Economics Letters 244 (2024) 112017

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Competition for publication-based rewards

Mike Felgenhauer[1]

Brunel University London, Marie Jahoda, Uxbridge, UB8 3PH, UK

ARTICLE INFO

JEL classification:
D82
D83

Keywords:
Experimentation
Persuasion
Manipulation
Information acquisition

ABSTRACT

This paper studies how more competition among researchers for publication-based rewards affects the quality of the publication process. Publishable results can be generated via costly informative sequential private experimentation or costly uninformative manipulation. By reducing expected rewards, competition may discourage manipulation in favor of experimentation, but not vice versa. It also reduces excessive experimentation. Both effects improve the quality of the publication process.

1. Introduction

Publication-based rewards encourage scientific misconduct by distorting researchers' incentives. But what is the effect of increased competition for such rewards on the quality of the publication process? This paper argues that more intense competitive pressure affects how much researchers engage in various forms of scientific misconduct that are not equally detrimental. The focus is on the interdependence of informative private experimentation with selective disclosure (such as p-hacking) and uninformative manipulation (such as faking data).[2] The paper finds that tougher competition reduces expected benefits from distortive behavior, but the extent differs for experimentation and manipulation. Increased competition can discourage manipulation in favor of experimentation, but not the reverse. Competition also reduces excessive experimentation. Both improves the quality of the publication process.

This paper belongs to the literature on persuasion with information acquisition (for example, Kamenica and Gentzkow (2011)). It is part of a branch in which experimental outcomes are observed in private and can be selectively revealed (as in Brocas and Carrillo, 2007; Henry, 2009). Some papers, as the present one, assume that experimentation is history-dependent, which implies that communication may not be fully revealing (as in Celik (2003)). The current paper is closely related to Felgenhauer and Xu (2021), in the following FX. In FX, a single sender can experiment and in addition engage in uninformative manipulation. The present paper compares the quality of the publication process with and without competition. Tougher competition affects the probability to obtain a reward and, thereby, researchers' behavior.[3]

2. Model

Without competition, one sender (researcher) can engage in costly sequential private experimentation with selective disclosure or manipulation in order to find an argument that supports an interesting claim, which can be used for persuading a receiver (editor) to choose a favorable action (publication). The sender may only obtain a reward (a job in academia) upon publication.[4] Competition is introduced by adding a second sender, who works on a different claim. There is a publication slot for each sender, but only one reward.

For each sender $i = 1, 2$ there is a state of the world $\omega_i \in \{\omega_{i1}, \omega_{i2}\}$, where ω_{i1} means that i's claim is "true" and ω_{i2} means it is "false". The states are unknown and independent, with $prob\{\omega_i = \omega_{i1}\} = 1/2$. The receiver chooses $a_i \in \{a_{i1}, a_{i2}\}$ for each sender i, where a_{i1} is "publication" and a_{i2} is "rejection".

If only sender i publishes, then he gets the reward with certainty. If both senders publish, then each obtains the reward with probability $1/2$. Otherwise, no sender gets the reward. Denote by ρ_i sender i's

E-mail address: mike.felgenhauer@brunel.ac.uk.

[1] This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

[2] Competition may also improve incentives to work hard/to be creative. This paper finds positive aspects of competition by exclusively focusing on misconduct.

[3] In research tournaments (as in Taylor (1995)), the size of the prize for the winner has an impact on behavior in the tournament. The idea that a change of the reward probability affects researcher behavior is similar.

