

Writing Economic Theory Papers

by **Simon Board and Moritz Meyer-ter-Vehn**

October 2014

These notes aim to distill our advice concerning how to write a decent theory paper. These notes are deliberately brief, making them easy to consult. In particular, it consists of the advice we most frequently give to our students.

First, we will follow our own advice and explain our marginal contribution over the other works out there. Varian (1997) is a wonderful essay, expressing his personal approach. In comparison, this paper seeks to be more of a “how to” guide. Thomson (2001) is a comprehensive and well written guide to all aspects of a young economist’s career. This paper is more focused on explaining what we look for in good economic theory paper.

Contribution of the Paper

When writing a theory paper, one must prepare to answer the most basic question: Why should I care? The answer usually falls into the following categories:

1. The paper asks a new question. A paper should answer a question. Occasionally the contribution of the paper will be to ask a novel question rather than writing down the model or doing the analysis. These are often the most cited papers - once you have a new hammer, lots of things start to look like nails. For example, Kamenica and Gentzkow (2011) ask what is the optimal way for a principal to design an information structure for an agent. The contribution is to ask not just whether a particular signal can make the principal better off, but what is the best the principal could do if they could choose any experiment.
2. The paper posits a new model. A model is a lens through which we see the world, and there are different lenses that allow us to focus on different aspects of the problem. Most models contain something new, but some are a sufficiently large conceptual step that it promises to bring new headway (even if the results of the paper are not that exciting). For example, Sannikov (2008) paper on the continuous time principal agent problem showed how we can model these classic problems in a different way that has the potential to be much more tractable.
3. The model speaks to an important application. For example, applying matching theory to school

choice is important because of the economic significance of the application; without this the comparison between Boston and Gale-Shapley algorithms would be a dry theoretical discussion.

4. The model identifies some interesting economic force. The model may be unrealistic in many dimensions, but it isolates a new feature of a family of problems that changes the way we think. Much of the value from economic theory is in honing our intuitions and understanding the qualitative nature of economic forces. For example, once one understands Spence's model of signalling, one can happily talk about bankers building large granite buildings; we understand that we need only check that single-crossing is satisfied.
5. The model develops new empirical predictions. The classic Friedman view is that a theory should be evaluated on the basis of its predictions. For example, a reasonable model of peer effects would posit that schools sort children of different ability assortatively into different classes, placing smarter kids into larger classes; this may then explain why regressions of class size on child performance have historically found relatively weak correlation.

One natural question is why we should care about "economic intuition" if we only care about predictions. One response is that a major part of economic theory is about categorization (e.g. defining a signaling model). A second comes from the Lucas critique which point out that if the foundational assumptions of a model are wrong, then when the environment changes the model may break down.

6. The paper makes a technical contribution. There are some long-standing problems in economics; if you can make progress on one by inventing a new proof method or just staring at the problem longer than others, then others will be naturally interested in the paper. For example, Sugaya (2014) proved the folk theorem in games with imperfect private monitoring.

Exercise: What is the contribution of the following papers:

- a. Maskin and Riley (1984) added risk averse buyers to the standard risk-neutral auction setup.
- b. Mailath and Samuelson (2001) consider a model of reputation where high-cost firms always produce low quality and low-cost firms are strategic.
- c. Crawford and Sobel (1982) consider a model of communication where a "sender" obtains a signal about the state of the world, sends a message to the "receiver" who takes an action. The two agents' ideal actions increase in the state, but the sender always want a higher action than the receiver.

While the main contribution of the paper could be any one of the six criteria above, these different objectives overlap. For example, if one is writing a paper that speaks to an interesting applied question,

the model should be reasonable and the economic forces persuasive within the context of that application. On the other hand, a more abstract paper where the contribution is the model need not look as “realistic” but should promise to deliver some new insights. There are also different tastes within the profession. Some researchers take models very seriously, thinking there are “correct” models and that our job is to work out the implications of these models; other researchers think that a model is a way of delivering an insight, so what matters is what we learn rather than the assumptions per se. A good introduction will make it clear to everyone what you see as the contribution.

While a paper can be great if it satisfies any of the above criteria, one should remember that we are economists and not mathematicians. This means that the ultimate aim is to speak about applications; mathematical tools are important, but they are a means to an end. Moreover, from a practical standpoint, there is a far bigger market for applied theory than high-brow theory. Nash only needed to define his equilibrium once, but it can be applied a thousand times.

The Introduction

A canonical Introduction has the following steps. Many good introductions differ from these steps, but this is a good benchmark.

1. Broad motivation (1-2 paras). Why is this area interesting? The paper can be motivated by an application (“Why does inequality rise after recessions?”) or via the literature (“Coase’s conjecture states that a monopolist who cannot commit cannot price above marginal cost”).
2. The contribution (1-2 paras). See above. It’s a good idea to have the line “The contribution of this paper is...” in order to force yourself to be explicit. For example, if your model and results are novel, then you might want a paragraph on both. Remember: people don’t know your model at this point, so everything should be self-explanatory.
3. Explain the model (1-2 paras). If readers are going to understand your results, you need to explain the setup. This does not need to be formal, but it needs to sketch out the main elements that drive your results.
4. Explain your results (1-2 pages). The aim is to give an overview of your main insights. Explain what you are doing, the driving forces behind your results and some important implications. This should inform readers who will only read the introduction (i.e. the majority of readers) and whet the appetite of those who are interested in this area.
5. Explain the relationship between your results and related work (1-2 pages). The aim is not to survey the literature, but to elaborate on your contribution. Never criticize others work explicitly:

it is unnecessary and those authors may be your referees. Instead of saying “Arrow and Debreu (1954) ignored market power” say “we develop Arrow and Debreu’s canonical model to include market power”.

