Striving for Solid Science: Preregistration and Direct Replication in Experimental Psychology

Contributors: Dorottya Rusz, Erik H. Bijleveld & Michiel A. J. Kompier
Pub. Date: 2018
Access Date: May 15, 2018
Academic Level: Advanced Undergraduate
Publishing Company: SAGE Publications Ltd
City: London
Online ISBN: 9781526439987
DOI: http://dx.doi.org/10.4135/9781526439987
©2018 SAGE Publications Ltd. All Rights Reserved.
This PDF has been generated from SAGE Research Methods Cases.

©2018 SAGE Publications, Ltd.. All Rights Reserved.

Abstract

Recently, experimental psychologists have been thinking a lot about how to do research in such a way that their findings can be replicated. As a result, it is becoming more and more common (a) to preregister one's own hypotheses and analysis plan online and (b) to conduct direct replications of one's own studies. In this Research Methods Case, we discuss our personal experiences with preregistration and direct replication. Illustrated by two projects from our own laboratory, we reflect on the costs and benefits of using preregistration and direct replication. Also, we discuss how preregistration and direct replication attempts may seem to harm personal career development, but at the same time can be inspiring and productive.

Learning Outcomes

By the end of this case, students should be able to

- Describe what transparent scientific practices are
- Evaluate the benefits of preregistration and direct replications
- Evaluate the costs of preregistration and direct replications
- Explain how and where to preregister their research plans
- Discuss how to deal with unsuccessful replications

Solid Science in Theory

When scientists try to answer their research questions, they are usually really excited about this. They feel that their topic of interest is extremely important, they are genuinely motivated to discover new pieces of knowledge, and, they invest blood, sweat, and tears into designing their studies. When a scientist makes a new discovery, maybe after months or even years of work, this tends to makes them really happy. After all, their discovery helps science—and helping science is what their job is all about.

Yet, "to help science" is not the only goal scientists have: in most cases, scientists are also trying to make a career of their own. They dream of having their work published in prestigious journals, of giving talks for huge audiences, and of getting cited frequently by other researchers—in a sense, to be a rock star in science. To accomplish such star status, a scientist, of course, needs to do important, significant discoveries. To become a rock star, a scientist needs their findings to be not just important and new, but first and foremost, statistically significant ($p < .05$). To become the Ariana Grande or the Justin Bieber of science, you better produce significant results.

©2018 SAGE Publications, Ltd.. All Rights Reserved.

This situation has been causing some concern in psychology (Nosek et al., 2015). The problem seems to be that most psychology journals accept only those findings that are statistically significant. As scientists are motivated to publish, they are also motivated to find statistical significance. Yet, unfortunately, there are many ways to stumble upon significant findings, for instance, by trying out different types of analyses, by looking at different subsets of data (e.g., only examining right-handed participants, or only taking into account female participants) or simply by examining a large number of variables. Although these practices may seem harmless (after all, to explore data is part of the job of almost any scientist), using these practices in an intransparent way causes trouble (Nuzzo, 2015). A recent study revealed that out of 100 studies published in three psychology journals, only about one third could be replicated (Open Science Collaboration, 2015). This means that there are many findings out there that we simply cannot assume to be true. Altogether, many psychologists are concerned that this problem stems, at least in part, from our motivation to find significant results, which we need to publish, which we need to make a career.

As a response to these concerns, pioneers of an *open science culture* (see Box 1) came up with some suggestions that could improve the way we do science (Munafò et al., 2017). They feel that, instead of focusing on how significant or novel results are, journals should care more about whether the study was done in a transparent way. Working in a transparent way involves, for example, *pre-registration* (see Box 1), which means that researchers make their hypotheses, experiments, and data analysis plans available to others *before* they start collecting data (Nosek et al., 2015). This is done to make sure that researchers do not change the hypotheses after already knowing the results, and to make sure that they are not trying many different analyses on the same data just to get significant results. Working in a transparent way also involves *directly replicating* (see Box 1) one's own (and other's) findings. If something is replicable, it means that we can truly rely on those findings (Open Science Collaboration, 2015). In sum, working in a transparent way might help science by producing results that we can actually believe in.

Box 1. Glossary.

Open science culture—refers to a culture where all scientist make their research ideas and processes available to others (Nosek et al., 2015).

