Experimentation and Manipulation with Preregistration

MIKE FELGENHAUER*

Brunel University London

September 20, 2021

Abstract

Preregistration requires scientists to describe the planned research activities before their project begins. Preregistration improves transparency in empirical research and is an institutional response to scientific misconduct. This paper studies the impact of a preregistration requirement in a model in which a sender can generate information for a receiver by running private experiments. The sender can also engage in uninformative manipulation. This paper argues that a preregistration requirement can discourage p-hacking, but also result in even more detrimental faked studies.

Keywords: Experimentation, persuasion, manipulation, information acquisition.

JEL classification: D82, D83

https://doi.org/10.1016/j.geb.2021.09.002

*Correspondence address: Brunel University London, Marie Jahoda, Uxbridge, UB8 3PH, UK; email: mike.felgenhauer@brunel.ac.uk. Declarations of interest: none.

1 Introduction

Replication studies suggest that scientific misconduct plays an important role in academic research (Open Science Collaboration 2015). Many papers, such as Glaeser (2008), identify different forms of scientific misconduct (such as data mining, data cleaning, etc.) and discuss institutional responses to each of them. A prominent response is improved transparency via preregistration. Preregistration requires scientists to describe the planned research activities before the project begins. This paper shows that preregistration affects the interaction of different forms of scientific misconduct. Such a transparency requirement discourages p-hacking, which is good, but also results in more faked studies, which is bad.

This paper compares academic publishing with and without preregistration in a simple model. Without preregistration, a sender (researcher) can sequentially run any number of costly private informative experiments and selectively reveal the outcomes to persuade a receiver (editor) to choose an action in his favor (publication). He can also produce an uninformative manipulated outcome (for example, by faking data) at a cost, where manipulation is possible at each experimentation history. Manipulation costs may, for example, result from expected punishment costs. The sender prefers publication to no publication, but less so if the results do not correctly reflect the decision-relevant state of the world (that is, if there is a false positive). With preregistration the sender can at most run one experiment and report its result truthfully, or engage in manipulation.

With each negative experimental outcome that the sender observes, the posterior deteriorates and the likelihood that a positive outcome is a false positive increases. Without preregistration running several experiments is sequentially rational if experimentation costs are sufficiently low to justify gambling for a positive outcome by experimentation instead of generating a positive outcome with certainty by manipulation. If the sender, who dislikes publishing a false positive to some extent, observes sufficiently many negative experimental outcomes, then unsuccessful stopping may be better than further experimentation or manipulation.

Preregistration encourages uninformative manipulation as follows. If this requirement is introduced, then gambling for a positive outcome is not possible after observing a negative outcome of the first experiment. Manipulation at this history may be better than unsuccessful stopping, since the likelihood of a false positive is moderate. A switch from informative experimentation with eventual unsuccessful stopping without preregistration to uninformative manipulation with preregistration may, thus, occur.

On the other hand, if the sender eventually manipulates without preregistration, then he also eventually manipulates with such a requirement. Again, a false positive is less likely after observing a single negative experimental outcome than after observing more than one such outcome. Hence, if eventual manipulation is attractive at a history that contains many negative outcomes, then it is even more attractive at the outset or after observing a single such outcome.

Preregistration, thus, may encourage a switch from experimentation with unsuccessful stopping to eventual manipulation, but the reverse does not happen. It follows that there are more researcher types that manipulate with preregistration than without this requirement, which is detrimental for the quality of academic publishing. A positive aspect of preregistration is that it eliminates excessive private experimentation. The quality of academic publishing clearly improves if a sender type does not manipulate with preregistration. From the above it follows that such a type also does not manipulate in the absence of preregistration. Whether additional manipulation or a reduction of private experimentation overall matters more for the publication quality depends on manipulation and experimentation costs, the opportunism of the sender and the quality of the experiments.

2 Literature

This paper is part of the economic literature on academic research (as, for example, Stern 2004; Aghion et al. 2008; Lewis and Ottaviani 2008; Olszewski and Sandroni 2011; Brodeur et al. 2016). Scientific misconduct, an important aspect in this literature, is commonly viewed as the result of incentives that researchers face (Glaeser 2008). Many scientists, as, for example, Ioannidis et al. (2014), advocate greater research transparency to reduce publication of false positives. Preregistration reduces a researcher's ability to cherry-pick hypotheses, data analyses or a good dataset (Coffman and Niederle 2015). However, preregistration is not uncontroversial in the scientific community (Coffman and Niederle 2015). A common criticism is that it discourages exploratory research or the use of novel research designs. Miguel et al. (2014) disagree and point out that prespecification frees exploratory analysis from being portrayed as formal hypothesis testing. Olken (2015) argues that empirical papers in economics not only investigate the result of a treatment, but also study the mechanisms underlying the treatment. The results then quickly become too complex / costly to specify in advance. Libgober (2020) shows that a receiver may prefer partial transparency (where only some aspects of an experiment are preregistered) to full transparency, as the sender may have to compensate in a dimension that is observable in order to credibly convey that he has chosen a not that detrimental action in the unobservable dimension. The current paper formally derives benefits (a reduction of excessive private experimentation) and downsides (additional manipulation) of preregistration.

