

OPINION

## Excessive trust in authorities and its influence on experimental design

*Tung-Tien Sun*

Most graduate curricula in biological sciences offer courses that cover various scientific disciplines, but they give relatively little formal instruction in experimental design. Students learn the latter primarily through hands-on experience in the laboratory, and some find this learning process bewildering and frustrating. So, what is the root of the problem, and how can young researchers get experiments to work more predictably and reproducibly?

Central to biological and other experimental sciences is the practice of formulating a hypothesis and subjecting it to vigorous experimental testing. Most graduate schools teach students how to survey the scientific literature and emphasize the scientific paradigms that define important and solvable problems<sup>1</sup>, and therefore provide a basis on which to formulate hypotheses. However, few graduate schools offer formal instruction on how to become a good experimentalist. Our focus on the intellectual aspects of research (asking the right questions) is completely appropriate, but it should always be accompanied by a solid training in laboratory investigation (getting high-quality data). Traditionally, we place our students in laboratories and hope that they will learn all there is to know about how to design and carry out good laboratory experiments from the example of senior investigators and colleagues. Although many students function well under such an informal, and rather variable, apprentice system and become good experimentalists, some new students find this learning

process bewildering and frustrating. This is exemplified by the following comments that I have heard. One student said, "I asked several senior graduate students and postdocs how to do an experiment. But everybody gave me a different answer. Whose advice should I trust?" Another said, "I followed a postdoc's detailed written protocol to the letter and I even used all his solutions that worked for him. But my experiment failed and I lost three months. I am really upset!"


So, what is the root of the problem, and is there anything we can do to help these students? I propose that excessive trust in authorities is a key underlying cause of many experimental failures. In this article, I discuss how an awareness of this problem and a few practical considerations, including the concept of risk assessment, can help to get most laboratory techniques to work more predictably and reproducibly. Many of the practical tips that I discuss are so instinctive and logical that most experienced scientists practice them naturally. However, I hope that young investigators will find these messages helpful and reassuring. Finally, although most of my comments are directed towards the design of biochemical experiments, they are generally applicable, with minor variations, to other experimental sciences (for further useful information, please refer to [REFS 2–7](#)).

### **Trust in authorities and risk assessment**

The American Heritage Dictionary defines authority as "an accepted source of expert information or advice", and there are many elements in an experiment that might be

regarded, sometimes uncritically, by a new experimenter as an 'authority'. These include professors' or other experts' opinions, books and published experimental protocols, labels on reagent bottles, and instrument readouts. When any of these are perceived as 'authorities', it means that the experimenter has decided that nothing can go wrong with such information or reagents. Accepting many (and sometimes all) elements as authorities simplifies a person's thinking, is convenient and gives a sense of security ("Well, the experts said so"). But each of these authorities constitutes a blind spot in our experimental design: when an experiment fails, we tend to overlook these elements as potential causes of the problem, and this compromises our ability to troubleshoot.

Although some purists might say that scientists should never trust any authority at any time, when carrying out experiments we, for practical reasons, invariably choose to trust some of the elements. The increased reliance on commercial experimental kits that frequently contain undisclosed, proprietary formulas, together with a less rigorous background in chemistry and other physical sciences, underpins a worrisome recent trend in which more of our students seem willing to trust authorities, sometimes excessively ("As long as the kit works, why should I care how?"). In this article, I emphasize the importance of deciding how much we can trust certain elements as authorities under a given circumstance. A practical question that we can ask ourselves is, how much can I afford to make (potentially preventable) mistakes in this experiment? For an experiment that takes only a day or two and uses inexpensive reagents, we might decide that we can afford to make such mistakes. But, for experiments that take months to do and/or require expensive or irreplaceable materials, we cannot afford to make potentially preventable (also known as stupid) mistakes. For example, the generation of transgenic or knockout animals takes months before we know the outcome, so it would be foolish to use casually



**Figure 1 | The importance of planning in experimental design.** This figure outlines a decision-making process that should take place before every experiment, and it emphasizes the importance of detailed planning and thorough understanding in experimental design. Rushing into experiments without thoughtful planning invites failure. Perhaps this is why someone once said, “Seventy percent of whether your experiment will work is determined before you touch the first test tube.”

characterized complementary DNAs or genes for this purpose. Another ‘experiment’ in which no preventable mistakes can be tolerated is the space-shuttle project. Although it has failed several times despite all humanly possible efforts, Gene Kranz — a former NASA flight director — summed up the attitude of NASA workers by saying, “failure is not an option” (the title of his 2000 New York Times bestseller book). It is therefore important to assess risk before carrying out every experiment: the outcome of this analysis (risk assessment) determines how much we can afford to blindly trust various elements that might be regarded as authorities (risk management; FIG. 1).

