Preselection and Expert Advice*

Mike FELGENHAUER and Elisabeth SCHULTE†
University of Mannheim

February 14, 2011

Abstract

We study the effects of preselection on an expert's incentive to give truthful advice. In a decision environment in which certain decisions yield more precise estimates about the expert's expertise, a mediocre expert's advice is biased. We show that this bias can be undone by the introduction of a preselection stage, where the decision maker himself sometimes studies the case, and thereby alters the expert's perception of the problem. We identify a parameter range in which the decision maker's choice is inefficient if it is not possible to commit to a certain preselection level.

Keywords: Reputation, cheap talk, safe haven.
JEL classification: D82; D83

*We thank Philippe Aghion, Jan Eeckhout, Georgy Egorov, Hans Peter Grüner, Oliver Kirchkamp, Benny Moldovanu, Nicola Persico, Ernst-Ludwig von Thadden, the seminar participants in Mannheim, Heidelberg, the ESEM 2009 in Barcelona, the 2010 CEME/NSF Decentralization Conference and ESWC 2010 in Shanghai for useful comments and suggestions.

†Correspondence address: Elisabeth Schulte, University of Mannheim, Department of Economics, D-68131 Mannheim, Germany; email: elisabeth.schulte@vwl.uni-mannheim.de; Tel: +49 621 181 3447. Elisabeth Schulte gratefully acknowledges support through the Collaborative Research Center 884.

1 Introduction

In many decision environments, ex post information about the decision quality depends on the chosen alternative. The cost of a project can typically only be observed if it is implemented, but not if the project is not carried out. Wrong hiring decisions become more easily obvious if a bad candidate is hired than if a good candidate is not hired. In a lobbying setting, Leaver (2009) shows that an interested party has an incentive to reveal the state of the world only for one type of decision in order to induce biased decision making. In the academic publication process, the quality of a manuscript can be estimated more precisely if it is published rather than if it is rejected, because published papers are usually more frequently read than unpublished ones.

Decision makers often seek the advice of experts before making a choice. Experts in turn tend to be concerned about being perceived as well informed. Such reputational concerns may lead to biased advice. Imperfectly informed experts tend to bias their advice towards the less informative alternative. We show how the introduction of a preselection stage affects such an expert's advice and we study the associated effects on the quality of the decision.

Preselection means that the decision maker may have a glance at the decision alternatives and discard some of them without consulting the expert. Such a procedure is often used in practice. In firms, a set of projects is pre-screened on one management layer and only a subset is proceeded to the next one. In hiring procedures, a committee selects candidates from the pool of applicants before a group of experts evaluates the candidates. In banks, granting a credit typically requires the approval on a higher hierarchy level, whereas the denial of a credit can often be decided on the lowest hierarchy level.

We choose the academic publication process as our leading example because many journals apply or consider to apply a desk-rejection policy. Desk-rejection means that the editor sometimes rejects a paper without asking for the advice of a referee. Such a procedure economizes on refereeing resources, but there is a risk that a good paper will be desk-rejected.[1] We show that the referee's strategic adjustment to the information processing procedure has additional effects on the quality of the journal. Indeed, with the introduction of a desk-rejection policy it can become more likely to accept a bad paper.

The application of a desk-rejection policy affects a referee's incentive to provide sincere advice, because it alters the problem. With a preselection stage, a manuscript's expected quality, conditional on being forwarded to the referee, is higher than in the case without preselection. In addition, the "safe haven" alternative (in this case rejecting the paper) becomes less safe, because the editor may hold a different view on the manuscript. For a referee who cares about being perceived as well informed by the editor, both effects make the recommendation to publish the paper more attractive.