[4] For example, a University may only want to hire a candidate with a publication. The University is able to observe publications, but not necessarily unpublished results. Frankel and Kasy (2022) assume that the public only observes published results.

https://doi.org/10.1016/j.econlet.2024.112017
Received 14 June 2024; Received in revised form 29 August 2024; Accepted 5 October 2024
Available online 5 October 2024

interim anticipated probability to obtain the reward upon publishing.[5] The benefit from the reward is lower in state ω_{i2} (published claim is false) than ω_{i1} (published claim is true). Sender i with common knowledge parameter $\theta \in (0,1)$ receives gross utility[6]:

	$\omega_i = \omega_{i1}$	$\omega_i = \omega_{i2}$
reward for sender i	1	θ
no reward for sender i	0	0

The outcome of sender i's experiment k is $\sigma_{ik} \in \{s_{i1}, s_{i2}\}$, where s_{i1} is a "publishable outcome" (argument that supports his claim) and s_{i2} is a "non-publishable outcome" (an argument against his claim is not interesting enough for publication). Outcome s_{ij} correctly reflects state ω_{ij} with probability $\pi \in (1/2, 1]$. Sender i privately observes experimentation history h_{it}. The posterior probability $prob\{\omega = \omega_{i1} \mid h_{it}\}$ if h_{it} only contains non-publishable outcomes is μ_{it}. Experimentation costs $c_E \in (0, \frac{\pi}{4} + \frac{1-\pi}{4}\theta)$ per experiment are subtracted from gross utility. The upper bound ensures that each sender either experiments or eventually manipulates, regardless of the other sender's behavior.

Sender i's message is $m_i \in \{s_{i1}, s_{i2}, \emptyset\}$, where $m_i = s_{ij}$ is feasible if $s_{ij} \in h_{it}$. Sender i may manipulate in private at costs $c_M > 0$, yielding an uninformative outcome s_{ij} and rendering $m_i = s_{ij}$ feasible. The receiver chooses a_{i1} if he observes $m_i = s_{i1}$ and a_{i2}, otherwise.[7]

Without competition, the sender first makes history-dependent choices to experiment further or to stop experimenting. Manipulation may occur after the final experiment. Next, he sends his message. With competition, senders move simultaneously. The receiver observes the message(s) and chooses his action(s). A sender's behavior is sequentially rational (given the correctly anticipated other sender's equilibrium strategy if there is competition).

The decision quality for sender i is the probability that the receiver's action is "correct": $DQ_i = \frac{1}{2} prob\{m_i = s_{i1} \mid \omega_i = \omega_{i1}\} + \frac{1}{2} prob\{m_i \neq s_{i1} \mid \omega_i = \omega_{i2}\}$. The average decision quality is $\sum_{i=1}^{2} DQ_i/2$ with competition and it is DQ_i without competition.[8]

3. Analysis

3.1. No competition

This subsection directly follows FX. If the sender observes a publishable outcome, then he sends a corresponding message.

Suppose there is no publishable outcome in history h_{it} and manipulation is not possible (or too costly). The sender stops unsuccessfully if

$$EU_{it}^{S} = \mu_{it}(\pi\rho_i - c_E) + (1-\mu_{it})((1-\pi)\rho_i\theta - c_E) < 0. \tag{1}$$

The LHS of (1) contains the payoffs from a single further experiment and then stopping after either outcome in state ω_{i1} (which realizes with probability μ_{it}) and state ω_{i2} (with probability $(1-\mu_{it})$). In ω_{i1}, this experiment yields a publishable outcome with probability π and, as there is only one sender, the sender anticipates a reward upon publication with probability $\rho_i = 1$. Experimentation costs are subtracted. The payoff in ω_{i2} is analogous.

Suppose (1) holds at a finite t and let T_{iE} be the lowest such t. The sender's continuation utility at h_{it} with $t < T_{iE}$ is[9]

$$EU_{it}^{E} = \mu_{it} \sum_{n=0}^{T_{iE}-t-1} (1-\pi)^n(\pi\rho_i - c_E) + (1-\mu_{it}) \sum_{n=0}^{T_{iE}-t-1} \pi^n((1-\pi)\rho_i\theta - c_E) \tag{2}$$

and $EU_{it}^{E} = 0$ if $t \geq T_{iE}$. The sender never stops experimenting unsuccessfully if inequality (1) is violated at $\mu_{it} = 0$.