### Failure in risk assessment

In addition to the opinions of professors, so-called ‘experts’ and textbook authors, who are clearly not infallible, commercial reagents and materials can also be a problem. For example, a young mouse geneticist was hired by a well-known university some years ago, and was given a laboratory and supervised several graduate students. The geneticist ordered batches of mice from a leading animal supplier to carry out breeding experiments. After several years of hard work, it was discovered that some of the supposedly inbred mice were impure or of the wrong strain. All of the geneticist’s painstaking and time-consuming work amounted to nothing, and led to no publications at the time of tenure review. The geneticist sued the animal supplier, but was told that the animals were shipped with a warranty that essentially stated, “This warranty limits our liability for replacement of the product. No other warranty of any kind, including fitness for a particular purpose, is provided.”

Knowing the situation, if you were the young mouse geneticist, what would you do differently? Risk assessment indicates that if you order some Balb/C mice for the purpose of isolating liver alcohol dehydrogenase, which takes only a few days to complete, you might decide to trust the mice (or, more correctly, the company that proclaimed that the mice are of a certain breed) as an authority. On the other hand, if you purchase the same batch of mice to carry out long-term breeding experiments — the outcome of which will determine your career and the thesis of your graduate student — would you still want to trust the animal supplier blindly as an infallible authority? The answer is clearly no. In this case, it would be extremely worthwhile to independently confirm the genetic purity of the mice. This example clearly highlights the importance of risk assessment, which determines how much we can trust the various elements under a given circumstance (and, in this example, leads either to the acceptance or the rejection of a batch of mice as an authority, depending on the circumstance).

### Maximizing experimental success

When you have decided, on the basis of risk assessment, that you must get an experiment or a new technique to work, what can you do

to improve your chances of success? As Gustave Flaubert said, “God is in the details.” To pay attention to even minute details, to the point that you become a perfectionist, will save you invaluable time in the long run.

**Thorough planning and understanding of the protocol.** You should plan your experiment carefully: write a step-by-step flowchart that details the entire procedure from the beginning to the end (see below), and thoroughly understand every step. You should be able to answer questions like: “What is the purpose of having EDTA here? What does it do and how does it work? Can it be substituted with EGTA?”; “How does SDS gel electrophoresis work? What are the structures of SDS, acrylamide, bisacrylamide, ammonium persulfate and TEMED, and what do they do in this gel system?”; and “What is the chemistry of this bifunctional crosslinking reaction? How does pH affect its crosslinking efficiency?” A thorough understanding of every step enables you to comprehend how the protocol works, and to troubleshoot more successfully if required. You can compare the flowchart you generate with the protocols from several sources to see whether they are consistent. Finally, you should check your detailed design of a particularly crucial or time-consuming

### Box 1 | The ‘N+(N-1) rule’: optimizing a reaction condition by titration

Titration of key elements			
Enzyme A	5 μg	1, 2, 5*, 10 μg	$N_1 = 4$
MgCl <sub>2</sub>	2 mM	1, 2*, 5 mM	$N_2 = 3$
KCl	25 mM	10, 25*, 50, 100 mM	$N_3 = 4$
pH	8	7, 8*, 9	$N_4 = 3$

<b>Conventional rule</b> $4 \times 3 \times 4 \times 3 = 144$	<b>N + (N-1) rule</b> $4 + (3-1) + (4-1) + (3-1) = 11$
------------------------------------------------------------------	-----------------------------------------------------------