In this paper the information that is used for persuasion can be acquired by experimentation. Therefore, this paper is part of the endogenous persuasion literature (as, for example, Kamenica and Gentzkow 2011; Henry and Ottaviani 2019).[1] It belongs to a branch of this literature in which experimentation occurs in private (as in Bro-

[1] The paper relates to strategic experimentation in "bandit problems" (see Bergemann and Välimäki 2008, for a survey). Experimentation is also investigated in the literature on the classical problem of sequential analysis (as in Wald 1947; Moscarini and Smith 2001).

cas and Carillo 2007; Henry 2009). Private experimentation with selective information disclosure is a natural assumption for many scientific methods, such as psychological experiments or regressions on a private database. Some papers, as the present one, assume that the sender's decision to continue experimenting is history-dependent (for example, Celik 2003; Felgenhauer and Schulte 2014). A consequence of this assumption is that the receiver in general cannot deduce the actual number of experiments and communication may not be fully revealing.[2] Closely related is Felgenhauer and Xu (2020), in the following FX. Their model gives the sender the additional option to manipulate at any experimentation history. They study how different face value requirements for persuasion (such as different p-value requirements for publication) affect the real value of a disclosed experimental outcome. The current paper uses the framework in FX. The contribution is to introduce preregistration and to compare the sender's behavior and welfare with and without preregistration.[3]

An implicit assumption in this paper is that the researcher has a less strong incentive to manipulate, the more strongly he believes that the claim he wants to publish is false. Without the dependence of the researcher's behavior on the distribution of the states (whether the claim is true or false), the level of manipulation would be the same with and without preregistration. In the literature on the cost of lying, Abeler et al. (2019) study an experiment in which each subject privately observes a state (good or bad), which can be misreported, and there is a monetary incentive to claim that the state is good, regardless of the actual state. They find that the more likely it is that the state is good, the more subjects misreport a bad state, that is, the propensity to lie depends on the distribution of states. Hence, there is empirical evidence that the incentive to misreport is distribution-dependent.

[2] There can be a commitment problem if experimentation is not history-dependent. At some experimentation history the sender may anticipate that persuasion is impossible by running the final experiments. Given that experimentation occurs in private, it is then not clear why he should run further experiments. In Baliga and Ely (2016) instead, the receiver's decisions are history-dependent.

[3] Without preregistration, the optimal face value requirement for a particular sender type in general cannot ensure that all types run at most one experiment.

Finally, uninformative manipulation (such as faking data) is related to the cheap talk literature (for example, Crawford and Sobel 1982). Strulovici (2017) investigates the effects of compensation schemes on experimentation and manipulation. The current paper instead assumes an informal relationship between the sender and the receiver, and, hence, no such compensation scheme. There are papers in the economic literature on academic research that attach a different meaning to the term "manipulation". For example, Di Tillio, Ottaviani and Sørenson (2017) study manipulation and selective disclosure in a randomized controlled trial (RCT). Manipulation occurs via selective sampling and selective assignment, such that a manipulated RCT is still informative.

3 Model

The model describes a scenario where a researcher can publish a paper that makes an interesting claim, but not the opposite of this claim. The claim can be true or false, which cannot be directly observed. The researcher needs a supporting argument for a publication. Arguments can be generated in private by running informative experiments or by uninformative manipulation and they can be selectively revealed.

The assumptions are as follows. There is a sender (researcher) and a receiver (editor). The receiver chooses $a \in \{a_1, a_2\}$, where a_1 is "publication" and a_2 is "rejection". There is an unknown state $\omega \in \{\omega_1, \omega_2\}$, where ω_1 means that the researcher's claim is "true" and ω_2 means that it is "false". The prior belief is $prob\{\omega = \omega_1\} = \mu_0$, with $\mu_0 \in (0, 1/2]$. The claim is more likely to be false, and, hence, "surprising".

Payoffs The sender prefers a_1 in each state, but less in state ω_2 (published claim is false) than in state ω_1 (published claim is true).[4] Sender type $\theta \in [0, 1]$, where θ reflects how much he dislikes to publish a false claim, obtains gross utility:[5]

[4] In Herresthal (2017), who compares private and public experimentation without the option to manipulate, the utility from choosing a particular action also depends on the state.

[5] The sender's gross utility is zero if he never stops experimenting. Experimentation and manipulation costs are subtracted from the sender's gross utility.

	$\omega = \omega_1$	$\omega = \omega_2$
$a = a_1$	1	θ
$a = a_2$	0	0

The receiver obtains utility 1 if she chooses a_1 in state ω_1 (published claim is true) and -1 if she chooses a_1 in state ω_2 (published claim is a false positive). Otherwise, her utility is 0. In the following, the receiver's ex ante utility is called the "publication quality".