The higher the desk-rejection rate is, the higher is the expected quality of a manuscript

[1] It is reasonable to assume that the editor has only limited time for pre-screening. Note also that if the editor has perfect information about the manuscript's quality, then the referee's advice is superfluous.

which enters the refereeing process and the more likely the recommendation to publish the paper is correct. Hence, the manuscript is recommended for publication with a higher probability. Consequently, a good manuscript is recommended for publication more often. As a downside, it also becomes more likely that a bad manuscript is recommended for publication, first, because the mediocre referee's (positive) information is imperfect and second, because he may even be induced to recommend publishing the manuscript upon the observation of a negative signal. However, an appropriately chosen preselection rule induces sincere advice.

At the preselection stage, the editor himself adds mistakes if his own information is not perfect. It is possible that he desk-rejects a paper that is in fact good. The organization of information processing determines the likelihood with which both types of mistake, rejecting a good paper or accepting a bad one, occur. The optimal preselection rule depends on the cost of information processing on both stages and on how the editor trades off one mistake against the other. The tougher the preselection stage is, the higher is the average quality of papers which reach the refereeing stage, but the lower is the average quality of a paper recommended for publication. As a consequence, the average quality of published papers may decrease when the journal applies a desk-rejection rule.

If the editor's own signal is too weak, then a good paper is rejected too often at the preselection stage. At the same time, a bad paper is recommended for publication more often at the refereeing stage. Then, both types of mistake are more likely to occur with a preselection stage than without. There are parameter constellations for which the editor applies a desk-rejection policy in order to reduce the expected cost of refereeing, but the referee's strategic adjustment to this policy can be detrimental. The editor would be better off if he could commit not to preselect at all. However, if the referee cannot observe the editor's information acquisition choice, he anticipates that the editor will preselect and behaves accordingly.

In our analysis of the effects of preselection on the referee's advice and on the quality of the decision in Section 2.1, we take the editor's information acquisition and decision making strategy as given. The analysis in this section is of separate interest, as the decision maker might indeed be committed to the rule.[2] In general, an editor has full discretion with respect to information acquisition and publication decisions. We therefore allow for strategic behavior on part of the editor in Section 2.2, assuming that information acquisition is costly at both stages, the preselection stage and the refereeing stage. We identify a set of parameter constellations for which the editor applies a desk-rejection policy with the intention to reduce the expected cost of refereeing, anticipating the referee's reporting behavior. The purpose of the analysis in Section 2.2 is twofold: First, we make sure that the editor's behavior that we postulate in Section 2.1 can indeed arise endogenously in equilibrium. Second, the analysis in Section 2.2

[2]Applying our model to the organization of a firm, an additional management layer has similar effects on reporting behavior on lower levels as a desk-rejection rule in the academic refereeing process. The firm organization is possibly not easily adjusted to each particular problem, such that the preselection rule is best thought of as exogenous.

allows us to show in Section 2.3, that there are parameter constellations for which the editor is strictly better off if he can commit to refrain from preselection, but in the unique equilibrium, he always preselects. We continue our analysis of the effects of preselection on the expert's advice in Section 2.4 with a discussion of the referee's participation and information acquisition behavior when both are costly. The final section concludes.

Related literature

The seminal papers on reputational concerns and their effects on expert advice are Holmström (1999) and Scharfstein and Stein (1990). In both papers, there is uncertainty but no asymmetric information about the expert's expertise. In Scharfstein and Stein (1990), the experts' information is correlated and reputational concerns lead to herding behavior.[3] In Holmström's paper, output is an indicator for expertise and reputational concerns deter investment in profitable opportunities and distort working incentives. With private information on the agent's side as in Chen (2010), the investment choice can serve as signaling device and reputational concerns cause overinvestment in a risky asset. Likewise, in the herding setup private information about expertise can prevent herding and may cause anti-herding (e.g., Avery and Chevalier, 1999, Effinger and Polborn, 2001, and Levy, 2004). It is common in this literature to assume that the expert is directly interested in a reputation for expertise.[4]

