© 2014. Published by The Company of Biologists Ltd | Journal of Cell Science (2014) 127, 2379–2381 doi:10.1242/jcs.155606

STICKY WICKET

Replicant III – Evolution

Mole

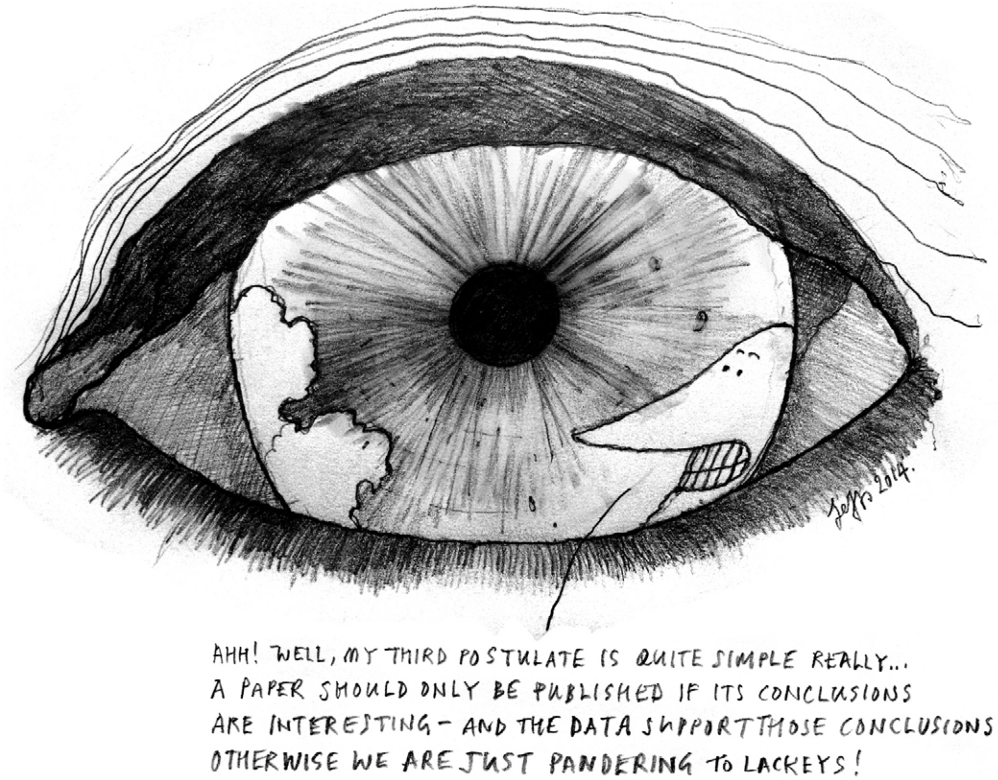

I'd love to embrace you, but first, I have to satisfy my sense of moral outrage.

Hey there. I'm watching Who Framed Roger Rabbit, and it just cracks me up. For those of you just joining us, I *was* watching Blade Runner, but it was just so depressing. And dark. *Noir* dark. And it was getting me down, with all of its weepy *replicants*. Or maybe it was the conversation we were having. About experimental findings that can't be replicated ('replicants,' you know?). And the various proposals to do something about this situation, many of which I suggest will bring our endeavor to a grinding halt. This is despite the stated objective of speeding up the process, because, it seems, we just don't have the time or money to let the scientific community do it ourselves. Maybe we should just give up. I was quit when I come in here, Bryant, I'm twice as quit now. (No, wait, that was Blade Runner, and I'm watching Roger Rabbit. Way funnier.)

One problem I see in all of this is the distinction between replication and reproduction of results. No, they aren't the same thing. If we replicate, we exactly copy the experiment. When we reproduce results, we vary the experiment, apply it to our own goals – we let it evolve. So what are we really talking about here?

One way we move findings forward is to apply them, say in drug discovery. If the experiment can be *replicated* and made into an assay, it may prove useful for those purposes. But often it doesn't work that way, hence, some of the cries of "replicant!" Indeed, it was the claim made by researchers in industry that major findings could not be replicated that propelled many of the ongoing discussions of the problem.