Experimentation and manipulation The sender can run any number of experiments without preregistration and at most one experiment with this requirement. Experiment j's outcome is $\sigma_j \in \{s_1, s_2\}$, where s_1 is called a "positive" outcome and s_2 a "negative" outcome. An outcome correctly reflects the state with exogenous probability $\pi \in (1/2, 1]$. That is, outcome s_i realizes in state ω_i with probability π. The probability that outcome s_i realizes in state ω_j, with $i \neq j$, is $1 - \pi$. The sender privately observes the experimentation history $h_t = \{\sigma_j\}_{j=1,\ldots,t}$. The history at the outset is $h_0 = \emptyset$. Denote by μ_t the posterior probability $prob\{\omega = \omega_1 \mid h_t\}$ if h_t does not contain positive outcomes. Each experiment costs $c_E > 0$. Let $c_E \leq \mu_1\pi + (1 - \mu_1)(1 - \pi)\theta$. As will be seen later, at higher costs the sender does not want to continue experimenting after observing a single negative outcome and preregistration is not binding. The sender may also manipulate in private at costs $c_M > 0$. Manipulation yields an uninformative outcome s_1.

Sender's message The sender sends a feasible message $m \in \{s_1, s_2, \emptyset\}$. Message $m = s_i$ is feasible if s_i stems from experimentation or manipulation.

Receiver's decision rule The receiver chooses a_1 if $m = s_1$ and a_2, otherwise.[6]

[6]Endogenizing the receiver's behavior, as in Felgenhauer and Schulte (2014) and Felgenhauer and Loerke (2017), complicates the analysis without substantially changing the results. Note, though, that "rejection" after observing a positive outcome would yield a higher publication quality if a sender type eventually manipulates or if the quality of the experiments is very low.

Timing Without preregistration the sender makes the history-dependent choice to conduct a further experiment or to stop experimenting at each experimentation history that he observes. After stopping he decides whether to manipulate. He then sends a message m to the receiver, who finally chooses a. The timing with preregistration is the same, but the sender may run at most one experiment.

4 Analysis

4.1 Experimentation, Manipulation and the Publication Quality Without Preregistration

It is sequentially rational for the sender to stop experimenting at a history h_t that contains a positive outcome.[7] He then obtains his preferred decision a_1 by sending message $m = s_1$. It may be sequentially rational to continue experimenting or to manipulate if h_t does not contain a positive outcome.

Suppose manipulation is not possible (or manipulation costs are sufficiently high such that it is not sequentially rational for the sender to manipulate). The sender's payoff from running one further experiment at some history h_t that does not contain a positive outcome and then stopping after either outcome is

$$EU_t^{E1} = \mu_t(\pi - c_E) + (1 - \mu_t)((1 - \pi)\theta - c_E). \tag{1}$$

Term $(\pi - c_E)$ is the expected payoff in state ω_1, which realizes with probability μ_t. In this term, π is the probability that the experiment yields a positive outcome, implying gross utility 1. Experimentation costs c_E have to be subtracted (in each state). Term $(1 - \mu_t)((1 - \pi)\theta - c_E)$ refers to state ω_2 and is interpreted analogously.

It is not sequentially rational to continue experimenting if $EU_t^{E1} < 0$. Suppose there

[7] The analysis in this subsection directly follows FX.

is a finite t at which this inequality holds and denote by T_E the lowest t at which this inequality holds. The sender's continuation utility at some history h_t, with $t < T_E$, is denoted by EU_t^E.[8] If $EU_t^{E1} > 0$ at all posteriors (in particular at the "worst" posterior $\mu = 0$), then the sender never stops experimenting unsuccessfully.

Now suppose that manipulation is also possible. Manipulation yields a positive outcome in each state and costs c_M. As the sender's gross payoff from a_1 is 1 in state ω_1 and θ in state ω_2, his continuation utility from manipulation at t is

$$EU_t^M = \mu_t(1 - c_M) + (1 - \mu_t)(\theta - c_M). \tag{2}$$

It is sequentially rational for the sender to experiment with eventual unsuccessful stopping at T_E (without any manipulation) if $EU_t^E \geq EU_t^M$ for all $t < T_E$ and $EU_{T_E}^M \leq 0$. Otherwise, eventual manipulation occurs at some T_M, with $T_M \leq T_E$. Both, EU_t^E and EU_t^M decrease in the number of negative outcomes that history h_t contains, that is, they decrease if μ_t decreases.

The sender's behavior affects the publication quality, which matters for the welfare analysis below. If the sender eventually manipulates or never stops experimenting unsuccessfully, then the publication quality is $\mu_0 - (1 - \mu_0)$.[9] Suppose the sender runs at most T experiments without manipulation. Action a_1 is chosen if not all of these experiments' outcomes are negative. This happens with probability $(1 - (1 - \pi)^T)$ in state ω_1 and with probability $(1 - \pi^T)$ in state ω_2, yielding receiver utilities 1 and -1 in the respective states. Thus, the publication quality is

$$\mu_0(1 - (1 - \pi)^T) - (1 - \mu_0)(1 - \pi^T). \tag{3}$$

Lemma 1 *(i) Suppose the sender runs at most $T \geq 2$ experiments without manipu-*

[8] A formula for EU_t^E is provided in the appendix.

[9] The receiver chooses a_1 in each state. Action a_1 is only correct in state ω_1 (yielding receiver utility 1), which ex ante realizes with probability μ_0. This action yields utility -1 in the other state, which realizes with probability $1 - \mu_0$.

lation. A decrease of T by one increases the publication quality. (ii) The publication quality from experimentation with unsuccessful stopping is higher than from manipulation or experimentation without unsuccessful stopping.