Some rules:

- Don’t make typos. When you are young, the readers are updating very quickly about your ability. Typos raise question marks, and can cause frustration if it clouds clarity. And a frustrated reader will likely reject your paper.
- Don’t make loose statements. Conjectures should be clearly labelled as such, and only included with caution. If a reader sees you make an incorrect statement they start to question everything else.
- Don’t include long-winded justifications of the model’s assumptions in the introduction. By all means include a sentence motivating an assumption that is non-standard, but you should not be overly defensive.
- Don’t spend too long on the extensions in the introduction. The reader barely understands the model at this point, let alone an extension. By all means give a headline result (“This shows our main results are robust to X ”) but don’t expect the reader to follow any details.
- Don’t be too repetitive. For example, if you cite an empirical study, then you don’t need it in the abstract, paragraph 3 and in the literature review.
- Be consistent. If you explain things using an intuition, then it is best to stick with this approach throughout the paper rather than switching every other result.

Model Section

The heart of a theory paper are its model and its theorems. The model section distills how you want to the reader to think about the economic problem at hand; the theorems are the most important and interesting insights you have derived in your model. If it takes you four pages to state your model, remember that a model should simplify the way we think about a problem. If you have ten theorems, you should think which three of them you would want the reader to take away if you had to.

A useful rule is: “one paper, one model”. The purpose of the model is to fix ideas, so changing it over the course of the paper is usually a source of confusion. When deriving your results you may find that some desired results hold in one variation of your model, while other desired results hold in other variations. Sometimes it is possible to capture those different variations by a parameter in the model and view the results as comparative statics in that parameter. Sometimes there is one canonical baseline model that can be extended in later sections. It is usually not a good idea to specify an overly general

model in which one can prove little of interest, and then introduce the crucial assumptions just before the theorem statements.

The standard way to present a game-theoretic model (and much economic theory these days is based on game theory) is to state as succinctly as possible:

1. Who are the players,
2. What can they do,
3. What do they know,
4. What are their preferences.

Delaying these definitions, changing this order, or interspersing them with too much interpretation and preliminary analysis is usually a source of confusion. Indeed, even in the Introduction, one should ensure they things are clear so the reader can understand your results.

Remarks. While laying out the model, you probably have some things you wish to say concerning the motivation and interpretation of some assumptions, and the necessity of others.

- Motivation of assumptions. If a motivation is short (1 line), then you can include it in the main model section. For longer justifications, you can include a “discussion” section at the end of the model section.
- Varying assumptions. Be careful when talking about other assumptions you could have made; this can cause confusion about the model you are analyzing. If an assumption is for convenience or definiteness, then you can say so in the model section (e.g. in a footnote) if it does not interrupt the flow, or in a “discussion” section at the end of the model section. Longer discussions of the role of assumptions can be delayed until a “discussion” section after the main result, an “extensions” section of the conclusion. In general, you should only mention something in the model section if it is likely concerns the reader. Otherwise, delay it until after your main result.

The Analysis

Writing this section is harder than it may seem. The analysis has to be crystal clear and also keep the reader’s attention. This means that you should ask what the reader want to learn from the paper, and present this information in the most precise way possible. Like any argument, you should first tell the reader where you are going, and tell them why each step is important or interesting. If your paper reads like page-after-page of mathematical manipulations you will quickly lose their attention.

When presenting a result, you should go through the following steps:

1. Before the result. Above the statement, give any relevant definitions, remind the reader of

anything they need to know and say what the result will say in words. If the result is self-explanatory, this may be one line.

2. State the Theorem. The statements should aspire to be self-contained and not require the reader to go searching for variable definitions. There is a balancing act between mathematical rigor (algebraic statements) and clarity (prosaic statements). In a presentation, you should tip the balance towards clarity; in the paper, the result should be fully rigorous. In practice, Lemmas can be written in a more mathematical way, while you should try to Theorems more comprehensible by make appropriate definitions before the Theorem. It is usually a bad idea to include a definition in a theorem statement because it tends to clutter the statement and reduce clarity. Also, try to avoid long equations in Theorem statements.
3. State the proof. If this is interesting or short, then include it in the main text. If it is standard and boring, then put it in the Appendix. If the analysis is key for the paper (e.g. your main derivation), then consider putting a heuristic version in the main text before the statement of the theorem. People often skip the proofs so, by putting something in the main text, you nudge your readers to read it.
4. State the intuition. Every result should be explained in simple English unless it is obvious or technical. The main takeaway from the paper is often the economic insight, so it is important to deliver a good explanation.
5. State any implications. If the result will be used later in the paper, or has empirical implications then say so. Of course, you should be wary about discussing tangential issues in the middle of the flow of an argument.

Other issues:

- Footnotes. Footnotes are for elaborations and clarifications. Footnotes are optional for a reader, so don't include necessary definitions. As a rule, a paper should have fewer footnotes than pages.

Bibliography

H. Varian (1997), "How to Build an Economic Model in Your Spare Time" in *Passion and Craft: Economists at Work*, edited by Michael Szenberg, University of Michigan Press.

W. Thomson (2001), "A Guide for the Young Economist", MIT Press