Preregistration—making the hypotheses, methods, and data analysis plans accessible to others before collecting data. Researchers can do this on a website (e.g., Open Science Framework, OSF) by uploading their research plans in a document that cannot

©2018 SAGE Publications, Ltd.. All Rights Reserved.

be changed from the moment it is uploaded (Munafò et al., 2017).

Direct replication—the process of repeating experiments to see if the results are the same across different situations (Maxwell, Lau, & Howard, 2015).

My research team and I really like the idea of working in this transparent way—that is, we like the idea of using both pre-registration and direct replication. So, we adopted this way of working during my PhD.

Solid Science in Practice

In this section, I guide you through two research projects of my PhD, in which one of my main goals was to work in a transparent way. I also reflect on how and when this way of working was difficult and challenging, but also when it was productive and inspiring, both in a personal and in a professional sense.

Case Study 1: Designing a New Computer Task

Coming Up With a Research Question

At the time I started my PhD, I became very interested in why people (including myself) constantly check their smartphones, even when they actually need to do something else, like studying or working. I was motivated to find out more about the psychological processes that drive this repetitive phone checking, which eventually leads to distraction from that which people are currently doing (e.g., studying a textbook). To investigate this question, I came up with a research proposal to examine the reason why people get distracted. I came up with a theoretical model that suggests that people become distracted, because they have an inherent tendency to seek rewards in their environment (see Box 2). So, people look at their smartphones—instead of focusing on studying or working—because it carries high social *rewards* to them (e.g., probably, it is more rewarding to get a like on your new Facebook profile picture than reading a boring textbook). When people are doing some difficult mental task (such as reading, studying, working), do they get distracted especially by things that carry value to them? This is the question we were trying to answer. We expected that people will get more distracted by things that carry high value to them than things that carry low value to them.

Box 2. Theoretical Background.

Distractions stemming from reward seeking behavior

©2018 SAGE Publications, Ltd.. All Rights Reserved.

When people carry out cognitive tasks, they often get distracted by irrelevant information (e.g., their smartphone ringing). In past research, these distractions have been thought of stemming from people's capacity limitations (Theeuwes, 2004). That is, people are, sometimes, unable to distinguish between what is relevant and irrelevant to the task, and mistakenly allocate attentional resources to irrelevant information.

However, people are not only driven by external information in the environment, but have motivational states too, which also play an important role in which information people attend to (Cohen, McClure, & Yu, 2007). These motivational states are responsible for regulating attention and action toward rewarding outcomes in the environment (e.g., food, money, likes on social media; Botvinick & Braver, 2015; Braver et al., 2014). Sometimes, these rewarding outcomes can be irrelevant to the task at hand (e.g., when you are doing your homework, but you get a Facebook invite for a party), so if people attend to them, it could lead to distraction from the task. In sum, it is possible that people do not get distracted from their tasks because they have capacity limitations, but because they constantly pursue rewards, which sometimes happen to be irrelevant to their tasks at hand.

This is the idea that we based on our experiments and we tested in Case Study 1. To be able to work on this project, I wrote a research proposal that needed to be accepted by a science committee in the university. After the committee accepted my proposal, I could start running my first experiments.

Implications

Distractions are highly prevalent, especially with current technological developments—people often get distracted at work (Jett & George, 2003), at school (Cheever, Rosen, Carrier, & Chavez, 2014), or even during driving (Caird, Johnston, Willness, Asbridge, & Steel, 2014), which can have fatal consequences. Yet, the psychological mechanisms underlying these distractions are not entirely clear yet. Our research could help to get a better understanding of the potential causes of distractions. This could inform policy makers and help in designing interventions at work, in the classroom, or in traffic.

Designing a New Computer Task

We designed a new computer task that could help us answer our research question (see Figure 1). This computer task had two parts. In the first part, participants learned that a certain color (e.g., blue) was associated with reward. More concretely, whenever they would see that color (e.g., blue), they earned some money. In the second part of our task, participants were asked to add up numbers. This was quite difficult, and participants really needed to concentrate to do this well. Yet, sometimes, they would again see the color they had previously learned to be related to earning money. Would this color distract our participants? Would this color make our participants *worse* at adding up the numbers? This is what our new task allowed us to test.

Figure 1. Our computer task to measure distraction by valuable cues.