Thus, less excessive private experimentation improves the publication quality (part (i)). The quality of a publication is higher if it is based on an informative revealed outcome than in case it is based on an uninformative revealed outcome (part (ii)).

4.2 Comparison of the Sender's Behavior With and Without Preregistration

Suppose that without preregistration the sender eventually manipulates at some $T_M > 1$. We have $EU_{T_M}^M > 0$, as otherwise it is not sequentially rational to manipulate at T_M. The sender then also eventually manipulates with preregistration either at the outset or after observing a negative outcome of the first experiment, as $EU_0^M > EU_1^M > EU_{T_M}^M$ (due to $\mu_0 > \mu_1 > \mu_{T_M}$) and as the continuation utility from stopping unsuccessfully at $t = 1$ is 0.

Suppose that without preregistration the sender experiments and does not manipulate. With such a requirement, depending on the size of EU_0^{E1}, EU_0^M and EU_1^M, he may now either run a single experiment and stop after either outcome without manipulation, or he may manipulate either at the outset or after observing a negative outcome of the first experiment.

Proposition 1 is a direct consequence of these effects.

Proposition 1 *Any sender type that eventually manipulates without preregistration also eventually manipulates with such a requirement. A sender type that does not eventually manipulate without preregistration (i) does not manipulate with such a requirement if* $EU_0^{E1} \geq EU_0^M$ $(\Leftrightarrow \mu_0(\pi - c_E) + (1-\mu_0)((1-\pi)\theta - c_E) \geq \mu_0(1 - c_M) + (1 - \mu_0)(\theta - c_M))$ *and* $EU_1^M \leq 0$ $(\Leftrightarrow \mu_1(1 - c_M) + (1 - \mu_1)(\theta - c_M) \leq 0)$ *(ii) eventually manipulates with preregistration if* $EU_0^{E1} < EU_0^M$ *or if* $EU_1^M > 0$.

As an illustration that the introduction of preregistration may induce a switch from experimenting with unsuccessful stopping to eventual manipulation, consider parameters ($\mu_0 = 0.5, \pi = 0.8$, $c_E = 0.15$, $c_M = 0.75$, $\theta = 0.7$). For these parameters the sender's benefit from manipulation after observing a negative outcome of the first experiment is greater than zero. However, experimentation costs are sufficiently low compared to manipulation costs to justify the gamble to obtain a positive outcome by chance via further experimentation. With each negative experimental outcome the posterior is depressed, which also makes manipulation less attractive and worse than stopping unsuccessfully at any $t > 1$. The sender eventually stops experimenting unsuccessfully without manipulation (at $T_E = 4$). With preregistration, further experiments are not possible after observing a negative outcome of the first experiment. Since the benefit from manipulation at this history is greater than zero, the sender switches to eventual manipulation if preregistration is introduced.

4.3 The Impact of Preregistration on Welfare

The components that matter for welfare are the publication quality, the ex ante probability that decision a_1 is made (in the following called the "publication probability") and expected costs. Each of the three components of welfare is first studied separately and then overall welfare is discussed.

Publication quality If the sender eventually manipulates with and without preregistration, then the publication quality in both cases is the same. If the introduction of such a requirement encourages a switch from experimenting with eventual unsuccessful stopping to manipulation, then the publication quality strictly deteriorates according to Lemma 1 (ii).[10] For example, for the above parameters ($\mu_0 = 0.5, \pi = 0.8$, $c_E = 0.15$, $c_M = 0.75$, $\theta = 0.7$), introducing preregistration causes a drop of the publication quality from 0.204 to 0 due to the switch from experimentation to manipulation. If there

[10] If there is a switch from experimenting without unsuccessful stopping to eventual manipulation, then the publication quality is the same.

is experimentation (without eventual manipulation) with and without preregistration, then the publication quality improves, as there is a reduction of excessive private experimentation according to Lemma 1 (i). The impact of preregistration on the publication quality depends on how opportunistic the sender type is.

Corollary 1 *There is a sender type* $\widetilde{\theta} = \min\{\frac{c_M - c_E - (1-\pi)\mu_0}{\pi(1-\mu_0)}, \frac{c_M - \mu_1}{1-\mu_1}\}$ *such that with preregistration the publication quality is (i) higher than without such a requirement if* $\theta \leq \widetilde{\theta}$ *and (ii) weakly lower than without such a requirement, otherwise.*

The critical sender type $\widetilde{\theta}$ follows from solving $EU_0^{E1} = EU_0^M$ and $EU_1^M = 0$ for θ (see Proposition 1). According to the corollary, preregistration improves the publication quality if the sender is sufficiently sincere and it weakly deteriorates the publication quality if the sender is sufficiently opportunistic. For a sincere type preregistration reduces detrimental excessive private experimentation and does not encourage manipulation. An opportunistic type eventually manipulates with preregistration, but not necessarily without this requirement.[11]

Publication probability The publication probability is one if the sender manipulates with and without preregistration. If the sender does not manipulate in both cases, then preregistration reduces the publication probability, as it reduces excessive experimentation and thereby the chance to generate a positive outcome. If the introduction of preregistration instead induces a switch from experimentation with unsuccessful stopping to eventual manipulation, then the publication probability increases with preregistration. In the former case the publication probability is smaller than one, but in the latter case the sender always produces a positive outcome.