Manipulation yields a publication and payoff $(\rho_i - c_M)$ in state ω_{i1} and $(\rho_i\theta - c_M)$ in state ω_{i2}. The continuation utility from manipulation at h_{it} is

$$EU_{it}^{M} = \mu_{it}(\rho_i - c_M) + (1-\mu_{it})(\rho_i\theta - c_M). \tag{3}$$

The sender eventually stops unsuccessfully at T_{iE} if $EU_{it}^{E} > EU_{it}^{M}$ for all $t \leq T_{iE}$. Otherwise, he manipulates at some T_{iM}, with $T_{iM} \leq T_{iE}$.

Experimentation with unsuccessful stopping implies a higher decision quality than experimenting without unsuccessful stopping or eventual manipulation. Less excessive private experimentation (without manipulation) improves the decision quality.[10]

3.2. Positive aspects of competition

Introducing competition is an exogenous institutional change and this change exogenously lowers the anticipated probability to obtain a reward upon publication from $\rho_i = 1$ (no competition) to some $\rho_i < 1$ (in any equilibrium with competition, as the other sender obtains a publishable outcome with some probability either by experimentation or manipulation). With competition, there can be multiple equilibria, which can be symmetric or asymmetric (an illustrative example can be found in the Appendix) and the size of ρ_i depends on the other sender's equilibrium behavior.[11]

Consider a sender i and let us study how an exogenous change of ρ_i affects i's incentives to manipulate and experiment.

Lemma 1. *Suppose there is a ρ_i^* such that $EU_{it}^{E}(\rho_i^*) = EU_{it}^{M}(\rho_i^*)$ at a h_{it} with $t \leq T_{iE}$ (and T_{iE} being finite) that does not contain a publishable outcome. For any ρ_i', with $\rho_i' > \rho_i^*$, we have $EU_{it}^{E}(\rho_i') < EU_{it}^{M}(\rho_i')$.*

At ρ_i^* the sender is indifferent between manipulating at t and experimenting further without manipulation (and, thus, the sender manipulates). Lemma 1 shows that an increase of ρ_i then makes manipulation more attractive. The reason is that experimentation without manipulation may be unsuccessful and implies a lower publication probability than manipulation. The effect of an increase of ρ_i is, thus, more pronounced on the benefit from manipulation than the benefit from experimentation.

We can now study the impact of competitive pressure on the senders' behavior and the average decision quality. Consider a single sender without competition who does not manipulate ($EU_{it}^{E}(\rho_i) > EU_{it}^{M}(\rho_i)$ at each $t \leq T_{iE}$ with $\rho_i = 1$). In response to tougher competition, a switch to eventual manipulation cannot occur: If that were possible instead, then, by continuity, this would imply that there

[5] ρ_i depends on the probability with which the other sender publishes.

[6] A sender's gross utility is zero if he never stops experimenting.

[7] The receiver's behavior can be endogenized (see Appendix). Messages s_{i2} and $\emptyset$ can be interpreted as not submitting a paper for publication.

[8] Two senders can generate "more" information than one. By focusing on the average decision quality, this paper finds positive aspects of competition without this effect.

[9] EU_{it}^{E} extends the left hand side of inequality (1). For example, $(1-\pi)^n$ is the probability that the next n outcomes are non-publishable in state ω_{i1}. The following experiment then yields payoff $(\pi\rho_i - c_E)$ in this state. If $t = T_{iE} - 1$, then EU_{it}^{E} reduces to the left hand side of (1).

[10] It is equal and highest at $T_{iE} = 1$ and $T_{iE} = 2$. It decreases in T_{iE} for $T_{iE} > 2$ (see Appendix).