But it is often the case that an experimental finding is not very easily replicated. This doesn't necessarily make it *wrong*, it just makes it tricky (and variable). Biology is complex (which is why we study it so hard), and it often gives up its secrets sparingly. Experiments that can be reproduced only variably can still be useful. And it is often the case that if the results are variable (and actually correct), we find ways to make the system behave more reliably. There is an apocryphal story about this in chemistry (I don't know if this is true – it was something I read long ago and probably forget the details). The story goes that when chemists were trying to figure out the elemental composition of compounds, the values that were obtained were variable and fractional (a bit of this and a bit of that). Until Dalton showed that the ratios were fixed, and subsequently, that the reported values

Correspondence for Mole and his friends can be sent to mole@biologists.com, and may be published in forthcoming issues.

An occasional column in which Mole and other characters share their views on various aspects of life-science research.

Journal of Cell Science (2014) 127, 2379–2381 doi:10.1242/jcs.155606

conformed to such ratios. Our understanding evolved, and in the process, our analysis improved.

But what about when the findings are variable *and* just plain wrong? When experiments show us that what has been published is wasting our time? The growing consensus is that far too many such results are being published and this time-wasting aspect is too costly, in time and money, to continue like this. Last time we talked about some of the ways that have been proposed to try to fix this situation, some of which seem pretty reasonable and some that don't. So I promised that I would suggest *other* ways to fix things, ways that have not been suggested so far. "Finally!", you say. "Are you telling us that you had some ideas about this and you could have told us any time?" "No," I say (quoting RR), "not until it was funny."

So here, at long last, is my proposal for how we can start to improve the increasing problem of *replicants* (things that can't be replicated, don't forget!), in the literature.

1. Change the literature. The biggest problem we have, and the biggest reason that things get published that subsequently turn out to be wrong, is based on the way we go about publishing our findings. We work on a problem, often for years, and through hypothesis, experimentation, more hypothesis, more experimentation, we eventually reach something we think is an interesting conclusion. We write up our work, presenting the case for reaching our conclusion, and submit it. And then everything changes. Instead of considering whether or not our conclusion is justified, we receive a 'deal with the devil' that says, do these experiments and if you get these results, we'll publish your work. And if you don't, go try somewhere else. So instead of years of evaluation and re-evaluation, it all comes down to making a list of experiments 'work.' I am *not* saying that researchers then fabricate the findings. (I insist that this is extremely rare – any of us who are serious about science would not do this, and if someone does, they should be booted out pronto. But this is not what we're talking about.) But we are backed into a corner – experiments that give variable results will give some that *might* be correct, or might not, but we are forced to choose the results that were required by the process, and hope hope hope that we chose the right ones. This is not the way to do science. But it is the way the system has evolved. We will only fix the problem when the journals explicitly ask reviewers if (a) the conclusions are interesting (if not, then don't tell the researchers what to do to make it 'more' interesting) and (b) if the data support those conclusions. If not, explain why, but do *not* say what result must be generated by an experiment that was not done – leave that to the researchers to better make their case. Yes, reviewers can suggest experiments, but until we treat these as optional, we cannot break this chain of demand and response that has gotten us into this situation of *reacting* to reviews. If the reviewers are forced to say whether or not they are convinced of an interesting conclusion by the work that has been presented, we are far more likely to read about results that have been obtained over a suitably long period of time, and that we can more likely use in our own work. "But Mole!" you say (I can still hear you Leon!), "this assumes that the work as originally submitted is reproducible!" And you're right to be concerned, which brings me to…