Expected costs Preregistration is a constraint for the sender. With preregistration

[11] An increase of $\widetilde{\theta}$ due to a parameter change implies that more types become "sincere". $\widetilde{\theta}$ increases in c_M, as experimentation becomes more attractive. Similarly, $\widetilde{\theta}$ decreases in c_E. With a higher π it is ex ante less likely to find a positive outcome by experimentation if the prior μ_0 is low (gambling by experimentation is less profitable) and more likely if μ_0 is high (gambling is more profitable). Hence, whether an increase of π increases or decreases $\widetilde{\theta}$ depends on the size of μ_0 (see appendix).

he is weakly worse off than without such a requirement. Suppose the sender manipulates with and without preregistration. As the publication probability is the same and the sender is weakly worse off with preregistration, expected costs with preregistration have to be weakly higher than without preregistration. With preregistration expected experimentation costs decrease and expected manipulation costs increase. If the sender does not manipulate with and without preregistration, then expected experimentation costs are lower with this requirement, as it prevents excessive experimentation. If the sender switches from experimentation to manipulation in response to preregistration, then again, the publication probability increases. As the sender is weakly worse off with the preregistration constraint, this implies that expected costs have to increase as well in order to overcompensate the higher publication probability (where expected experimentation costs decrease and expected manipulation costs increase).

The following proposition summarizes these effects.

Proposition 2 *If a sender type eventually manipulates without preregistration, then the introduction of such a requirement does not change the publication quality and the publication probability, but it increases expected costs. If a sender type experiments without preregistration, then the introduction of such a requirement improves the publication quality, but it lowers the publication probability and expected costs under the conditions of Proposition 1 (i),* $EU_0^{E1} \geq EU_0^M$ *and* $EU_1^M \leq 0$*. If a sender type experiments without preregistration, then the introduction of such a requirement lowers the publication quality, but it increases the publication probability and expected costs under the conditions of Proposition 1 (ii),* $EU_0^{E1} < EU_0^M$ *or* $EU_1^M > 0$*.*

Let welfare be a linear combination of the above components, with λ_R being the weight attached to the receiver's ex ante utility (the publication quality), λ_P being the weight for the publication probability and λ_C being the weight for expected costs, with $\lambda_R + \lambda_P + \lambda_C = 1$.

For any weights, welfare is weakly lower with preregistration if the sender manipulates without preregistration. The reason is that, according to Proposition 2, the publication quality and the publication probability are not affected by preregistration, but expected costs are higher with preregistration. Suppose in the following that the sender does not manipulate without preregistration.

Consider first the case where the sender and the receiver have the same weight in the welfare function ($\lambda_R = \lambda_P = \lambda_C$). If the sender manipulates with preregistration, then both, the sender and the receiver are worse off with this requirement: the sender, as preregistration is a constraint, and the receiver, as manipulation yields the lowest publication quality. Suppose the sender does not manipulate with preregistration. The receiver prefers preregistration, as it prevents excessive private experimentation, but the sender dislikes this constraint. The size of this conflict of interest depends on the sender's opportunism: the more opportunistic the sender is, the more excessively he is willing to experiment without preregistration (which he only does because it gives him additional utility), but which is worse for the receiver, as more excessive private experimentation depresses the publication quality. The receiver gains more from preregistration, the more opportunistic the sender is, but the sender then loses more with this constraint.[12] In the appendix it is shown that with sufficiently low sender opportunism, welfare is maximized with preregistration, but with high opportunism welfare can be maximized without this requirement. This point also depends on the specification of the receiver's preferences. For example, if the editor's disutility from publishing a false positive is sufficiently high instead, then welfare is maximized with preregistration even if the researcher is completely opportunistic.

The case where the sender is sufficiently small such that the impact of his payoff on welfare is negligible can be approximated by $\lambda_R = 1$. The effects of preregistration on welfare then directly follow from Proposition 2. An extreme case is $\lambda_R = 0$ and $\lambda_P = \lambda_C$, where welfare corresponds to the sender's payoff. Introducing preregistration

[12]In addition, the sender's gross utility per se increases in θ and, thus, by construction, his relevance in the welfare function.

lowers welfare, as the sender is weakly worse off with this constraint. If welfare is some convex combination of the sender's and the receiver's payoff, with $\lambda_R \in (0,1)$ and $\lambda_P = \lambda_C$, and the sender does not manipulate with preregistration, then the effect of this requirement on welfare clearly depends on the size of λ_R, due to the conflict of interest between the sender and the receiver. If the sender instead manipulates with preregistration, then, analogous to above, the sender and the receiver both prefer no preregistration. In the context of this model's interpretation, the sender's gross payoff increases in the probability of a publication, regardless of whether it is true or false. It may be viewed as reasonable to attach a low weight to the sender's gross payoff. Let us approximate this case by setting $\lambda_P = 0$. Now only the publication quality and expected costs matter for welfare.