[11] At least one equilibrium exists with competition. A detailed equilibrium characterization, equilibrium conditions or equilibrium selection do not matter for the major result, as long as tougher competition for the reward lowers the chance to obtain the reward upon publication ρ_i. This is guaranteed for the parameters here where the number of senders is increased from 1 to 2. Increasing the number of senders further would require equilibrium selection such that ρ_i decreases in the number of senders.

were some $\rho_i^* < 1$ such that $EU_{it}^E(\rho_i^*) = EU_{it}^M(\rho_i^*)$ at some $t \leq T_{iE}$.[12] But, according to Lemma 1, we then have $EU_{it}^E(\rho_i) < EU_{it}^M(\rho_i)$ at $\rho_i = 1$. This, however, contradicts that a single sender without competition does not manipulate (which implies $EU_{it}^E(\rho_i) > EU_{it}^M(\rho_i)$ at $\rho_i = 1$).[13] Therefore, if a sender without competition experiments without manipulation, then senders also experiment without manipulation in any equilibrium with competition. Furthermore, a sender without competition experiments more excessively than with competition, as the benefit from running a further experiment EU_{it}^S in (1) increases in the probability to obtain a reward upon publication ρ_i. Lemma 2 summarizes these findings.

Lemma 2. *Consider a sender without competition and suppose that it is sequentially rational for this sender to experiment without manipulation. Then with competition, (i) both senders experiment without manipulation in equilibrium, and (ii) the expected number of experiments run by each sender is weakly lower than for a sender without competition.*

As a consequence of Lemma 2, the average decision quality with competition is higher than without competition given that the sender without competition does not manipulate. If the sender without competition eventually manipulates instead, then competition cannot deteriorate the average decision quality further: The worst that could happen is that both senders manipulate, yielding the same average decision quality as without competition. The major result directly follows.

Proposition 1. *The average decision quality in any equilibrium with competition is weakly higher than without competition.*

Each sender's expected gross utility from a publication with competition is lower than without competition. Higher competitive pressure cannot encourage a switch from experimentation to manipulation, but the reverse may occur. Competition also reduces excessive experimentation. Both effects have a positive impact on the average decision quality.

4. Discussion

The paper studies increased competitive pressure, but not its causes. The causes may also affect the outside option. Consider a job market in a pandemic, where Universities hire less. The pandemic also reduces job opportunities in other sectors. Even though there is tougher competition for jobs in the academic sector, they may be even more valuable if the situation elsewhere is worse. This could be viewed as an increase of the *reward for sender i* from 1 to $1 + \tau$ in state ω_{i1} and from θ to $\theta + \tau$ in state ω_{i2} with some $\tau > 0$. The consequences of this increase are similar to an exogenous increase of ρ_i, which suggests that the detrimental forms of misconduct become more attractive. The situation should be different for researchers with tenure, whose jobs are relatively safe, but for whom promotions etc. should be harder to obtain due to the pressure on University budgets. Here, the more severe forms of misconduct should become less attractive.

The paper argues that researchers' incentives to engage in different forms of misconduct depend on the probability with which the reward is awarded. Tougher competition reduces the award probability and improves incentives. An alternative approach would be to directly modify the award probability. A University could, for example, reduce the probability to hire a candidate with a publication in order to reduce misconduct. But then, by chance, the University may end up without hiring someone who does work that is associated with the job. The model aims to describe situations where the latter effect is first order.

Competition here affects a researcher's probability to obtain a reward upon publication, but not the publication probability. An alternative view is that researchers compete with their papers for limited journal space (and that they consider the publication of their work as a reward). In this case, ρ_i could be interpreted as the probability to publish upon presenting an argument supporting a publishable claim (that is, upon sending a message containing a publishable outcome s_{i1}). More competition then also reduces ρ_i. The incentives to engage in excessive private experimentation and manipulation are analogous to above. Future work could explore competition where it matters to publish first.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Appendix