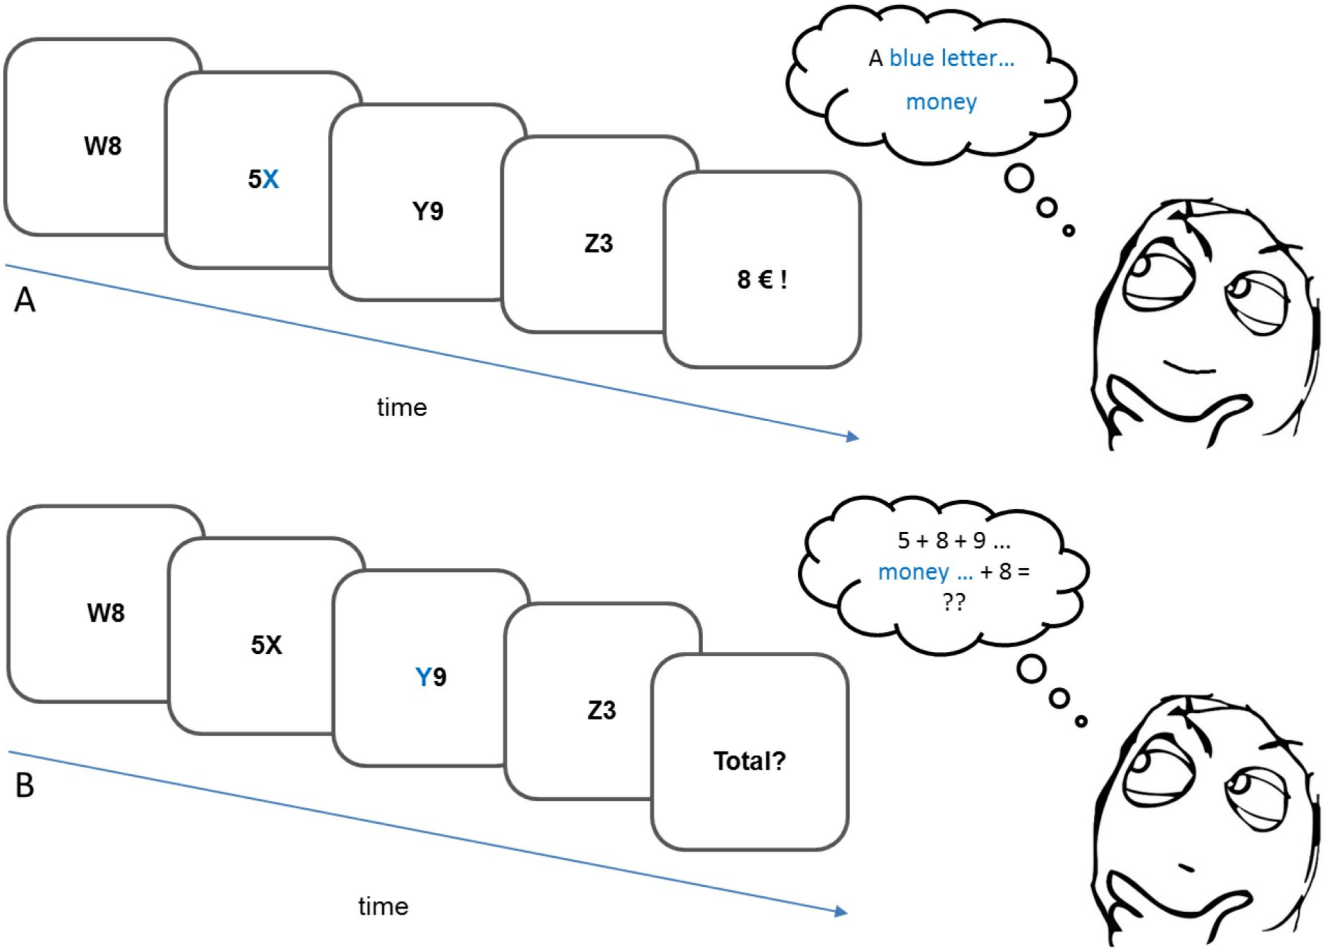

Note: (a) The first part of the task was a learning phase. Sometimes, but not always, one letter (e.g., X) was colored blue. This meant that participants could earn money. In this way, participants learned that seeing a blue letter means earning money, so they learned to associate the color blue with high rewards. (b) In the second part of the task participants were adding up the numbers (8+5+9+3). Sometimes, one of the letters was blue, just like in the learning phase, but this time they needed to be ignored, so they were distractors. We expected that these blue letters will make people less able to concentrate on their task.

