Advice from an easily confused biologist

Thomas Lesaffre

September 2025

Every study in evolutionary biology starts with one or more scientists asking a complicated question about how living organisms have been and are being evolved. These scientists then look for an answer using sophisticated tools that are often specific to their subfield, study system or even lab, and what they find is often not trivial, requiring a lot of thinking to make sense of. Once they think they have cracked it, they are faced with the difficult task on weaving a meaningful story that shines light on their initial question, in a way that can be understood by readers and listeners that will be easily confused and so put off by poor story-telling. This is especially true for theoretical studies, because most biologists have a conflicted history with mathematics and are thus particularly prone to giving up when the story is difficult to follow. In this document, I have compiled a handful of advice on telling a good theoretical story that I often find myself giving to fellow early career researchers, from one weaver to another. I have organised this document along three steps in the life of a theoretical study, namely (1) building a model, (2) writing about it and (3) giving a good theory talk. I hope this will be of some help.

1 Building a model

Start weaving early. The first key to ultimately telling a good story is to start thinking about it as soon as you start working on a project. The earlier you start thinking about the story you are going to tell the better it will be in the end, because you will have given more time to your ideas to mature.

Have a clear *biological* question. We are theoretical biologists, not mathematicians. Our goal is not to produce fancy new mathematical proofs or to apply elegant technics. Although it might be pleasant to be able to do so, this should never be the motivation to start a project. Our job is to think conceptually about evolution, and mathematical modelling is merely a tool we use for this purpose. So as you start your project, you should be able to explain what is the biological question you will be addressing and why it is interesting, based on the current state of knowledge in evolutionary biology. I encourage you to keep pondering over this as a form of fruitful doubt: as you make your way through your project, your outlook on your work will be modified by your findings and interactions with fellow scientists, which will enrich your thinking and make your story better in the end.

Start simple and extend based on what you find. Think of the simplest biological scenario that captures what you are looking to investigate and start there. See if you can recapitulate known results, and later extend from this simple model based on your findings, always keeping in mind the story you are ultimately going to tell. For each extension you pursue, you should again be able to explain why it is interesting.

At all times, have clear biological scenario in mind. Being perfectly clear about the biology that you are modelling is absolutely essential. You *must* be able to answer these questions at all times:

(i) What is the demography, or in other words, how is population size regulated?

(ii) What is the life cycle, that is, what is the sequence of events and actions individuals go through over a time step?

(iii) What are the evolving traits, and what is their genetic basis?

One way to make sure you are clear enough about your model is to think of how you would implement it in an individual-based simulation, where the actions and genetics of each individual are modelled explicitly. If you cannot think of how to do this, then you must go back to the drawing board and clarify things until you can.

Pay attention to notations. Finally, a major aspect of clarity in modelling is to define good notations. There is many spoken and unspoken rules, or conventions that may vary between subfields, but there is general ones that we should all follow. Useful guidelines are given in the paper by Edwards and Auger-Méthé (2019) provided on the summer school website.

2 Writing about your model

A theoretical paper is always composed of two parts, a main text that tells the story and an appendix that supports it.

2.1 A good paper starts with a good appendix

Ultimately, the goal of an appendix section is to allow future readers (including future you!) to re-derive all the results that you present in the main text, and to provide support for every claim made about your model in the paper. But more importantly, it is a precious tool for you to use as you progress through your project. Therefore, the appendix is the first section that you should write.

Start writing your appendix early. Start writing as soon as you have a clear biological question and a baseline model to tackle it. The first section that you write should be a detailed description of your model. This section might ultimately be moved to the main text, but writing it early will force you to think clearly about the biological scenario you are modelling, again answering the questions from above: What is the demography? The life cycle? What are the evolving traits and their genetic bases? How would I implement all of this in an individual-based simulation?

Write up results as they come. Once you start analysing your model, write up the results you obtain in the appendix every time you gain a new biological insight, however small it may be. This will help you spot mistakes and potential gaps in your reasoning. Equations should be explained term-by-term to help the reader through them. It is often a bit of a pain to do so, but it forces you to be precise and everybody (including future you and your future students) will thank you for it. Ideally, you should accompany these explanations with a clean *Mathematica* notebook (or something equivalent) that checks that your calculations are correct. For instance, when you compute a selection gradient and re-arrange it to present terms in a meaningful way, it is good to have *Mathematica* perform a comparison between what you get from the symbolic calculation, and the expression you write in the appendix,

```
Sg = D[W, y] /. y -> x; (* The selection gradient *)
Sg == [Whatever your expression is] // Simplify
```

which should return ‘True’.

Make your paper as self-contained as possible. Readers should not need to go through several older papers to understand the scenario you are modelling, so instead of writing statements like "Mate choice is modelled as in Bob (1999)", which are really annoying to readers because they then have to go read what Bob was up to in 1999, explain how mate choice occurs in your model and politely cite Bob at the end.

