretical claims, the more free-standing the theory has to be. In other words, even if a reader were only to read the conceptual part of the paper, he or she would be convinced of the internal logic of the conceptual argument. So what, then, is the use of cases if the theory could be free-standing? I believe that there are at least three important uses for case research: motivation, inspiration, and illustration.

First, cases are often a great way to motivate a research question. If one’s conceptual argument is about why A leads to B, a case can be a persuasive way of demonstrating why this is an important phenomenon. One can offer a purely theoretical motivation, but one that is grounded in a real-life situation is usually much more appealing. Likewise, although it is true that individual cases cannot prove a theory (“A *always* leads to B”), individual cases can sometimes suffice to falsify theories, as a single counterexample is enough. For instance, you might say, “Existing theory claims that A leads

to B, but here is a case where A did not lead to B,” pointing to the fact that the theory is not quite right, or at least that there might be something missing in the theory, motivating further research and justifying more refined conceptualization.

There are two caveats, though, to this game of poking holes in existing theories. First, theories and models are always simplifications. If they were as complex as reality, they would not be useful. Indeed, the value of theories is to cut through idiosyncracies and unearth similarities across cases. Thus, we almost always will be able to find instances in which a theory does not hold precisely. The burden of the author is then to convince the reader that this violation is really an important and insight-provoking violation of the theory. Second, the other challenge is to come up with a new conceptual framework that does not overdetermine the phenomenon. This is a particular pitfall of case research. As one immerses oneself in the intricate details of a particular case, many variables may appear to be crucial. Yet theory is only helpful if it can rise above the idiosyncratic case. Thus, one will have to make choices and simplifications in order to create useful theory.

The immersion in rich case data enables, however, the second main use of cases: as inspiration for new ideas. Indeed, the goal of inductive theory generation features quite prominently in many case-based research papers. If only limited theoretical knowledge exists concerning a particular phenomenon, an inductive research strategy that lets theory emerge from the data can be a valuable starting point.

Yet, as noted above, I believe that cases can also help sharpen existing theory by pointing to gaps and beginning to fill them. Thus, the near-ubiquitous claim that “not much is known, hence we engage in grounded theory building,” does not seem to me a necessary condition for the justification of case research. Moreover, such claims of existing ignorance at times do not ring true (and in the worst case can be taken by reviewers as describing the author’s, rather than the field’s, state of knowledge). It can also get writers tied up in knots about professing to have entered the field with no preconceptions. In my view, an open mind is good; an empty mind is not. It is true that one wants to retain the capacity to be surprised, but it seems useful (and inevitable) that our observations be guided and influenced by some initial hunches and frames of reference (see also Suddaby, 2006).

The third valuable use of cases in the context of making a conceptual contribution is to employ them as illustration. At first this may sound like a mundane use, but it really is not. Pure conceptual

arguments quite often have two shortcomings. First, a proposed theory may posit that construct A leads to outcome B, but since A is a “construct,” the reader often wonders what A is in real life. How would one measure A? How would one know that the empirical variable that one has obtained really captures A? By seeing a concrete example of every construct that is employed in a conceptual argument, the reader has a much easier time imagining how the conceptual argument might actually be applied to one or more empirical settings.

The second shortcoming of purely conceptual arguments is that the underlying mechanisms are often completely speculative: A leads to B because forces XYZ operate. It is then up to the reader to decide whether the proposed causal model and its factors X, Y, and Z are plausible. If you can show with an example that X, Y, and Z actually operate and create the relationship between A and B, this is a quite powerful use of a case. In fact, getting closer to constructs and being able to illustrate causal relationships more directly are among the key advantages of case research vis-à-vis large-sample empirical work. In large-sample work, the distance between conceptual constructs and measurable variables is often rather large. (For instance, the number of constructs for which the R&D-to-sales ratio has served as a proxy must be in the hundreds.) If the econometrics is done convincingly, the reader may well believe that an empirical relationship has been documented between proxies E and F. But whether this relationship really reflects the underlying conceptual arguments concerning constructs A and B is often quite unclear, not to mention whether it is really forces XYZ that cause the correlation between E and F.

The ability to get closer to theoretical constructs is particularly important in the context of longitudinal research that tries to unravel the underlying dynamics of phenomena that play out over time. As scholars have increasingly begun to appreciate the role of dynamic processes (e.g., path dependency or evolutionary processes), rich longitudinal research is needed to provide the details of how these processes actually play out.

In terms of final paper structure, the distinction between using cases for inspiration versus illustration is mainly one of sequence. For example, an inductive paper employing a case as inspiration might start with the case and then focus on theory. In contrast, a paper using a case as illustration might more usefully present the case after the theory (which in turn may be preceded by a motivating case example). Although for expository purposes these differences in presentation may be useful, the

research itself tends to be much more iterative, going back and forth between data and theory.

To illustrate the different uses of cases, let me briefly touch on two of my own papers. In my work involving the study of Liz Claiborne (Siggelkow, 2001), my research was driven by the conceptual question, “What is the relationship between inertia and tight internal fit among a firm’s activities when the firm faces external shocks to its environment?” Liz Claiborne was only one of a number of companies that I studied at the time. The framework proposed in the paper emerged more from a conceptual exercise than from my exposure to Liz Claiborne’s experiences. However, the case turned out to be a very helpful illustration and was used in that manner after the conceptual framework was presented. Clearly my research on Liz Claiborne had an influence on my thinking, but it was not the primary inspiration for the eventual framework.