Our First Discoveries

©2018 SAGE Publications, Ltd.. All Rights Reserved.

When we had finished creating our task, we had a nice computer program that showed people all the stimuli, exactly as we wanted. So, we invited 41 participants to our laboratory, and all of them did our new task, the first experiment (Study 1).

Our findings were interesting, but a little bit mixed. On the one hand, we found that people were indeed distracted by the color they had learned to associate with rewards. So, when they saw a letter in that color, they more often made a mistake in adding up the numbers. This finding was significant and nicely in line with our hypothesis. On the other hand, we did not find this effect right away. It was only present among people who felt that adding up the numbers was difficult. Participants who felt this was easy were not bothered by the colored letters at all. These people were still very well able to add up the numbers.

Making Improvements to Our Task

Based on what we found, we thought that our task might have been too easy. After all, a large part of our participants was not really bothered by distractors; overall, they performed really well. We thought that maybe it would be possible to find an effect of the colored letters for all participants. Certainly, this would be a stronger support for our hypothesis. After all, now we did find some support, but we really had to search for it in our data.

As a next step, then, we made the task more difficult: participants had less time to add up numbers, so they had to be faster and they needed to add up more numbers than before. After these changes in our computer task, we invited 47 participants to our lab to take part in our second experiment (Study 2). This time, the findings were straightforward: in line with what we expected, we found that all participants made mistakes when they saw the color they had learned to associate with rewards. It seemed like that making the task more difficult actually improved our paradigm: we could find the effect not only in some, but, this time, in a clear majority of participants. Looking at these results was, of course, very nice and made us very enthusiastic about the project. Yet, we needed to decide how to proceed.

Try to Publish or Try to Replicate?

At this point, there were two possible ways to go either (a) try to publish these good-looking results or (b) preregister and try to directly replicate our findings. The first option sounded appealing as publishing your results at a very early stage of your career is pretty cool—it makes you feel like your work is valuable to others, which works as a positive reinforcer that motivates you to produce even nicer findings. However, what if these results would not be replicable? You might harm science by putting something out there, which could confuse other researchers in the future.

©2018 SAGE Publications, Ltd.. All Rights Reserved.

Doing a Preregistration

Because we did not want to harm science, and because we were simply very curious about whether our findings were replicable, we decided to directly replicate our study. To be transparent, we also preregistered our next experiment. Also, we hoped that the preregistration and replication attempt would help the publication process; we expected that other people, including journal editors and reviewers, would appreciate our solid science approach.

So, we preregistered our study before data collection. We did this on a website called Open Science Framework. On this website, researchers can make their materials openly accessible to everyone (you can check out my preregistration by clicking on https://osf.io/y74kx/). In the preregistration document, we outlined

1. A short description of the study,
2. The hypotheses,
3. The design (independent and dependent variables),
4. The planned sample,
5. The computer task, and
6. The analysis plan.

Then, we uploaded and "froze" this document on the website before we started to collect data. This "freezing" is to make sure that no information in the document is changed during data collection and after, so that others can make sure that we did everything according to what we had planned before doing the study.

Trying to Replicate Our Finding

Then, we invited 93 participants in our lab to take part in the third experiment (Study 3). This time, however, the results were not in line with what we expected: people were not bothered by those colors that they associated with rewards. This was very surprising, as we did not change anything in the methods and the findings were very strong in the previous study. It was also very disappointing to see these results, as it raised a lot of questions: Is my research reliable and important? Am I going to be able to publish inconsistent results? How much this situation will slow me down in my progress? All in all, this was quite a stressful situation to deal with.

Dealing With a Non-Replication

The biggest problem was that we did not know which results to believe in. Should we trust the second experiment, which showed that people got distracted by colors that they associated rewards with? Or should we trust the third experiment that did not show this effect? In other

©2018 SAGE Publications, Ltd.. All Rights Reserved.

words, we still did not really know if our idea made sense or not.

To answer this question, we decided to try to understand our results better by using an alternative approach on our existing data from two experiments (Dienes, 2014). This approach revealed that in our data there is indeed evidence for our hypothesis that people get more distracted by things that carry some reward to them (colors associated with rewards). This gave us some more confidence about our idea, so, despite our inconsistent results, we still decided to try to publish it. We did anticipate that this might be difficult, as it is still the case that scientific journals currently prefer strong, conclusive, and significant findings. Yet, we still hope that the theoretical foundations of our ideas, and the quality and transparency of our methods will help us in publishing our work.