By definition, titration means to determine the optimal concentration of a component in a reaction mixture by measuring the effects of systematically changing its concentration. In the example given in the main text, several of the reaction components are particularly important — that is, 5 μg of Enzyme A, 2 mM MgCl<sub>2</sub>, 25 mM KCl and pH 8.0 (whatever the buffer). We can therefore carry out the titration by comparing the effects of varying: the enzyme concentration at 1, 2, 5\* and 10 μg ( $N_1 = 4$ ; meaning that for variable number one — that is, the enzyme — we will test four concentrations); the MgCl<sub>2</sub> concentration at 1, 2\* and 5 mM ( $N_2 = 3$ ); KCl at 10, 25\*, 50 and 100 mM ( $N_3 = 4$ ); and pH at 7, 8\* and 9 ( $N_4 = 3$ ) (see figure; asterisks highlight the originally recommended amount). To do all these titrations properly — that is, for us to change only one variable at a time (a cardinal rule) — this would require a total of  $N_1 \times N_2 \times N_3 \times N_4$  (or  $4 \times 3 \times 4 \times 3 = 144$  reactions). This is, of course, impractical. Fortunately, as we know that all the data points in a biochemical reaction almost always form rather smooth curves, instead of jumping all over the place (barring poor pipetting skills), we do not need to carry out all of these 144 reactions. Rather, we can take a shortcut using a procedure that we call the ‘N+(N-1) rule’, which tells us the minimal number of reactions that we need to do in a titration. In the above example, we can add up the four Ns, subtracting one from each of them except the first one:  $[(N_1) + (N_2-1) + (N_3-1) + (N_4-1)]$ . So, we need a total of  $4 + (3-1) + (4-1) + (3-1) = 11$  reactions only (see BOX 2 for how these 11 reactions can be designed).

experiment with your mentor or experienced colleagues before you carry out the experiment, as they might be able to give you some advice/suggestions that could save you a lot of time in the long run.

#### Importance of positive and negative controls.

As an experiment can be ruined if any of its many reagents and steps fail, you must include vigorous controls to ensure that the experiment is working the way you expect — no matter whether it is a two-day or six-month experiment. A control is an experiment that is conducted in exactly the same way as, and carried out side-by-side with, the other experiments that involve the unknowns (that is, using the same reagents and going through identical steps), except that, for the controls, you know precisely what the outcome should be. The failure of your control experiments raises a red flag and indicates that you must troubleshoot before you can move on.

There are two types of control. A positive control is an experiment that is expected to produce certain positive data. If such a control produced unexplainable negative data, then your negative data are meaningless. Conversely, a negative control is an experiment that is expected to produce certain negative data. So, if such a control produced unexplainable positive data, then your positive data are meaningless. As there is no way to predict whether you will get positive or negative results before you do an experiment, you need both types of control for every experiment. Otherwise, your results could be confusing or, worse, misleading. As a set of controls is required for every experiment (such controls are a constant ‘overhead’), it pays to maximize the number of experimental tubes or reactions to reduce the overall overhead cost and, more importantly, to allow the comparison of a wide set of data that were generated within any given experiment.

Failures in control experiments indicate that something in the procedure has gone wrong. It could be: an enzyme or antibody that worked perfectly yesterday but that was ‘cooked’ overnight, unbeknown to you, owing to refrigerator malfunction; a microscope filter that was changed by a previous user; or a new batch of a chemical that you have not tested before. These examples make it almost comical to hear questions or comments like, “I did my controls last week. You mean I have to do them every time?” (answer: “Yes, absolutely.”) or, “I will do the experiment first. If it works then I will do the controls.” (response: “Without proper