The decision quality decreases in T_{iE}: Suppose sender i does not manipulate. We have $DQ_i = \frac{1}{2}prob\{m_i = s_{i1} \mid \omega_i = \omega_{i1}\} + \frac{1}{2}prob\{m_i \neq s_{i1} \mid \omega_i = \omega_{i2}\} = \frac{1}{2}(1-(1-\pi)^{T_{iE}}) + \frac{1}{2}\pi^{T_{iE}} = \frac{1}{2}(1-(1-\pi)^{T_{iE}} + \pi^{T_{iE}})$, where $(1-(1-\pi)^{T_{iE}})$ is the ex ante probability that there is a publishable outcome in state ω_{i1} and $\pi^{T_{iE}}$ is the ex ante probability that there is no publishable outcome in state ω_{i2}.[14] We have $DQ_i = \pi$ at $T_{iE} = 1$ and $T_{iE} = 2$. Let us compare DQ_i with T_{iE} and $T_{iE} - 1$ for $T_{iE} > 2$. We have $\frac{1}{2}(1-(1-\pi)^{T_{iE}} + \pi^{T_{iE}}) < \frac{1}{2}(1-(1-\pi)^{T_{iE}-1} + \pi^{T_{iE}-1}) \Leftrightarrow -(1-\pi)^{T_{iE}} + (1-\pi)^{T_{iE}-1} < \pi^{T_{iE}-1} - \pi^{T_{iE}} \Leftrightarrow (-(1-\pi)+1)(1-\pi)^{T_{iE}-1} < \pi^{T_{iE}-1}(1-\pi) \Leftrightarrow \pi(1-\pi)^{T_{iE}-1} < \pi^{T_{iE}-1}(1-\pi) \Leftrightarrow (1-\pi)^{T_{iE}-2} < \pi^{T_{iE}-2}$, with $\pi > 1-\pi$ by assumption.

Proof of Lemma 1. Consider a given t (which does not change throughout the proof) as described in the Lemma. At ρ_i^* we have $EU_{it}^E(\rho_i^*) = EU_{it}^M(\rho_i^*)$:

$$\mu_{it} \sum_{n=0}^{T_{iE}-t-1} (1-\pi)^n(\pi\rho_i^* - c_E) + (1-\mu_{it}) \sum_{n=0}^{T_{iE}-t-1} \pi^n((1-\pi)\rho_i^*\theta - c_E)$$
$$= \mu_{it}(\rho_i^* - c_M) + (1-\mu_{it})(\rho_i^*\theta - c_M)$$

Consider a $\rho_i' = \rho_i^* + \tau$, with $\tau > 0$. Now it is established that $EU_{it}^E(\rho_i^* + \tau) - EU_{it}^M(\rho_i^* + \tau) < 0$. An increase of ρ_i may (i) increase T_{iE} or (ii) not change T_{iE} if the sender experiments without manipulation. The proof of (ii) is analogous to (i).

(i) An increase of ρ_i increases T_{iE} to some T_{iE}', with $T_{iE}' > T_{iE}$: $EU_{it}^E(\rho_i^* + \tau) - EU_{it}^M(\rho_i^* + \tau) = \mu_{it} \sum_{n=0}^{T_{iE}-t-1}(1-\pi)^n(\pi(\rho_i^* + \tau) - c_E) + (1-\mu_{it}) \sum_{n=0}^{T_{iE}-t-1} \pi^n((1-\pi)(\rho_i^* + \tau)\theta - c_E)$
$+\mu_{it} \sum_{n=T_{iE}-t}^{T_{iE}'-t-1}(1-\pi)^n(\pi(\rho_i^*+\tau)-c_E)+(1-\mu_{it}) \sum_{n=T_{iE}-t}^{T_{iE}'-t-1} \pi^n((1-\pi)(\rho_i^*+\tau)\theta-c_E)$
$-(\mu_{it}((\rho_i^* + \tau) - c_M) + (1-\mu_{it})((\rho_i^* + \tau)\theta - c_M))$
where the third line captures the additional experiments after T_{iE}. $= \mu_{it} \sum_{n=0}^{T_{iE}-t-1}(1-\pi)^n\pi\tau + (1-\mu_{it}) \sum_{n=0}^{T_{iE}-t-1} \pi^n(1-\pi)\tau\theta + \mu_{it} \sum_{n=T_{iE}-t}^{T_{iE}'-t-1}(1-\pi)^n(\pi(\rho_i^* + \tau) - c_E) + (1-\mu_{it}) \sum_{n=T_{iE}-t}^{T_{iE}'-t-1} \pi^n((1-\pi)(\rho_i^* + \tau)\theta - c_E) - (\mu_{it}\tau + (1-\mu_{it})\tau\theta)$, where the equality holds since $EU_{it}^E(\rho_i^*) = EU_{jt}^M(\rho_i^*)$. $= \mu_{it} \sum_{n=0}^{T_{iE}'-t-1}(1-\pi)^n\pi\tau + (1-\mu_{it}) \sum_{n=0}^{T_{iE}'-t-1} \pi^n(1-\pi)\tau\theta + \mu_{it} \sum_{n=T_{iE}-t}^{T_{iE}'-t-1}(1-\pi)^n(\pi\rho_i^* - c_E) + (1-\mu_{it}) \sum_{n=T_{iE}-t}^{T_{iE}'-t-1} \pi^n((1-\pi)\rho_i^*\theta - c_E) - (\mu_{it}\tau + (1-\mu_{it})\tau\theta)$.