Summary

- We could have tried to publish our research after we found the first significant results,
- However, we chose to do research in a more transparent way (i.e., to use preregistration and direct replication) because we think that this is the correct way of doing science.
- Benefits: we did our research in a thorough way; we contribute to the advance of open science culture.
- Costs: It took more time to do the study; the paper will be more difficult to publish.

Case Study 2: Dealing With Inconsistent Results

Distractions in Real Life

As you have seen it in Case Study 1, I am generally interested in how people get distracted by things that carry some sort of reward to them. At one point during my project, I became curious whether this reward-driven distraction also exists in real-life situations, with real life things that are rewarding, such as food. So, I started another project, in which I wanted to find out how people get distracted by food that are rewarding, such as a piece of strawberry cheesecake, and food that are less rewarding, such as Brussels sprouts. We expected that when people need to do some difficult task, they would be more bothered by a picture of a piece of delicious food popping up on the screen (e.g., a piece of cheesecake)—simply because it carries more value to them.

Creating a Computer Task With Food Distractions

To test our idea, we designed a new computer task (see Figure 2). In the first part of the task, participants saw pictures of different sorts of food on the screen (e.g., a cheesecake, some Brussels sprouts, etc.). Their task was to tell how much they liked each food items. The

©2018 SAGE Publications, Ltd.. All Rights Reserved.

computer program selected the most delicious and the least delicious items, based on the participant's own ratings. In the second part of the task, participants saw a matrix of letters and numbers on the screen. Their task was to read this matrix and indicate for each character whether it was a number or a letter. To be successful in this task, participants had to continuously pay attention, as it was easy for them to lose track of where they were (see Figure 2, bottom). At random moments, a food picture appeared next to the matrix. We expected that people would perform worse—because of a lapse of attention—more often when this food item was delicious (i.e., cheesecake) compared to when it was not (e.g., Brussels sprouts).

©2018 SAGE Publications, Ltd.. All Rights Reserved.

Figure 2. Our computer task to measure distraction by food cues.