#### Box 2 | The ‘N+(N-1) rule’: constructing a table to titrate key components

To titrate the components listed in BOX 1, you can make a table with 11 rows for a total of 11 reactions (as calculated in BOX 1; see table). Above these rows enter: Enzyme A (with 4 columns indicating the test concentrations — that is, 1, 2, 5\* and 10  $\mu\text{g}$ );  $\text{MgCl}_2$  (1, 2\* and 5 mM); KCl (10, 25\*, 50 and 100 mM); and pH (7, 8\* and 9) (see BOX 1; asterisks highlight the originally recommended amount). To titrate the enzyme, select 1  $\mu\text{g}$  for tube 1, 2  $\mu\text{g}$  for tube 2, 5  $\mu\text{g}$  for tube 3, and 10  $\mu\text{g}$  for tube 4; for all four tubes, we will select (or gamble by using) the recommended ‘standard’ values for all other reagents. The results from these four tubes will titrate the enzyme concentration by changing only a single variable (the enzyme concentration) at any one time. Note that tube 3 (highlighted by<sup>†</sup>) is unique, in that all four components happen to be of the standard value (this is the only reaction that would have been done if we had decided not to titrate). The result of this tube can therefore be used, repeatedly, when we titrate all other components (therefore ‘N-1’). Next, we can titrate  $\text{MgCl}_2$ , by selecting 1 mM (tube 5), skipping the recommended 2\* mM (this tube would be identical to tube 3), and selecting 5 mM (tube 6). As noted before, for these two tubes we will keep all the other variables constant by selecting the recommended values. Do the same for titrating KCl (tubes 7–9) and pH (tubes 10 and 11). You can then carry out the experiment, and enter the results (in this case, the incorporation of a radioactive precursor in counts per minute (cpm)) in the last column (see BOX 3 for how to plot the results).

Reaction	Enzyme A ( $\mu\text{g}$ )				$\text{MgCl}_2$ (mM)			KCl (mM)				pH			Results (cpm)
	1	2	5*	10	1	2*	5	10	25*	50	100	7	8*	9	
1	+					+								+	30
2		+					+							+	60
3 <sup>†</sup>			+					+						+	65
4				+					+					+	68
5			+		+					+				+	50
6			+				+				+			+	60
7			+			+		+						+	75
8			+			+				+				+	50
9			+			+					+			+	40
10			+			+				+			+		62
11			+			+				+				+	55

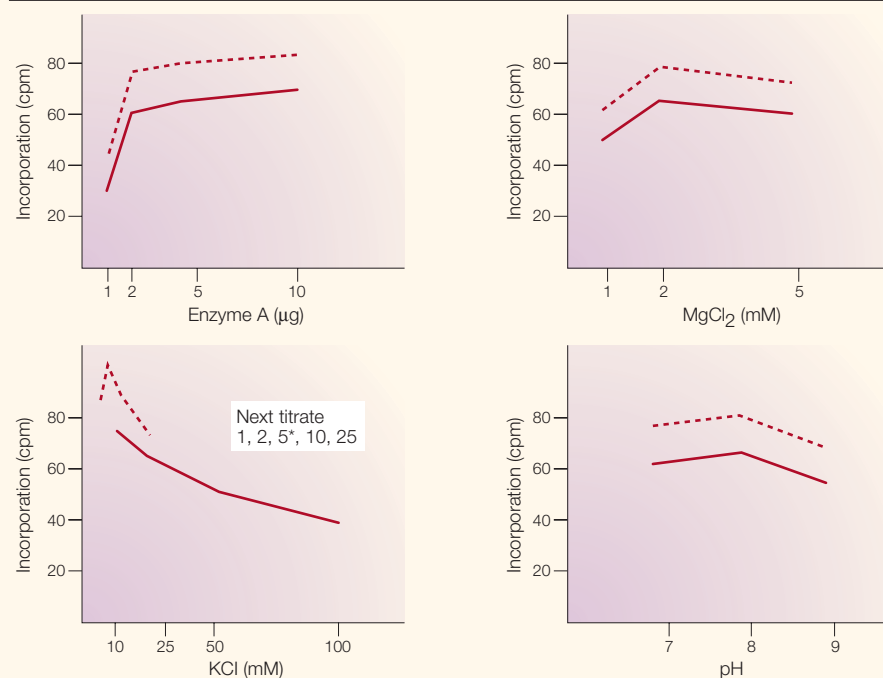
controls, how could you know that your experiment has ‘worked?’) or, “This seems to be a straightforward procedure that everybody and his brother-in-law can get to work. I trust that everything will work properly, so I don’t need to do controls.” (response: “In God we trust, everybody else must show controls.”).