2. Change the way we work in our institutions, and in our labs. Because of the gauntlet we run in publishing work, this reaction to reviewers (how often have we heard of colleagues who submit work to 'see what the reviewers want,' even holding back valuable results, knowing that they have to give them 'something.') we encourage a reactive behavior in our trainees. In many (most?) institutions students are told that they have to publish a certain number of papers to graduate, and therefore we advise them to 'get this result' for an experiment they are told to do by their advisors and/or committees. How often do we read, in drafts or even published papers, that an experiment was performed to confirm, or prove, an idea. But experiments are performed to *test* ideas. If we start to support this simple concept, that science is about testable ideas, we may begin to infuse our trainees with the currently crazy notion that *whatever a well controlled, reproducible experiment tells us, is a result that is useful*, we may well be on our way to helping to build a better system. Long, long ago, when I was a mere Molet, I knew a fellow student who received her Ph.D. by carefully showing that a published result could not be reproduced, despite her arduous efforts, and she performed very well-designed experiments that showed that the original conclusions were wrong. She was equipped to advance to what resulted in a long career. "But wait a minute, Mole!" (I can still hear you, L) "trainees need to publish for their career advancement. Are you saying that we tell them not to publish?" Of course not. But we need them to publish work that is useful, and it will only be useful if the results can be reproduced by others (not necessarily replicated, but at least reproduced.) Even negative results. But to do that, we have to…

3. Change what sort of scientific conclusions are publishable. Remember, I've suggested that a conclusion that is interesting and is supported by the data should be published. Yes, I'm back to changing the literature, but in another way. Because not only are we required to support an interesting conclusion with data, but we have to create the illusion (and it is *always* an illusion) that the story we are telling is the *whole story*. Well, stop it. If we can recognize that no story is ever complete, we can focus on what we actually think we know. If we can openly state what we *don't* know, without fear of collapsing our effort, we will make a large stride forward. Once upon a time, it was very desirable to publish a really intriguing phenomenon – it stimulated discussion, speculation, and research. Now, for some reason, this is frowned upon; every story has to wrap everything up tidily.

And this particularly applies to negative results, as we've already talked about. While there is an outcry for the need to be able to publish negative results, we cannot do so in this system. Because even if a negative conclusion is interesting, there is nevertheless an automatic response (engendering reaction) that says, "Okay, but what *is* going on then?" So it is never enough to say that the process does *not* work this way, we have to show how it *does* work. But the problem of negative results goes deeper than that. First of all, what makes a negative result interesting? Generally, this is only because someone else published the opposite – usually in a prominent journal – escalating the need to show they are wrong (in part so others won't waste too much effort on this, and partly because we tend to get a kick out of it). And unless the prominent journal has a policy that they will publish evidence that the conclusion of one of their papers has problems (even if the journal says they do, they *hate* to do this), the negative result usually finds itself in a much lower tiered journal, where it might be ignored.

This issue of publishing negative results is not so simple. There are rules. First, it is not enough to state that "we tried to do the experiment and couldn't make it work." As I've said before, I can't make *anything* work. So, if you are going to convince me that your negative result is valid, you have to control your experiment with things that do work. Second, and perhaps most usefully, it is valuable (but not always necessary) to approach the conclusion from another direction – design an experiment that

Journal of Cell Science (2014) 127, 2379–2381 doi:10.1242/jcs.155606

tests it, and ideally one that produces a positive result of the negative of the conclusion. This is the black swan approach: if the conclusion is that all swans are white, a description of a black swan is a positive result that disproves the conclusion. (My friend, Black Swan, has a great deal to say on this and the ensuing paradoxes, but we'll get to her another time.)

All of this said, there is a growing list of places to publish good quality negative results, many of them in the form of online, open access journals. If we can agree that such publications have value, then we can promote them further (and as I've suggested, cite them). But this brings me back to our trainees. Realistically, negative results should be carefully pursued only when it is a part of an overall research plan, and not an exercise in "cleansing" the literature. Anyone who devotes all of their time to showing what is wrong will not find out what is new (and correct), and that's what we really do in this business. So no, I am not saying that trainees should strive to produce negative conclusions, but when these occur in the course of a robust research effort, then yes, they should get credit for this.