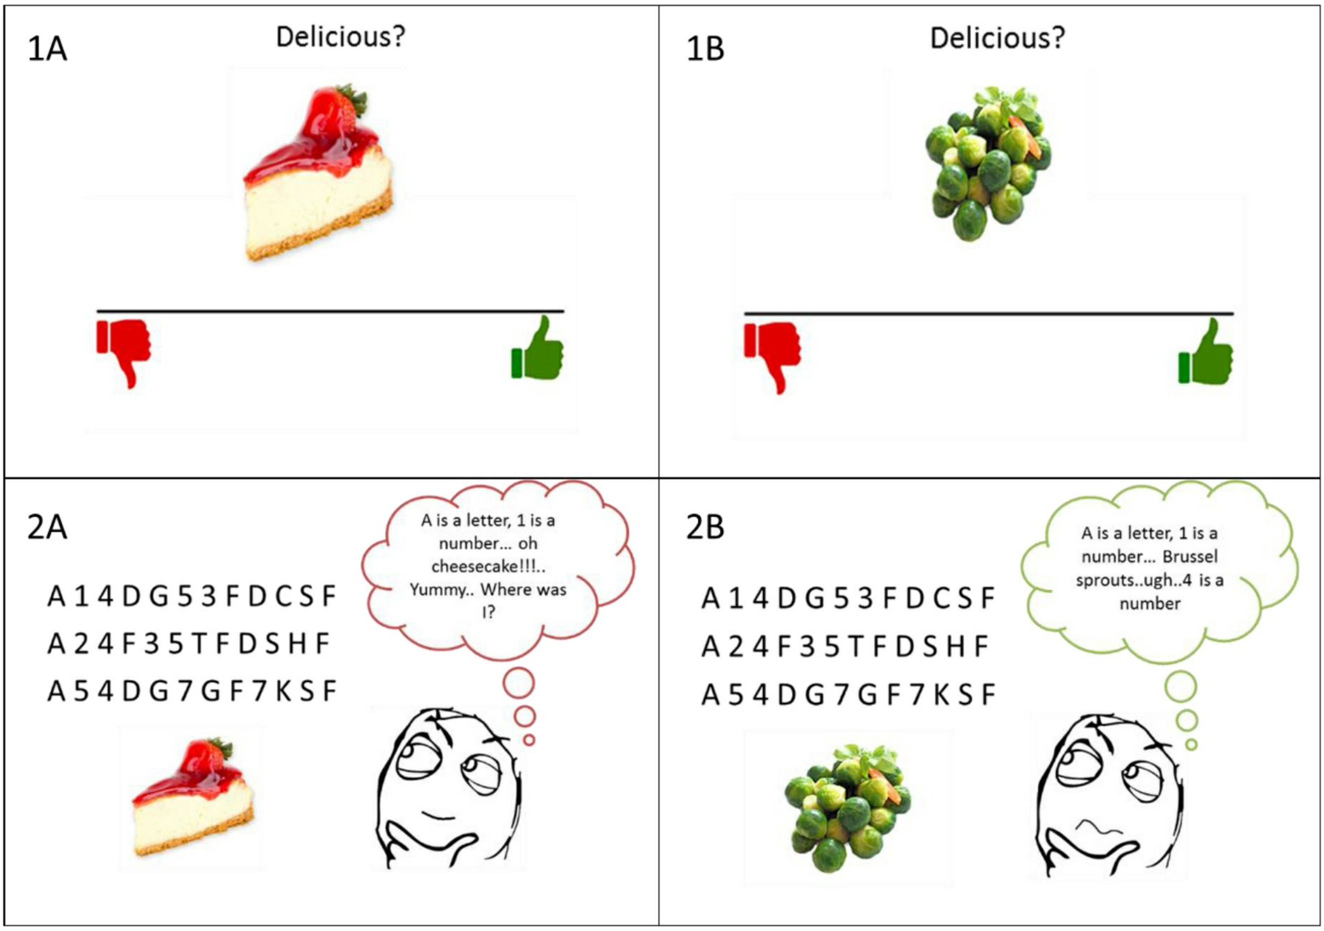

Note: In the first part of the task, our participants indicated whether they find different food items, such as a cheesecake (a) or Brussel sprouts (b) delicious. They could do it by clicking on the visual analog scale under the food items. In the second part, their task was to go through the matrix of letters and numbers and press key 1 when they see a number and press key 2 when they see a letter. Sometimes, in random moments, the previously rated food items appeared below the matrix. Sometimes, this food item was the most liked item (e.g., (c): cheesecake), sometimes the least liked item (e.g., (d): Brussel sprouts). We expected that participants will make more mistakes and respond slower when they see a cheesecake than when they see Brussel sprouts.

Three Experiments—Confusing Results

First, we wanted to try out our new task, so we invited 21 people to our lab (Study 1). The results looked quite promising: in line with our expectations, we found that people made more mistakes on the task (i.e., they lost focus) when they saw something delicious appearing on the screen. This was very motivating—but because we had only a few participants, we did not want to publish these results directly. First, we wanted to see whether we could find the same results

©2018 SAGE Publications, Ltd.. All Rights Reserved.

again. So we invited 51 participants to the lab and ran the same study again (Study 2).

This time, surprisingly, we found the exact opposite: people actually performed *better* (made less mistakes) when they saw a delicious food item appearing on the screen. These results were unexpected, but could have some plausible explanation. For instance, it could be that these delicious food items—even though they were irrelevant to the task—made people more motivated to perform well. In sum, so far, two experiments showed contradicting results; nevertheless, the second experiment had more participants, so we trusted these results better than the results from the pilot study.

Again, there were two possible ways to go: (a) publish our findings directly or (b) try to replicate the results of the second study. The first option sounded tempting, as the results from the second experiment were somewhat new: to our knowledge, not so many previous studies had shown that task-irrelevant stimuli can make people perform better. However, we knew that the possibility existed that our results would not be replicable. So, publishing them like that could, in the future, confuse other researchers, and therefore, harm scientific progress. So, before publishing, we wanted to see if we could really trust our results, so we decided to go for the more solid option: trying to replicate Study 2.

As a next step, we preregistered a direct replication of the second experiment. Similar to Case Study 1, we registered our third experiment at the Open Science Framework before data collection. Then, we invited another 64 people to participate in our task (Study 3). The results, again, were quite shocking: this time, participants were not bothered by the delicious food items at all; they always performed well on the task. This was unexpected and disappointing: all three experiments showed different results, which was basically impossible to make sense of. So, the situation was worse than in Case Study 1, in which at least we found some consistency in our results across experiment (Studies 1 and 2 showed similar results). Even though we worked in a transparent and open way, this research project was not going well.