#### Titration key components: the N+(N-1) rule.

The titration of several key components is an effective way to get a new laboratory technique to work optimally with a minimal amount of time and effort. As an example, suppose that you want to carry out a new (hypothetical) enzymatic assay. Although the reaction mixture contains many components, on the basis of your thorough understanding of the protocol you realize that several of them are particularly important: 5  $\mu\text{g}$  of Enzyme A, 2 mM  $\text{MgCl}_2$ , 25 mM KCl, and pH 8.0 (whatever the buffer). To get this new technique to work, you need to keep several points in mind. First, it is impossible for you

to duplicate precisely this recommended reaction condition as it was used in the hands of the original investigators. Second, if you decide to try this technique using only the recommended ‘standard’ condition, and you obtain a result of, say, 150 cpm (counts per minute; as a measure of the amount of incorporated radioactive precursor) versus a negative control of 30 cpm, how good is this result? Can you rule out the possibility that the incorporation could have been as high as 500 cpm if you had used the truly optimal condition? Finally, and most importantly, if the experiment produces a rather poor incorporation of, say, 50 cpm, which is only slightly above the 30 cpm of the negative control, what went wrong and what can you do to improve the outcome? Titration is a good way to answer all of these questions and allows you to move forward with confidence. In BOXES 1–3, I describe how to use an ‘N+(N-1) rule’ to titrate many variables using a minimal number of tubes or reaction mixtures.

## Box 3 | The 'N+(N-1) rule': plotting the titration results



Continuing from BOXES 1,2, plotting the results of tubes 1, 2, 3 and 4 shows the titration of Enzyme A; tubes 5, 3 and 6 of MgCl<sub>2</sub>; tubes 7, 3, 8 and 9 of KCl; and tubes 10, 3 and 11 of the pH. If the plots confirm that the recommended values work best in your hands, you can now confidently move on to do some 'real' experiments. But, not infrequently, things might not work as expected. As an example, in a hypothetical titration (see solid lines in the graphs), although the recommended values of Enzyme A, MgCl<sub>2</sub> and pH indeed work optimally, the KCl curve shows no optimum — the optimum seems to be lower than even the lowest value (10 mM) that you tested, despite the reported optimum of 25 mM. This could mean that the original paper was wrong or, perhaps more probable, that there is something wrong with your KCl solution. If the solution is correct, one thing you could do is to repeat the titration, except that the next time you would titrate KCl at 1, 2, 5\*, 10 and 25 mM (which partially overlaps with the previous titration of 10, 25, 50 and 100 mM to provide a continuity between the two sets of data). Note that in this second titration, we would use (by guessing) 5 mM KCl (highlighted by an asterisk) as our estimated optimal concentration when titrating the other three components (if this guess is way off we might need to do a small-scale, third titration of the KCl alone to elucidate its optimal concentration, but this is rarely necessary). Also note that the use of this more optimal concentration of KCl can (barring interactions among the components) result in a higher incorporation of radioactive precursor in counts per minute (cpm) in all of these titrations (dashed lines). This kind of systematic approach gives you confidence in your reaction condition, allows you to improve or troubleshoot the reaction, and can greatly enhance your ability to get an experiment to work predictably and reproducibly.

productivity of their research. You can generate better hypotheses if you know the literature and present models well. It might also help if you can express your hypothesis in a schematic diagram to highlight the relationships among the various elements, so that you can better predict what will happen if you introduce variables into the experimental system.


Second, instead of following detailed, published experimental protocols directly, you can convert the protocol into a flowchart that contains a minimal number of words (FIG. 2). However, it should still contain sufficient details such that you could give it to a colleague and he/she would know exactly what to do to complete the entire experiment. Writing a flowchart serves several useful purposes. It gives you a chance to think about each step as you write it down. This enables you to anticipate potential problems before starting an experiment, so that you do not have to make rushed, last-minute decisions while the experiment is taking place. In addition, without all the extraneous words, a flowchart is much easier to follow while you are carrying out the experiment. The flowchart also allows you to use a coloured pen to tick off every step right after you have carried it out (see red ticks in FIG. 2). This makes it easier for you to know precisely where you are in the procedure and helps to avoid careless mistakes. Finally, the flowchart gives you plenty of space to record any observations or last-minute changes in the experimental conditions, which can sometimes lead to unexpected findings (FIG. 2).