Proposition 3 *Let $\lambda_P = 0$. If a sender type experiments without preregistration, then welfare is higher with such a requirement under the conditions of Proposition 1 (i), $EU_0^{E1} \geq EU_0^M$ and $EU_1^M \leq 0$, and lower under the conditions of Proposition 1 (ii), $EU_0^{E1} < EU_0^M$ or $EU_1^M > 0$.*

Under the conditions of Proposition 1 (i) the sender also experiments with preregistration, which increases the publication quality and reduces expected costs according to Proposition 2, and, hence, preregistration improves welfare if $\lambda_P = 0$. Under the conditions of Proposition 1 (ii) the sender manipulates with preregistration, which decreases the publication quality and increases expected costs according to Proposition 2, and, hence, preregistration deteriorates welfare if $\lambda_P = 0$. Preregistration only improves welfare if the sender does not manipulate with this requirement. According to Proposition 1 (i), this is the case if manipulation costs are sufficiently high such that manipulation after observing a negative outcome of the first experiment is worse for the sender than stopping unsuccessfully ($EU_1^M \leq 0$) and if experimentation costs are sufficiently low such that running a single experiment (gambling for a positive outcome) is better than manipulating at the outset (positive outcome with certainty), that is, if

$EU_0^{E1} \geq EU_0^M$.

5 Conclusion

This paper studies benefits and downsides of preregistration in a simple model in which decision-relevant information can be acquired via sequential private experimentation and where in addition uninformative manipulation may occur. Introducing a preregistration requirement may cause a switch from informative experimentation to uninformative manipulation, but not vice versa. A researcher who dislikes to publish a false positive to some extent, but still eventually manipulates without preregistration at a poor posterior, also manipulates with preregistration at a better prior / posterior where a false positive is less likely. A positive aspect of preregistration is that it reduces excessive private experimentation for a sender who does not manipulate with such a requirement. Such a sender also does not manipulate without preregistration.

Whether preregistration improves the quality of the academic publishing process depends on the opportunism of the sender, the precision of the experiments and manipulation and experimentation costs. Some of these parameters are contingent on the scientific method. For example, experimentation costs may be influenced by method specific restrictions on admissible arguments (Felgenhauer and Schulte, 2014). For a given method, experimentation costs should also decrease over time due to technological progress and methods may evolve over time. The likelihood of replication studies should influence manipulation costs. The level of opportunism should depend on institutional incentives. A proposal to introduce preregistration should take these parameters into account and include measures (such as restrictions on admissible arguments) to keep them in a range where preregistration is indeed beneficial. Actions to circumvent preregistration, such as pre-trials, should also be made expensive.

The paper is framed in the context of academic publishing. Another important application is the approval process for new drugs in the pharmaceutical industry, where it is often legally required to preregister medical experiments. The effects identified here can be directly applied to the preregistration of such experiments.

A measure (such as improved transparency) to reduce a particular kind of scientific misconduct (such as p-hacking) can encourage other forms of scientific misconduct (such as faking data). Future work could further explore the interdependencies between different forms of scientific misconduct (possibly contingent on the scientific method) and appropriate institutional responses.

APPENDIX

Continuation utility from experimenting The sender's continuation utility from experimenting without manipulation at some history h_t, with $t < T_E$, that does not contain a positive outcome is

$$EU_t^E = \mu_t \sum_{n=0}^{T_E-t-1} (1-\pi)^n(\pi - c_E) + (1-\mu_t) \sum_{n=0}^{T_E-t-1} \pi^n((1-\pi)\theta - c_E), \qquad (4)$$

with $\mu_t = \frac{\mu_0(1-\pi)^t}{\mu_0(1-\pi)^t+(1-\mu_0)\pi^t}$. The continuation utility in state ω_1, which realizes with probability μ_t, is $\sum_{n=0}^{T_E-t-1}(1-\pi)^n(\pi - c_E)$. Component $\sum_{n=0}^{T_E-t-1}(1-\pi)^n\pi$ is the probability at t to find a positive outcome (yielding gross utility 1) in state ω_1 if the sender overall runs at most T_E experiments.[13] Component $\sum_{n=0}^{T_E-t-1}(1-\pi)^n c_E$ are expected experimentation costs at t if the sender overall runs at most T_E experiments in this state.[14] $\sum_{n=0}^{T_E-t-1}\pi^n((1-\pi)\theta - c_E)$ is the continuation utility in state ω_2.

Proof of Lemma 1: (i) In (3), term $(1-(1-\pi)^T)$ increases in T and $-(1-\pi^T)$ decreases in T. The critical case is $\mu_0 = 1/2$, in which case the publication quality simplifies to $\frac{1}{2}(1-(1-\pi)^T) - \frac{1}{2}(1-\pi^T) = \frac{1}{2}(\pi^T - (1-\pi)^T)$. Now it is shown that decreasing T by 1 increases the publication quality, that is $\frac{1}{2}(\pi^T - (1-\pi)^T) \leq \frac{1}{2}(\pi^{T-1} - (1-\pi)^{T-1})$, which is equivalent to $\frac{1}{2}((\pi^{T-1} + (1-\pi)^{T-1})\pi - (1-\pi)^{T-1}) \leq \frac{1}{2}(\pi^{T-1} - (1-\pi)^{T-1})$. The inequality holds if $(\pi^{T-1} + (1-\pi)^{T-1})\pi \leq \pi^{T-1} \Leftrightarrow \frac{\pi}{(1-\pi)} \leq (\frac{\pi}{(1-\pi)})^{T-1}$.