Suurmond, Visser and Swank (2004) have shown that reputational concerns in combination with decision-dependent ex post information lead to biased advice, but may increase information acquisition incentives. In Milbourn et al. (2001), career concerns boost the expert's information acquisition incentives even without a distortion in reports. In Levy (2005), ex post revelation of information is endogenous. Some decisions are more likely to be double-checked than others. Asymmetric ex post revelation of information mitigates the incentive to contradict the (exogenous) prior in order to signal ability. In our paper instead, there is an exogenous asymmetry in the informativeness of the decision as in Suurmond, Visser and Swank (2004), but the expert's prior depends on the preselection rule. The preselection mechanism provides a counterbalance to the bias induced by the asymmetry in ex post information revelation.

Prat (2005) shows that transparency may be bad for the decision maker. In his model, if the decision maker cannot observe the expert's actions, they reflect the correlation between the expert's information and the state. If the decision maker can observe the expert's actions, they reflect the correlation between the signal and the expert's type (about which players are symmetrically informed). In our model, the decision maker can choose to become (imperfectly)

[3] See also Ottaviani and Sørensen (2001). In Ottaviani and Sørensen (2006), they show that reputational concerns yield a bias towards reporting the prior mean in a continuous version of the model in which information about the state is revealed ex post. They show that in equilibrium, information transmission is binary.

[4] In Sobel (1985), Bénabou and Laroque (1992), Morris (2001) and Lagerlöf and Frisell (2007), there is an instrumental preference for a reputation for honesty, which increases the sender's impact on the decisions in later periods.

informed prior to consulting the expert. As the expert anticipates that the decision maker is informed, he uses the correlation between the decision maker's signal and the state, which can dilute the informativeness of his advice.

2 The model

An editor must decide about the publication of a manuscript. The manuscript's quality $\omega \in \{g, b\}$ may be either good or bad. Assume, for simplicity, that both states are equally likely. The editor derives state-dependent utility from accepting or rejecting the manuscript, as given in the following table:[5]

	$\omega = g$	$\omega = b$
accept	u_{ag}	u_{ab}
reject	u_{rg}	u_{rb}

Table 1: The editor's payoffs

The editor would like to accept good papers and to reject bad papers, i.e.,

$$\begin{aligned} u_{ag} &> u_{rg}, \\ u_{rb} &> u_{ab}. \end{aligned}$$

At a cost ξ, the editor can buy a signal, which reveals the true paper quality with probability $1 - \varepsilon$. Let η denote the probability with which the editor acquires a signal. We call (η, ε) a *preselection rule*. We will say the preselection rule becomes "tougher", if c.p. $\eta \uparrow$ or $\varepsilon \downarrow$ (or both). We denote the editor's signal with $s_e \in \{\emptyset, g, b\}$, where $\emptyset$ means that the editor did not acquire a signal. Next, the editor can either decide to accept or reject the paper without further information acquisition, or to consult a referee in which case the editor has to pay c.[6] For the time being, we take the referee's participation for granted. We show in Section 2.4 that if the cost of reviewing for the referee is sufficiently low, it is optimal to always participate.

If the manuscript enters the refereeing process, the referee receives a signal about its quality. The referee may be either mediocre or smart, where q is the ex ante probability that he is smart. The referee privately knows his type. For simplicity, we assume that a smart referee privately observes the paper quality ω perfectly. A mediocre referee privately observes an imperfect signal $s_r \in \{g, b\}$ about ω which is correct with probability $p \geq 1/2$. The referee recommends either to accept the paper for publication ($x = g$) or to reject the paper ($x = b$).[7] Denote by γ_θ the

[5] Information acquisition costs have to be subtracted.

[6] The cost c could either arise from the need to search for a referee or be a transfer to the referee. The payment does not condition on the referee's report and is hence not relevant for the referee's reporting strategy.

[7] Laband (1990) presents evidence that referees increase the quality of a manuscript by providing feedback to the authors. We focus on the screening role of referees.