Third, prepare all reagents and solutions beforehand. After you have finished designing an experiment, calculate the total amount of each reagent and solution that you will need for the entire experiment. Prepare sufficient solutions (at least for those that do not require expensive or labile reagents) to allow you to repeat the experiment several times. If you need to use certain reagents, such as restriction enzymes or fluorescent antibodies, that are parts of the common laboratory stock, it is a good idea to 'squirrel away' the appropriate amounts of these materials before you begin your experiment. Otherwise someone else might use them up before you need them, without ordering replacements (a terrible thing to do indeed).

Fourth, thoroughly document your results. Data should be entered, as a rule, directly into your notebook. Entering data temporarily on scrap paper for later transcription risks the loss of precious data. Printouts, X-ray films, photos and other documents should be clearly indexed (date and experiment number) and taped securely in the appropriate places in the laboratory notebook.

**A well-kept notebook.** Although a laboratory notebook is usually thought of simply as a place to enter data once experiments are complete, it is an equally important place for experimental design. A well-kept notebook should contain, under the heading of an individual experiment, many useful items including the date, the hypotheses and/or questions, key references and notes, experimental protocols (in the form of flowcharts), reagents and solutions, results and discussion. Several particularly important points are discussed below.

First, it is important to state clearly the hypothesis that you want to test. The importance of a hypothesis was emphasized by Peter Medawar, who said, "No experiment should be undertaken without a clear preconception of the forms its results might take; for unless a hypothesis restricts the total number of possible happenings or conjunctions of events in the universe, the experiment will yield no information whatsoever."<sup>3</sup> Indeed, for two scientists who are equally competent experimentalists, it is their ability to pose good questions and hypotheses that determines the quality and



**Figure 2 | The flowchart of an experimental protocol.** A flowchart allows you to tick off, with a coloured pen, every step that you have done, and provides space for recording observations (for example, 'cloudy'). It also allows you to record unintended, last-minute changes in the protocol (like changes in the incubation condition from a planned 37°C for 30 min to 42°C for 36 min, owing to an overheated water bath and a delay in stopping the reaction). This way, at the end of the experiment, you know precisely what actually happened in every single step. After you have repeated this experiment several times, you might find that one repetition worked better than others. If it happens to be the one with an overheated water bath and an extended incubation time, you might have made an important chance discovery. Francis Bacon said, "Truth emerges more readily from errors than from confusion.". In the above example — in which you have accidentally altered the incubation step and have obtained a better result — if you have kept a good record in your laboratory notebook, which allows you to trace back exactly what happened, you have learnt from an 'error'. But if you do not have a good record and are puzzled by the apparent non-reproducibility of your data, you will remain 'confused'. OD, optical density; rpm, revolutions per minute.

Large X-ray films or other data, such as dried gels, which cannot fit into the notebook, should be clearly labelled and filed.

Finally, write a thoughtful discussion. This is perhaps the most important and fun part of the experimental record. If things did not work as well as you had hoped, you can discuss what you can do to improve your experimental protocol, which will lead to a revised flowchart. If the experiment worked well technically, are the data consistent with your hypothesis? If so, can you think of some other independent and perhaps even more stringent test? It is important to bear in mind here that you can never 'prove' a hypothesis — you can only build up supporting evidence so that you become progressively more 'comfortable' with its validity. However, the outcome of a single, crucial experiment can disprove a hypothesis ('asymmetry of proof')<sup>3</sup>. In the case that your data contradict the hypothesis, can you modify some parts of the hypothesis to account for the new findings, or do you need to replace the original hypothesis with a new one? In all these cases, you have made significant advances. A thorough and in-depth analysis of the data can therefore generate many interesting questions

and ideas, which will lead to new hypotheses that you can test in your next experiment.

### A matter of scientific attitude

People differ widely in the degree to which they trust authorities. Important factors that determine how much a person trusts authorities include their educational and cultural background, personal experience and, to some extent, natural inclination (that is, the person's genetic make-up). As a complex product of all of these factors, everyone develops their own individual way of decision-making. At one extreme are people who blindly and unconditionally trust all authorities. These people can run into trouble at every step of the way, not just in experimental design, but also in other decisions they make in life. At the other extreme are people who do not trust any authority and are sceptical of everything. These people often dwell on every little piece of information, and spend inordinate amounts of time on minor details. They fail to prioritize, and can be inefficient, indecisive and bogged down by details. What we hope for is a balance between these extremes — a balance in which a person can adjust, in a logical and flexible way, how much they trust various authorities depending on the circumstances and on the basis of risk assessment.