We have $\mu_{it} \sum_{n=0}^{T_{iE}'-t-1}(1-\pi)^n\pi\tau + (1-\mu_{it}) \sum_{n=0}^{T_{iE}'-t-1} \pi^n(1-\pi)\tau\theta < (\mu_{it}\tau + (1-\mu_{it})\tau\theta)$, since $\sum_{n=0}^{T_{iE}'-t-1}(1-\pi)^n\pi < 1$ and $\sum_{n=0}^{T_{iE}'-t-1} \pi^n(1-\pi) < 1$ for any finite T_{iE}'. Finally, it is shown that $\mu_{it} \sum_{n=T_{iE}-t}^{T_{iE}'-t-1}(1-\pi)^n(\pi\rho_i^* - c_E) + (1-\mu_{it}) \sum_{n=T_{iE}-t}^{T_{iE}'-t-1} \pi^n((1-\pi)\rho_i^*\theta - c_E) < 0$. The reason is that the sender with ρ_i^*

[12] T_{iE} decreases if ρ_i decreases (see Lemma 2 (ii)).

[13] Mixing between manipulation and experimentation at some history h_{it} requires $EU_{it}^M = EU_{it}^E$.

[14] Note that $(1-(1-\pi)^{T_{iE}})$ is equal to $\pi+(1-\pi)\pi+(1-\pi)^2\pi+\cdots+(1-\pi)^{T_{iE}-1}\pi = \pi \sum_{t=0}^{T_{iE}-1}(1-\pi)^t$ (where π is the probability to find a publishable outcome in state ω_{i1} with the next experiment, which in turn is only run if all previous experiments yielded non-publishable outcomes).

prefers stopping unsuccessfully to running a single further experiment at T_{iE} and any more experiments are worse than a single further experiment. Suppose the sender runs a further experiment, that is, $T'_{iE} = T_{iE}+1$: $\mu_{it}\sum_{n=T_{iE}-t}^{T'_{iE}-t-1}(1-\pi)^n(\pi\rho_i^*-c_E)+(1-\mu_{it})\sum_{n=T_{iE}-t}^{T'_{iE}-t-1}\pi^n((1-\pi)\rho_i^*\theta-c_E) = \mu_{it}\sum_{n=T_{iE}-t}^{T_{iE}-t}(1-\pi)^n(\pi\rho_i^*-c_E)+(1-\mu_{it})\sum_{n=T_{iE}-t}^{T_{iE}-t}\pi^n((1-\pi)\rho_i^*\theta-c_E) = \mu_{it}(1-\pi)^{T_{iE}-t}(\pi\rho_i^*-c_E)+(1-\mu_{it})\pi^{T_{iE}-t}((1-\pi)\rho_i^*\theta-c_E)$.