Solution

Making a decision was quite difficult in a personal way, as we had invested so much time and effort in this project. On the professional side, however, it was quite easy to make a decision. We could not make sense of all these inconsistent results, and we did not learn a lot from these studies. Also, we knew that it would be difficult, perhaps impossible, to publish such inconsistent results. So, in the end, we decided to leave this project behind.

Summary

©2018 SAGE Publications, Ltd.. All Rights Reserved.

- We could have tried to publish our results after each experiment. After all, Study 1 was in line with the expectations and Study 2 showed a finding that was rather novel.
- We chose to do solid science instead of trying to publish premature findings.
- Benefits: we prevented non-replicable findings from entering the scientific literature;
- Costs: we spent time and effort with nothing to show for it.

Take Home Messages

A Good Deed for Science

There are a lot of advantages of registering your research plans, starting with the most important one: you do something good for science (Markowetz, 2015). Despite that we had limited success so far, we still believe that working in a transparent way can contribute to a good collective outcome. For instance, we experienced firsthand that even when you see p = .001, you cannot be sure to replicate your findings (Case Study 1). This is disappointing and annoying. It feels like wasting a lot of effort. Despite this, we still feel that it was worth it: science benefits from these replication attempts regardless of the outcome (Munafò et al., 2017). If the replication is successful, it means that your results might be trustworthy. If the replication is unsuccessful, you make sure that your results are not trustworthy, which prevents the literature to become distorted by false positive findings. Either way, you are contributing to good science. After all, we all would like to base our research ideas on previous findings that we can truly believe in.

Ease Your Work

On the more self-serving part, preregistering our plans saved some time later, when we were writing up the results. Also, we expect that it will save some time in the future, when we wish to follow up on our past experiments. Finally, when editors, and reviewers, and other readers study our case, they might gain more insight about the whole research process and line of reasoning by looking at our preregistration materials.

No Strings Attached

Some people believe that when they pre-register their analysis plan, they are no longer allowed to explore their data in other (non-preregistered) ways. In our view, this is a misunderstanding. Of course, you can do every analysis that you think makes sense, as long as you report in your paper which analysis was pre-registered, and which one was not. So, we have never felt that doing pre-registration put us in chains.

Solid Science and Personal Reality

©2018 SAGE Publications, Ltd.. All Rights Reserved.

Researchers who promote an open science culture often say that working in a transparent way is good for you and your reputation. However, as you have seen in Case Studies 1 and 2, this does not seem to be always true. Sometimes, the decision to work in a transparent way leads to unfortunate outcomes. For instance, it took a long time to submit our first paper for publication (Case Study 1)—also, this case contains inconsistent findings, which is always difficult to sell to journal editors and reviewers. These can be problematic for my future career: ideally, as a researcher, you would like to publish a couple of articles at an early stage in your career (during your PhD that usually takes 3-5 years), which will help you to get a good job in academia. However, it seems like that the decision to focus on transparent working rather than producing significant results might have added some extra level of difficulty for me to climb up on the academic career ladder.

Yet, we are very optimistic and motivated to work in a transparent way. Although the older generation still remains unaware or skeptical (Bishop, 2017; Simmons, Nelson, & Simonsohn, 2011)—perhaps, this is because they already have steady jobs—more and more (especially younger) researchers started to embrace open science culture. For example, some scientific journals have started publish studies that have transparent methods, regardless of whether the results are new or significant (Munafò et al., 2017). If more people adopt the same attitude, a new generation of scientists can create a culture in which solid science is the norm, rather than the exception.

Conclusion

This research method case shows that working in a transparent way sometimes seems to hinder your progress in science; for instance, one can spend a lot of time on doing replications that in the end turn out to be unsuccessful. Also, it can be demotivating if you put a lot of effort in something that has no tangible result. Nonetheless, we still believe that this is the correct way to go and retrospectively we are happy with all decisions we made throughout the research process. We think that if more and more people do preregistration and replication that will eventually lead to a better science.

Exercises and Discussion Questions

1. Explain why direct replications are important to science.
2. How can preregistration help science?
3. Discuss the potential drawbacks of preregistration and replications.
4. Imagine that you analyze your data for your bachelor thesis, which you are about to write. You have not found support for your hypotheses, so you explore the data out of curiosity. In

©2018 SAGE Publications, Ltd.. All Rights Reserved.

this data exploration, you find that if you analyze only female participants, your hypotheses are confirmed. How do you report this in your bachelor thesis? Why?