The degree to which a person trusts authorities tells us a lot about their decision-making process and logic. As this process is developed by an individual over a period of years or even decades, it is a deeply ingrained behavioural pattern. To this extent, the pattern becomes an important part of a person's scientific attitude or, to put it more broadly, personality. So, students who excessively trust authorities tend to make all kinds of technical errors, which can usually be traced back to them placing inappropriate trust in authorities and to the accompanying faulty logic in decision-making. Such individuals usually have only a superficial understanding of the protocols and only carry out controls haphazardly. They tend not to be able to troubleshoot when they run into technical problems, and their data are often non-reproducible. They frequently have problems in communicating with their mentors and other scientists. With repeated help from their mentors and co-workers, they can — after much struggle — identify existing technical problems and rectify them, but new problems always emerge. This can, of course, be a very frustrating experience for the student and mentor alike. However, recognizing that excessive trust in authorities can be the root of many such problems will enable mentors to better understand these problems and to better help their students.

### Concluding remarks

A key idea that I hope I have developed in this article is self-reliance: you are mainly responsible for the success of your experiments. If experiments fail, blaming someone else afterwards (the so-called authorities) for having given you the wrong protocol or bad advice, or for providing you with the wrong reagents, solutions or animals, does not help. Francis Bacon put this well. He said, "If a man will begin with certainties, he shall end with doubts; but if he will be content to begin with doubts, he shall end with certainties." If, when we start an experiment, we inappropriately trust many factors as authorities, and think that nothing can go wrong with these (that is, we begin with certainties), our experimental design tends to be superficial and clumsy, and the experiments frequently fail. When this happens, we will not understand what has gone wrong and we will end with doubts. On the other hand, if we try to foresee as many problems as possible and we do not blindly trust authorities (that is, we begin with doubts), our experiments will be more likely to work, and we will end with certainties.

*Tung-Tien Sun is at the Epithelial Biology Unit, Departments of Dermatology, Pharmacology and Urology, New York University Cancer Institute, New York University Medical School, 550 First Avenue, New York 10016, USA. e-mail: sunt01@med.nyu.edu*

doi:10.1038/nrm1429

1. Kuhn, T. S. *The Structure of Scientific Revolutions* (ed. Neurath, O.) (Univ. Chicago Press, Chicago, 1970).
2. Wilson, E. B. *An Introduction to Scientific Research* (McGraw-Hill, New York, 1952).
3. Medawar, P. B. *Advice to a Young Scientist* (Basic Books, New York, 1981).
4. Grinnell, F. *The Scientific Attitude* (Westview Press, Boulder, Colorado, 1987).
5. Barker, K. *At the Bench: A Laboratory Navigator* (Cold Spring Harbor Laboratory Press, 1998).
6. Ramón y Cajal, S. *Advice for a Young Investigator* (MIT Press, Boston, 1999).
7. Phillips, E. M. & Pugh, D. S. *How to Get a PhD: A Handbook for Students and Their Mentors* (Open University Press, Buckingham, 2000).

### Acknowledgements

I thank S. Miller (Johns Hopkins University, Baltimore, Maryland), R. Behringer (The University of Texas M. D. Anderson Cancer Center, Houston, Texas), F. Watt (London Research Institute, ICRF) and many of my colleagues and students at New York University (NYU) Medical School for critically reading the manuscript. This article is based on the Borden Lecture that I delivered during the 2003 Joint Annual Meeting of the British Societies of Cell Biology and Developmental Biology. The lecture is also a part of a graduate course that I teach, entitled 'Scientific methods: survival skills for young biomedical researchers', which covers experimental design (discussed in this article), critical analysis of scientific literature, scientific writing, oral presentation and scientific integrity. The course has been a requirement for all first-year graduate students of the Sackler Institute of Graduate Biomedical Sciences at NYU Medical School since 2001.

### Competing interests statement

The author declares that he has no competing financial interests.

### Online links

#### FURTHER INFORMATION

Tung-Tien Sun's laboratory: <http://www.med.nyu.edu/sun/>  
Access to this links box is available online.