The sender's expected utility with ρ_i^* from a further experiment at T_{iE} is worse than stopping unsuccessfully at T_{iE}:

$$\mu_{iT_{iE}}(\pi\rho_i^*-c_E)+(1-\mu_{iT_{iE}})((1-\pi)\rho_i^*\theta-c_E)<0. \tag{4}$$

Plugging in $\mu_{iT_{iE}} = \frac{\mu_{it}(1-\pi)^{T_{iE}-t}}{\mu_{it}(1-\pi)^{T_{iE}-t}+(1-\mu_{it})\pi^{T_{iE}-t}}$ yields $\frac{\mu_{it}(1-\pi)^{T_{iE}-t}}{\mu_{it}(1-\pi)^{T_{iE}-t}+(1-\mu_{it})\pi^{T_{iE}-t}}(\pi\rho_i^*-c_E)+\frac{(1-\mu_{it})\pi^{T_{iE}-t}}{\mu_{it}(1-\pi)^{T_{iE}-t}+(1-\mu_{it})\pi^{T_{iE}-t}}((1-\pi)\rho_i^*\theta-c_E)<0 \Leftrightarrow \mu_{it}(1-\pi)^{T_{iE}-t}(\pi\rho_i^*-c_E)+(1-\mu_{it})\pi^{T_{iE}-t}((1-\pi)\rho_i^*\theta-c_E)<0$. Therefore, $EU_{it}^E(\rho_i^*+\tau)-EU_{it}^M(\rho_i^*+\tau)<0$. □

Endogenous receiver behavior

Replication studies suggest that misconduct matters. Yet, publications cannot be uninformative in expected terms if editors care about the value of contributions. Hence, even if manipulation cannot be directly observed, but is anticipated to occur, there must also be informative contributions for papers to be worth publishing. The extension below incorporates this idea with endogenous receiver behavior.

Without loss of generality consider only one sender i. The receiver's utility is

	$\omega_i=\omega_{i1}$	$\omega_i=\omega_{i2}$
$a_i=a_{i1}$	1	$1-p_d$
$a_i=a_{i2}$	p_d	1

with $p_d \in (1/2,1)$. At the optimum he only chooses $a_i = a_{i1}$ if his posterior belief passes the "threshold of doubt" p_d, that is, the posterior that $\omega_i=\omega_{i1}$ must be greater than p_d. Suppose there are two sender types, ϑ_1 and ϑ_2 (each with ex ante probability 1/2). Type ϑ_1 has preferences as above and ϑ_2 is an "honest" type that runs a single experiment (without manipulation).[15] The sender privately observes his type.

(i) Consider parameters such that the sender above manipulates. Given that p_d is sufficiently close to 1/2, there is an equilibrium in which type ϑ_1 manipulates here as well and where the receiver chooses a_{i1} only if $m=s_{i1}$: The receiver's decision rule is as above, and, hence, type ϑ_1's behavior is a best response and uninformative. Type ϑ_2's behavior is such that his message $m=s_{i1}$ correctly reflects ω_{i1} with probability $\pi>1/2$. Hence, $m=s_{i1}$ is more likely in state ω_{i1} and $m\neq s_{i1}$ is more likely in state ω_{i2}. The receiver anticipates each type's behavior and forms equilibrium beliefs about ω_i. There are no off-the equilibrium path events. Thus, the receiver's behavior is a best response if p_d is sufficiently low.

(ii) Consider parameters such that the sender above does not manipulate. By an analogous argument as in (i) there is an equilibrium where type ϑ_1's and the receiver's behavior are as above if p_d is sufficiently low.

Example of symmetric and asymmetric equilibria with competition: Suppose c_M is sufficiently high (no manipulation) and $\pi=0.75$, $\theta=0.8$, $c_E=0.11$. First, two equilibria are described and then the equilibrium conditions are checked:

There is a symmetric equilibrium where each sender runs at most 6 experiments, with $\rho_i = 0.545$ for each sender. There is also an asymmetric equilibrium where sender i runs at most 3 experiments and sender $-i$ never stops unsuccessfully. In this equilibrium, $\rho_i=0.5$ (as sender $-i$ publishes almost with certainty) and $\rho_{-i}=0.609$. Given that $-i$ experiments longer than in the symmetric equilibrium, the prospects to get a reward upon publication for i are lower than in the symmetric equilibrium. Due to this lower benefit from a publication, i stops unsuccessfully earlier than in the symmetric equilibrium. Given that i experiments less in the asymmetric equilibrium, the prospects for getting a reward for $-i$ upon publication are better than in the symmetric equilibrium and experimenting longer is optimal for $-i$. Indeed, ρ_{-i} is sufficiently high such that $-i$ prefers to continue experimenting even if he knows that his state is adverse.