My work involving Vanguard (Siggelkow, 2002) was quite different. It again started out with a conceptual question: “How do firms evolve toward tight internal fit among their activities?” The first use of Vanguard was to motivate the research. The current system of Vanguard’s choices concerning its activities, resources, and other organizational elements was extremely complex. It was fairly persuasive to argue that it was unlikely that such a complex system had sprung into being in one fell swoop. (Interestingly, although I have always used this set-up for verbal presentations of this work, this “motivational” use of the case did not survive the review and rewriting process. The paper’s eventual motivation became purely theoretical.)

The existing literature did not offer much guidance on how to describe the evolution of such systems in a systematic manner. Hence, I felt the best way to gain some traction was to let the case speak, albeit through my chosen lens of thinking about firms as systems of interdependent choices. Consequently, the Vanguard case became the primary inspiration for the ensuing constructs that allowed me to more systematically describe the evolution of firms as systems of interdependent choices. In the discussion part of the paper, I applied the constructs to show how one can compactly describe a number of possible developmental paths and employed Vanguard’s history as an illustration. Thus, the Vanguard case was used in a number of different ways in this paper.

To summarize, regardless of how cases are eventually used, research involving case data can usually get much closer to theoretical constructs and provide a much more persuasive argument about

causal forces than broad empirical research can. One should use this advantage. However, one will not be able to say, "You should believe my theory that A leads to B, because I show you an example here." That is asking too much of a single case study, or even of a few cases. The theory should stand on its own feet. One needs to convince the reader that the conceptual argument is plausible and use the case as additional (but not sole) justification for one's argument.

Focusing on the conceptual argument has a further beneficial side effect: it provides guidance to the writer about what details to exclude. A common weakness of case-based research papers is lack of selectivity and presentation of only those details that relate to the conceptual arguments. This is a charge easier levied than remedied. Once one gets immersed in a case and has spent considerable time reading about or observing an organization, it can easily feel that *everything* is "so interesting" and, as a result, should be shared with the reader. The unfortunate truth is that in most cases readers will not be as interested in the particular case as you are (unless you have a talking pig). Rather, readers will be much more interested in the conceptual argument, because it is this argument that can shape their future thinking and allow them to see the world in a slightly different light. Having a salient example of this new insight (i.e., your case) is a bonus, but it is not where most readers will see the biggest added value from your paper.

Thus, once one has crystallized one's conceptual argument, it is helpful to go back to the case and ask which details really help tell the story and illustrate the mechanisms behind one's arguments. In particular, the persuasiveness of the arguments is greatly strengthened if serious attention is given to alternative explanations—and why these alternatives are unlikely to hold. It is hard to overdo this part of a paper. The more robustness checks one can offer, the more convinced readers will become of the newly proposed mechanisms. To focus on the conceptual contribution can also have an implication for the structure of the paper. For instance, a chronological report of the data may seem most natural, yet it may actually not be the most helpful way to support one's conceptual points. Grouping the data differently may be more effective.

If a case paper thus stands and falls with its conceptual contribution, an even thornier question arises: "What is a valuable contribution?" I will not attempt to provide an answer, but I will point out a particular challenge faced by case-based papers: the problem of "ex post obviousness." At the end of

a paper, a reader will ask, "How surprised was I by this finding?" If the answer is "not at all," the author usually has a problem.

Granted, this is a problem also found with large-sample empirical papers. Yet these papers can at least claim that they have confirmed that A leads to B in a large sample—and with $p < 0.05$ certainty! Case-based research does not have recourse to this fallback. Case-based evidence is more at the level of an existence proof: Here is one example of how A leads to B. If the reader can reply, "I'm not really that surprised that you can find in the world at least one example of A leading to B," the value of the contribution of the paper can be in doubt. And by the way, dressing up the findings in theoretical high-brow language, although commonly attempted (why report a talking pig if you could report "an isomorphism between auditory signals emitted by members of the species *sus domesticus* and *homo sapiens*") is not a way out of ex post obviousness.

I would like to conclude with a final thought on the issue of contribution that may be controversial—but at least I will have succeeded in being provocative. In my view, it is much harder to make a paper interesting whose findings or conclusions only address theory. A paper should allow a reader to see the world, and not just the literature, in a new way. An acid test would be the following: Imagine that someone who is interested and knowledgeable about the phenomenon you study, but who does not know the literature, were to read your paper. Would this reader find your paper and its results interesting? Or are the paper's contributions only of interest to those who can appreciate the references and refinements to prior theory? If theory talks only to theory, the collective research exercise runs the danger of becoming entirely self-referential and out-of-touch with reality, of coming to be considered irrelevant.

This warning applies, of course, to all types of research, not just the case-based. It is striking, though, to see that this charge can sometimes be levied against case-based papers, which with their direct and intense exposure to the world would seem to be more immune to this problem than most other forms of research. However, some writers seem to strain themselves to construct arcane conceptual arguments to justify their research, in the process losing sight of the truly interesting empirical observations they have made. I am all for theory development; yet theory as a purely self-referential exercise—rather than as an attempt to better understand the

world—strikes me in the end as a poor allocation of time and effort.

REFERENCES

- Ramachandran, V. S. 1998. *Phantoms in the brain*. New York: Harper Collins.
- Siggelkow, N. 2001. Change in the presence of fit: The rise, the fall, and the renaissance of Liz Claiborne. *Academy of Management Journal*, 44: 838–857.
- Siggelkow, N. 2002. Evolution toward fit. *Administrative Science Quarterly*, 47: 125–159.
- Suddaby, R. 2006. What grounded theory is not. *Academy of Management Journal*, 49: 643–642.

Nicolaj Siggelkow (*siggelkowwharton.upenn.edu*) is an associate professor of management at the Wharton School, University of Pennsylvania. His research focuses on the strategic and organizational implications of interactions among a firm's choices of activities and resources. His work analyzes organizational design, the evolution of firms' systems of choices, and firms' adaptations to environmental changes. He received his Ph.D. in business economics from Harvard University.