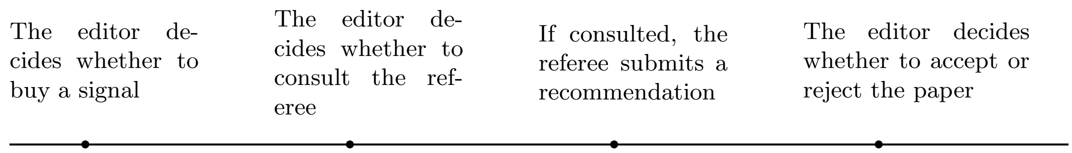

Figure 1: Timing of actions

probability that a referee of type $\theta \in \{smart, mediocre\}$ sends a positive report ($x = g$) after observing a positive signal and similarly denote by β_θ the probability that a referee of type θ sends a negative report ($x = b$) after observing a negative signal. $\gamma_\theta = \beta_\theta = 1$ corresponds to type θ giving sincere advice. After obtaining the referee's advice, the editor decides whether to publish the manuscript. The sequence of actions is illustrated in Figure 1.

The editor uses all the observable information and his belief about the referee's reporting strategy $(\hat{\gamma}_\theta, \hat{\beta}_\theta)$ to update his assessment of the paper's quality. If the manuscript is published, its true quality becomes obvious. If it is rejected, there is no opportunity for further learning about the quality. The editor forms an opinion about the referee's expertise, which is measured in terms of the probability that the referee is smart. We denote this probability by $\pi = prob\{smart|x, s_e, \omega, \hat{\gamma}_{smart}, \hat{\beta}_{smart}, \hat{\gamma}_{mediocre}, \hat{\beta}_{mediocre}\}$. π depends on the editor's ex post information. Denote with π_{gg} the probability which the editor assigns to the referee being smart, given that he recommended publication ($x = g$) and the manuscript turned out to be good ($\omega = g$), and with π_{gb} if it turned out to be bad. Accordingly, denote with $\pi_{b\emptyset}$ the probability which the editor assigns to the agent being smart, given that he recommended a rejection ($x = b$) and that the editor has not observed any further information ($s_e = \emptyset$). π_{bg} denotes the probability which the editor assigns to the agent being smart, given that he recommended a rejection and the editor has observed a good signal.

The referee cares about π. This could be for instrumental reasons, as the editor may for instance be a key player in the profession and his opinion may be relevant on the job market. It could also be the case that the referee just enjoys being thought of as smart. We represent the referee's preferences by the following utility function:

$$U = \pi.$$

The referee does not care about the editor's choice per se. He is only interested in leaving a good impression. If in addition the referee cares about the quality of the decision (from which we abstract in this paper, but which is certainly relevant in reality), there is an additional incentive to tell the truth. The magnitude of the effect that we are pointing at in this paper then depends on the relative importance of the referee's payoff components.[8]

[8]Conversations with colleagues revealed that almost everybody would be embarrassed if a paper which they recommended for publication contained a serious flaw. Some are also concerned that a paper they advised to

The equilibrium concept is Perfect Bayesian Nash Equilibrium. We restrict attention to equilibria in which the referee behaves sincerely if he is smart.[9] In the following, (γ, β) refers to the mediocre referee's reporting behavior and $(\hat{\gamma}, \hat{\beta})$ to the editor's belief about it. The proofs of our statements can be found in the appendix.

2.1 The referee's recommendation

In our analysis of the mediocre referee's behavior, we take the editor's strategy as given. For the time being, we can think of the editor as a non-strategic player. We assume that he acquires information with probability η and desk-rejects the paper upon the observation of a negative signal. Otherwise, he asks for the referee's recommendation and follows his advice. In Section 2.2, we identify conditions under which this behavior is indeed optimal.

As we assume a fixed editor behavior in this section, the term "equilibrium" refers to a combination of (γ, β) and $(\hat{\gamma}, \hat{\beta})$ which are mutually consistent and respect Bayes' Law. Lemma 1 characterizes the potential equilibria with and without preselection. The case without preselection is captured by $\eta = 0$.