But after all of this, with changes in what we can publish, what we value in our trainees, and how we perceive the value of our work, what if we still find that a lot of *stuff* is still wrong? How can we stop this flood of useless information? Well:

4. Maybe 'wrong' can be a little bit 'right.' A paper is published with an exciting conclusion, a positive conclusion, and the data strongly support it. We find the conclusion insightful and it drives our own thinking and affects the studies we are doing. But someone shows that an experiment in that paper cannot be replicated, or even reproduced. Terrible! But is it? Should we expunge the work from the literature and from our own efforts? I'll give you an example from my own experience. Several years ago, Dr. Toad published a series of papers that described the function of a molecule in a process I'm interested in. Since that time, a number of us (collectively, a 'bunch' but the collective for scientists might be more interestingly coined – I would say a 'mole' but that might hint at a very large number – perhaps a 'molette of scientists?' Sorry, Molette.) have come to doubt this conclusion. Along the way we have produced some pretty compelling negative results. But in doing so, we have also uncovered some (if I can say so) really interesting things that we are pursuing. So I would have to say that while I don't think Dr. Toad was entirely right, his studies have proven to be very useful to our own efforts, and I think (again, if I can say so) we're learning a great deal.

There's another aspect to reproducibility and 'correctness' that comes into play as well. Some important, vibrant areas of research are not very conducive to replication/reproduction of results. My friend, Red Fox, pointed this one out: if we find that commensal bacteria play a major role in, say, a behavioral experiment, it may be extremely difficult to reproduce this result in another lab (at least, until *all* of the details are understood). But should we readers be made to wait in ignorance of an interesting result because someone feels that the observations must be made independently?

Okay, Leon, okay – I know. I haven't addressed the last, very important issue in all of this. There are some researchers who, for whatever reason, publish a lot of rather prominent stuff and while they may get a reputation for being wrong, their careers flourish. And for those of us who are very careful to try to be as correct as possible, even at the expense of publishing the attention-grabbing, career-propelling, high-impact paper, it just isn't *fair*. These people should be punished, or at least *outed*, and there should be *consequences for being wrong*. Right? Well, here's the thing.

5. We have to choose what's most important. While it may be very cathartic to propose dire consequences for being wrong (at least, for those of us who think we're always right), I think that whatever we do to embark on such an endeavor will follow the law of unexpected consequences. (I'm not really sure that this is a law, really, but it's always stated as one. Basically it is a corollary of an actual law that we don't always know what's going to happen. Then again, I'm not sure that's a law either, or else there'd be a fine for getting things right. But I digress.) I suggest that it is really, really important that we sometimes take chances that we'll be wrong. Personally, I know I'm wrong all the time. In the vast majority of cases, we find out really quickly that I was wrong. Sometimes it takes years and years, because I was almost right. Look, I have a friend, Professor Hamster, who can't stand to be wrong. Any experiment she thought about, she could see several, or many, reasons why it wouldn't work. And if she did it anyway (not always the case) and it worked (which was sometimes the case), she worried it was wrong anyway, and went back to the start. She's struggling. We have to take chances, I suggest, and even if this means that some things we ultimately publish are actually wrong, I think (provided we follow 1–4, above) we'll get an awful lot right. So we have to choose – can we live with a few scientists who get a reputation for being wrong a lot (who, as often happens, get some things right; as we say in the lab, "Just because Toad says it, doesn't mean it's wrong!") but whose careers do very well? Or do we clamp down on them, and any others, who get something wrong, leading to much more careful work but no risk taking. We never have the whole story, we never have it all right, and there are always conditions under which a process will not work, or will work differently. Yes, it's hard.

By the way, some of you may wonder who Leon is. He's in Blade Runner, and he's a replicant. Are you telling me you still haven't seen it?

Holden: You're in a desert, walking along in the sand, when all of a sudden you look down...

Leon: What one?

Holden: What?

Leon: What desert?

Holden: It doesn't make any difference what desert, it's completely hypothetical.

Leon: But, how come I'd be there?

Asking questions is a good thing. Not being afraid of being wrong is a good thing, as long as we really *try* to get it right. Sometimes we need to think about Blade Runner-esque dystopia, so we don't go there, but other times we can think about Roger Rabbit, and have some fun being foolish, and yeh, even silly. Science is Toon-Town, where everything is possible, and it's cyberpunk LA, where nothing works as it's supposed to. We feel around in the dark for a light switch, that's all, and sometimes we turn on the garbage disposal. It's what we do. Maybe we can do it better, but I am continuously amazed by how much we *do* do. (Roger wanted to make a terrible joke that Leon won't understand, but me, I'm going to bed. Nighty night.)