Equilibrium conditions: The probability that sender i obtains a reward upon publication ρ_i depends on $-i$'s publication probability. Thus, in a pure strategy equilibrium, ρ_i depends on when $-i$ stops experimenting unsuccessfully T_{-iE}[16]:

$$\rho_i = \frac{(\frac{1}{2}(1-(1-\pi)^{T_{-iE}})+\frac{1}{2}(1-\pi^{T_{-iE}}))}{2}+(\frac{1}{2}(1-\pi)^{T_{-iE}}+\frac{1}{2}\pi^{T_{-iE}}) \tag{5}$$

The tables below show the expected utility from running a single further experiment EU_{it}^S at each t in the two equilibria.

Symmetric equilibrium, where $\rho_i = 0.544555664$ for each i: We have $T_{iE}=6$, as $EU_{it}^S>0$ for all $t<6$ and $EU_{it}^S<0$ at $t=6$ (where the first experiment is run at $t=0$ after observing the prior $\mu_{i0}=1/2$).

t	μ_{it}	EU_{it}^S
0	0.5	0.14866394
1	0.25	0.073787537
2	0.1	0.028861694
3	0.035714286	0.009607762
4	0.012195122	0.00256364
5	0.004098361	0.000138615
6	0.001369863	−0.000678586

Asymmetric equilibrium: Here, $\rho_1=0.5$ and $\rho_2=0.609375$. Sender 2 never stops unsuccessfully, as his expected utility from a single further experiment given that the state is adverse is $0.11875>0$. The next table shows the values for sender 1, with $T_{1E}=3$, where $EU_{it}^S>0$ for all $t<3$ and $EU_{it}^S<0$ at $t=3$:

t	μ_{1t}	EU_{1t}^S
0	0.5	0.1275
1	0.25	0.05875
2	0.1	0.0175
3	0.035714286	−0.000178571

References

Brocas, Isabelle, Carrillo, Juan D., 2007. Influence through ignorance. Rand J. Econ. 38 (4), 931–947.

Celik, Gorkem, 2003. Interested experts: Do they know more? working paper.

Felgenhauer, Mike, Xu, Fangya, 2021. The face value of arguments with and without manipulation. Internat. Econom. Rev. 62 (1), 277–293.

Frankel, Alexander, Kasy, Maximilian, 2022. Which findings should be published? Am. Econ. J. Microecon. 14 (1), 1–38.

Henry, Emeric, 2009. Strategic disclosure of research results: The cost of proving your honesty. Econ. J. 119, 1036–1064.

Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. Amer. Econ. Rev. 101 (6), 2590–2615.

Taylor, Curtis R., 1995. Digging for golden carrots: An analysis of research tournaments. Amer. Econ. Rev. 85 (4), 872–890.

[15] Assuming an honest type is simplifying. What matters is that this type does not manipulate at the optimum. Assuming a type that exclusively experiments with unsuccessful stopping (but that runs more than one experiment) would not qualitatively change the arguments.

[16] As the states are independent, i's experimentation is uninformative about $-i$'s state. $\frac{1}{2}(1-(1-\pi)^{T_{-iE}})+\frac{1}{2}(1-\pi^{T_{-iE}})$ is the probability that $-i$ obtains a publishable outcome (in which case i's probability to obtain a reward upon publication is 1/2) and $(\frac{1}{2}(1-\pi)^{T_{-iE}}+\frac{1}{2}\pi^{T_{-iE}})$ is the probability that $-i$ does not obtain a publishable outcome (in which case i obtains a reward with certainty upon publication).