---

[13] For example, if $t = 0$ and $T_E = 2$, then $\sum_{n=0}^{T_E-t-1}(1-\pi)^n\pi = (1-\pi)^0\pi + (1-\pi)\pi$, where $(1-\pi)^0\pi = \pi$ is the probability that the first experiment yields a positive outcome in state ω_1, and $(1-\pi)\pi$ is the probability that the first experiment yields a negative outcome and the second experiment yields a positive outcome in state ω_1. After the first experiment we have $t = 1$ and $T_E = 2$, which yields $\sum_{n=0}^{T_E-t-1}(1-\pi)^n\pi = (1-\pi)^0\pi = \pi$, which is the probability that the next (and final) experiment yields a positive outcome at this history in state ω_1.

[14] For example, if $t = 0$ and $T_E = 2$, then $\sum_{n=0}^{T_E-t-1}(1-\pi)^n c_E = (1-\pi)^0 c_E + (1-\pi)c_E$, where $(1-\pi)^0 c_E = c_E$ is the cost of running the first experiment in state ω_1, and $(1-\pi)c_E$ is the probability that the first experiment yields a negative outcome (only in this case the second experiment is run) times the costs of the second experiment in state ω_1. After the first experiment we have $t = 1$ and $T_E = 2$, which yields $\sum_{n=0}^{T_E-t-1}(1-\pi)^n c_E = (1-\pi)^0 c_E = c_E$, which is the cost from running the second (and final) experiment at this history in state ω_1.

As $\pi \in (1/2, 1]$ by assumption, we have $\frac{\pi}{(1-\pi)} > 1$. Hence, $\frac{\pi}{(1-\pi)} \leq (\frac{\pi}{(1-\pi)})^{T-1}$ if $T \geq 2$. By assumption, $c_E \leq \mu_1 \pi + (1-\mu_1)(1-\pi)\theta$, which implies $T \geq 2$ if the sender does not manipulate.

(ii) It has to be shown that $\mu_0(1-(1-\pi)^T) - (1-\mu_0)(1-\pi^T) \geq \mu_0 - (1-\mu_0)$. Collecting terms yields $(\frac{\pi}{(1-\pi)})^T \geq \frac{\mu_0}{(1-\mu_0)}$, which is true as $\pi \in (1/2, 1]$ and $\mu_0 \in (0, 1/2]$. Q.E.D.

The effect of π on sender type $\widetilde{\theta}$ in Corollary 1: Let parameters be such that $\widetilde{\theta} = \min\{\frac{c_M - c_E - (1-\pi)\mu_0}{\pi(1-\mu_0)}, \frac{c_M - \mu_1}{1-\mu_1}\} \in (0, 1)$. Suppose $\frac{c_M - c_E - (1-\pi)\mu_0}{\pi(1-\mu_0)} < \frac{c_M - \mu_1}{1-\mu_1}$ and consider a marginal increase of π (the precision of the experiment) in which case this inequality still holds. The effect of π on $\widetilde{\theta}$ depends on the sign of $(\mu_0 - c_M + c_E)$, since $\frac{d}{d\pi}(\frac{c_M - c_E - (1-\pi)\mu_0}{\pi(1-\mu_0)}) = \frac{\mu_0 - c_M + c_E}{\pi^2(1-\mu_0)}$ (where experimentation may only be sequentially rational if $c_M > c_E$). An increase of π increases $\widetilde{\theta}$ if μ_0 is greater than $(c_M - c_E)$ and it decreases this type otherwise.

The effect of the sender's opportunism on welfare if $\lambda_R = \lambda_P = \lambda_C$: Suppose the sender does not manipulate with preregistration. It is now shown that with sufficiently low sender opportunism welfare is maximized with preregistration, but with high opportunism welfare can be maximized without this requirement. Let $c_E \in (0, 1-\pi)$. Assumption $c_E < 1-\pi$ guarantees that the sender experiments until he finds a positive outcome if $\theta = 1$: the sender does not stop experimenting unsuccessfully even if he knows that the state is adverse, as $\mu(\pi - c_E) + (1-\mu)((1-\pi)\theta - c_E) > 0$ simplifies to $c_E < 1-\pi$ if $\mu = 0$ and $\theta = 1$.