Explain your methods. When you start describing your analyses, give a quick explanation of what your method is going to be and direct the reader to key references for details. Then, when you compute quantities that will be used for biological interpretation, explain what they represent and how to interpret them before doing so. For instance, with the selection gradient, explain that it gives the direction and strength of directional selection acting on a trait: a positive (resp. negative) value indicates that selection favours an increase (resp. decrease) in the trait from the current state of the population.

Have a separate section for your simulations. If you are running individual-based simulations, you should have a separate section that describes the program in addition to the biological model description you gave at the beginning. In that section, you should write in words exactly what the program does (i.e., a pseudo-code that explain in what distributions values are drawn etc...). Further, make sure to annotate your code thoroughly as you write it (you will thank yourself for this when you come back from your long weekend at the beach). Once the paper is ready, make the program available online and give it a perennial identifier (DOI) using tools like Zenodo (which has the advantage of being integrated with GitHub). Also, it is worth taking the time to learn an efficient programming language. Most biologists only get taught R, which is suitable for statistics but absolute garbage for simulations. If you are a biologist that has only ever used R, learning SLiM is a smart time investment, as a lot of the syntax mimics R (and there is great workshops organised by the creators of the language all over the world, you can ask Vitor about it!).

2.2 The main text

If you have followed the advice above, then by the time you have enough results for a paper you should also have a fairly streamlined version of the story you are going to tell in mind. At that point, you are ready to write the main text.

Find the simplest route to introducing and answering your biological question. As a rule, you should try to make the main text as concise as possible. This should be relatively easy if you have been thinking about this from the start of your project. I find that the best theory papers are those that take you seamlessly from one section to the next, so that the main text can be read from start to finish easily – like a story.

Again, describe the model biologically. I cannot repeat this enough, you must be completely clear about the biological scenario you are going to consider. Demography, life cycle, traits and genetics must be covered with the same level of detail as in your appendix section. You may end up simply migrating the detailed description you gave in your appendix into the main text, but there is no harm in repeating yourself. Thankfully, appendices do not have a word limit.

Only present equations that are necessary. When you do present an equation in the main text, you should describe each term and explain their biological meaning.

Call variables and parameters by their biological meaning. In the text, always refer to variables and parameters by what they are biologically, instead of their symbols. For instance, do not write "Inbreeding depression increases as h decreases", but "(...) as deleterious mutations become more recessive (h decreases)".

Refer to the appendix. Every time you make a claim about the model or show an equation, make a reference to the section and equation(s) of the appendix that support the point or derive the result you are showing.

3 Giving a good theory talk

Once you have a story ready, you will likely have the opportunity to present your work at a conference, in a seminar or at a group meeting. Below, you will find important advice about how to give a successful theory talk, i.e., one where most of the audience and not just one or two theoreticians will listen to you until the end. These guidelines should also be useful to you for preparing the talk your group will be giving on Saturday.

A talk is a teaching exercise. The most important advice I can give is about the mindset you should have when designing and giving a talk. Many PhD students freshly out of their Masters still feel like they are a student being judged by professors who know the truth about what they are presenting, but you are not a student any more! You are an academic. You know your topic better than most people in the room – apart maybe from that one grumpy old professor at the back, but you are not preparing your talk for them! –, you are the only one who knows the original work that your are presenting and you are going to teach your colleagues about it. Really think of your talk as a teaching exercise.

One or two take-homes maximum. If you want your temporary students to remember anything about your talk, you need to focus on a small number of results to have time to go through them properly. Many students and postdocs try to jam their entire life's work in a 10 minutes talk, resulting in catastrophic failure. You are not here to prove your worth – you made it this far, you are a worthy scientist! – you are here to teach the audience something interesting, so a simple and well-supported message is key.

Work backwards from your final message. Once you have identified the message you want people to remember, work backwards from your conclusion slide to figure out exactly how much information you need to make your point.

Describe the model! I will repeat it again, you must give a complete description of the biological scenario you are considering. Giving a complete description of the model is the only way people can gain an intuitive understanding of what is happening in your model. If they are not able to do so, they will not trust your conclusions. "*Trust me bro*" is not a valid teaching method. You should also introduce relevant parameters and variables as you describe the model, but only those that will be useful to make your point. This may mean showing only a truncated version of your full model, which is fine and may lead to easy questions for you.

Avoid equations. Try to avoid showing equations at all cost, otherwise you risk losing most of the audience because biologists are often unfortunately allergic to mathematics. If you absolutely need an equation to make your point, it has to be – or at least look – simple, and you must walk people through it. Explain each term, and use it to give an intuitive understanding of what you are trying to convey. Use as few symbols as possible, and refer to quantities by their biological meaning.

Explain your plots and stay close to the biology. When showing a plot, explain what is on the axes and add information gradually. Always relate the results to their biological interpretation, never losing sight of your teaching goal.

Keep it short! No one ever got annoyed if a talk was nine minutes instead of ten, but the opposite is really not true.

References

A. Edwards and M. Auger-Méthé. Some guidance on using mathematical notation in ecology. *Methods in Ecology and Evolution*, 10(1):92–99, 2019.