5. Imagine the following scenario: one of your classmates tells you that he investigated how mood affects mathematical performance. He expected that negative mood leads to worse performance. He found support for his hypotheses *only* when he excluded participants who were tired on the day of the experiment. His supervisor tells him that he could publish these significant findings in a scientific journal. What is your advice to your classmate?

Further Reading

Wagenmakers, E. J., & **Dutilh, G.** (2016). Seven selfish reasons for preregistration. *APS Observer*, *29*(9). Retrieved from https://www.psychologicalscience.org/observer/seven-selfish-reasons-for-preregistration#.WVYVHFN97R

References

Bishop, D. V. M. (2017). *Reproducible practices are the future for early career researchers*. Retrieved from http://deevybee.blogspot.nl/2017/05/reproducible-practices-are-future-for.html

Botvinick, M., & **Braver, T.** (2015). Motivation and cognitive control: From behavior to neural mechanism. *Annual Review of Psychology*, 66, 83–113. Retrieved from https://doi.org/10.1146/annurev-psych-010814-015044

Braver, T. S.Krug, M. K.Chiew, K. S.Kool, W.Westbrook, J. A.Clement, N. J., … Somerville, L. H. (2014). Mechanisms of motivation-cognition interaction: Challenges and opportunities. *Cognitive, Affective, & Behavioral Neuroscience*, 14, 443–472. doi:http://dx.doi.org/10.3758/s13415-014-0300-0

Caird, J. K.Johnston, K. A.Willness, C. R.Asbridge, M., & **Steel, P.** (2014). A meta-analysis of the effects of texting on driving. *Accident Analysis and Prevention*, 71, 311–318. doi:http://dx.doi.org/10.1016/j.aap.2014.06.005

Cheever, N. A.Rosen, L. D.Carrier, L. M., & **Chavez, A.** (2014). Out of sight is not out of mind: The impact of restricting wireless mobile device use on anxiety levels among low, moderate and high users. *Computers in Human Behavior*, 37, 290–297. doi:http://dx.doi.org/10.1016/j.chb.2014.05.002

Cohen, J. D.McClure, S. M., & **Yu, A. J.** (2007). Should I stay or should I go? *How the human brain manages the trade-off between exploitation and exploration. Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362, 933–942. doi:http://dx.doi.org/10.1098/rstb.2007.2098

©2018 SAGE Publications, Ltd.. All Rights Reserved.

Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5, 1–17. doi:http://dx.doi.org/10.3389/fpsyg.2014.00781

Jett, Q. R., & **George, J. M.** (2003). Work interrupted: A closer look at the role of interruptions in organizational life. *Academy of Management Review*, 28, 494–507. doi:http://dx.doi.org/10.5465/AMR.2003.10196791

Markowetz, F. (2015). Five selfish reasons to work reproducibly. *Genome Biology*, 16(1), 274. doi:http://dx.doi.org/10.1186/s13059-015-0850-7

Maxwell, S. E.Lau, M. Y., & **Howard, G. S.** (2015). Is psychology suffering from a replication crisis? *What does "failure to replicate" really mean? American Psychologist*, 70, 487–498. doi:http://dx.doi.org/10.1037/a0039400

Munafò, M. R.Nosek, B. A.Bishop, D. V. M.Button, K. S.Chambers, C. D.Percie, N., … Wagenmakers, E. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1, 1–9. doi:http://dx.doi.org/10.1038/s41562-016-0021

Nosek, B. A.Alter, G.Banks, G. C.Borsboom, D.Bowman, S. D.Breckler, S. J., … Yarkoni, T. (2015). Promoting an open research culture. *Science*, 348, 1422–1425. doi:http://dx.doi.org/10.1126/science.aab2374

Nuzzo, R. (2015). How scientists fool themselves—and how they can stop. *Nature*, 526, 182–185. doi:http://dx.doi.org/10.1038/526182a

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349, aac4716. doi:http://dx.doi.org/10.1126/science.aac4716

Simmons, J. P.Nelson, L. D., & **Simonsohn, U.** (2011). False-positive psychology. *Psychological Science*, 22, 1359–1366. doi:http://dx.doi.org/10.1177/0956797611417632

Theeuwes, J. (2004). Top-down search strategies cannot override attentional capture. *Psychonomic Bulletin & Review*, 11, 65–70. doi:http://dx.doi.org/10.3758/BF03206462