Consider the highest level of opportunism $\theta = 1$. The sum of the receiver's and the sender's ex ante utilities without preregistration, where the sender does not stop experimenting unsuccessfully and presents a positive outcome regardless of the state, is then equal to $\mu_0 - (1-\mu_0) + 1 - (\frac{\mu_0}{\pi} + \frac{1-\mu_0}{1-\pi})c_E = 2\mu_0 - (\frac{\mu_0}{\pi} + \frac{1-\mu_0}{1-\pi})c_E$, where $1-(\frac{\mu_0}{\pi} + \frac{1-\mu_0}{1-\pi})c_E$ follows from EU_0^E in (4) with $\theta = 1$ and $T_E \to \infty$. With preregistration it is $\mu_0(1-(1-\pi)) - (1-\mu_0)(1-\pi) + \mu_0\pi + (1-\mu_0)(1-\pi) - c_E = 2\pi\mu_0 - c_E$. We have $2\mu_0 - (\frac{\mu_0}{\pi} + \frac{1-\mu_0}{1-\pi})c_E > 2\pi\mu_0 - c_E$ if c_E is sufficiently low, as $\pi < 1$. Hence, for a high level of sender opportunism, no preregistration can maximize welfare.

Next, consider a low level of opportunism. For each c_E, there is a θ sufficiently low such that $T_E = 2$ and he is just indifferent between running the second experiment and stopping unsuccessfully after a negative outcome of the first experiment (which is the case if $\mu_1(\pi - c_E) + (1 - \mu_1)((1 - \pi)\theta - c_E) = 0$). The sender's ex ante utility with and without preregistration is, thus, the same. As the receiver is better off with preregistration, where a revealed positive outcome is based on a single experiment (and not potentially two experiments without preregistration), this requirement maximizes welfare. Hence, for a sufficiently low level of sender opportunism preregistration maximizes welfare.

REFERENCES

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond. 2019. "Preferences for Truth-Telling." *Econometrica*, 87(4): 1115-1153.

Aghion, Philippe, Mathias Dewatripont, and Jeremy C. Stein. 2008. "Academic Freedom, Private-Sector Focus, and the Process of Innovation." *RAND Journal of Economics*, 39(3): 617-635.

Baliga, Sandeep, and Jeffrey C. Ely. 2016. "Torture and the Commitment Problem." *Review of Economic Studies,* 83: 1406-1439.

Bergemann, Dirk, and Juuso Välimäki. 2008. *Bandit Problems.* The New Palgrave Dictionary of Economics, Steven Durlauf and Larry Blume, eds., 2nd ed. Macmillan Press.

Brocas, Isabelle, and Juan D. Carrillo. 2007. "Influence through Ignorance." *RAND Journal of Economics*, 38(4): 931-947.

Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg. 2016. "Star Wars: The Empirics Strike Back." *American Economic Journal: Applied Economics*, 8(1): 1-32.

Celik, Gorkem. 2003. "Interested Experts: Do They Know More?" working paper.

Coffman, Lucas C. and Muriel Niederle. 2015. "Pre-analysis Plans have Limited Upside, Especially where Replications are Feasible." *Journal of Economic Perspectives*, 29(3): 61-80.

Di Tillio, Alfredo, Marco Ottaviani and Peter Norman Sørensen. 2017. "Persuasion Bias in Science: Can Economics Help?" *The Economic Journal*, 127(605): 266-304.

Felgenhauer, Mike, and Petra Loerke. 2017. "Bayesian Persuasion with Private Experimentation." *International Economic Review,* 58(3): 829-856.

Felgenhauer, Mike and Elisabeth Schulte. 2014. "Strategic Private Experimentation." *American Economic Journal: Microeconomics*, 6: 74-105.

Felgenhauer, Mike and Fangya Xu. 2020. "The Face Value of Arguments With and Without Manipulation." *International Economic Review*, forthcoming.

Glaeser, Edward L. 2008. "Researcher Incentives and Empirical Methods." In The Foundations of Positive and Normative Economics: A Hand Book, eds. A. Caplin and A. Schotter, 300–319. New York: Oxford University Press.

Henry, Emeric. 2009. "Strategic Disclosure of Research Results: The Cost of Proving Your Honesty." *The Economic Journal*, 119: 1036-1064.

Herresthal, Claudia. 2017. "Hidden Testing and Selective Disclosure of Evidence." working paper.

Ioannidis John P. A. et al. 2014. "Increasing Value and Reducing Waste in Research Design, Conduct and Analysis." *The Lancet*, 383(9912): 166-175.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590-2615.

Lewis, Tracy R., and Marco Ottaviani. 2008. "Search Agency." working paper.

Libgober, Jonathan. 2020. "False Positives and Transparency." *American Economic Journal: Microeconomics*, forthcoming.

Miguel, Edward et al. 2014. "Promoting Transparency in Social Science Research." *Science*, 343(6166): 30-31.

Moscarini, Giuseppe, and Lones Smith. 2001. "The Optimal Level of Experimentation.", *Econometrica*, 69(6): 1629-1644.

Olken, Benjamin A. 2015. "Promises and Perils of Pre-Analysis Plans." *Journal of Economic Perspectives*, 29(3): 81-98.

Olszewski, Wojciech, and Alvaro Sandroni. 2011. "Falsifiability." *American Economic Review*, 101(2): 788-818.

Open Science Collaboration. 2015. "Estimating the Reproducibility of Psychological Science." *Science*, 349 (6251): aac4716.

Stern, Scott. 2004. "Do Scientists Pay to Be Scientists." *Management Science*, 50(6): 835-853.