Lemma 1 *Consider any preselection rule (η, ε). In any equilibrium in which the smart referee is sincere, either $\gamma = 1$, $\beta = 1$, or both.*

Lemma 1 asserts that there is no equilibrium in which the mediocre referee is insincere for both signals. Either there is a bias towards rejection ($\gamma < 1$, $\beta = 1$), a bias towards acceptance ($\gamma = 1$, $\beta < 1$), or truthtelling ($\gamma = 1$, $\beta = 1$). The next lemma establishes equilibrium uniqueness.

Lemma 2 *Consider any preselection rule (η, ε). In the class of equilibria in which the smart referee is sincere, a unique equilibrium exists.*

As every preselection rule is associated with a unique equilibrium, we can now study the effects of a change of the preselection rule on the referee's advice. We first consider the benchmark case without preselection $\eta = 0$, i.e., all papers enter the refereeing process without prior screening.

2.1.1 No preselection

Lemma 3 *Consider $\eta = 0$ (no preselection). For any $p < 1$, $\exists\ \overline{q}(p)$: If $q > \overline{q}(p) = \frac{2p}{1+p}$, then $\gamma = 0$, $\beta = 1$.*

reject may be published in a higher-ranking journal. However, few seem to follow the fate of papers that were rejected. This suggests that the former kind of embarrassment is more prevalent than the latter. Similar effects on reporting behavior as pointed out in this paper arise from such an "embarrassment asymmetry".

[9] Note that in this class of equilibria, the probability that the editor assigns to the referee being smart if he recommended publication and the manuscript is of bad quality is zero.

Without preselection, recommending a rejection is the safe haven alternative for the mediocre type. If q, the ex ante probability that the referee is smart, is sufficiently large, then always recommending a rejection is not too suspicious because the smart type recommends a rejection in the case that the manuscript is of bad quality. Recommending a publication yields a higher reputation, but only if this advice is correct, which is too unlikely from the mediocre referee's point of view. Hence, the mediocre referee prefers hiding in the safe haven to giving sincere advice.

Lemma 4 *Consider $\eta = 0$ (no preselection). No equilibrium with $\gamma = 1$ exists. In equilibrium, $\beta = 1$.*

Without preselection, there is no sincere equilibrium due to the drawing power of the safe haven alternative. If the mediocre referee observes a good signal, then he recommends a rejection with a positive probability. Moreover, he recommends a rejection whenever observing a bad signal. Hence, the referee's advice is biased towards rejection.

Proposition 1 *Consider $\eta = 0$ (no preselection). In equilibrium there is a bias towards rejection with $\gamma < 1$ and $\beta = 1$.*

Note that this result is sensitive to the assumption that the manuscript is good with a rather low prior probability.[10] If the prior probability is very high, as it may for instance be if the paper's authors are from top universities, the bias may be in the opposite direction.[11] Consider the extreme case, in which the paper is of high quality with certainty. In this case, $\gamma = 1, \beta = 0$ is optimal for the mediocre referee. For a very high expected paper quality, even upon the observation of a bad signal, the mediocre referee assesses the probability that the paper is good as high enough to recommend the paper for publication. For lower values, the bias towards rejection is even more pronounced. For convenience, we adhere to the assumption that the manuscript is good with probability 1/2 for the rest of the paper, yielding a recommendation biased towards rejection without preselection. The following section shows how a simple preselection mechanism may induce sincere reporting.

2.1.2 Preselection

Introducing a preselection stage changes the picture. We assume that the editor looks at the paper with a certain probability ($\eta > 0$), in which case he obtains an informative signal (i.e.,

[10] However, our results can be extended for more general prior distributions. In particular, for any q, there exists a $\lambda(q)$ such that if the prior probability that the manuscript is good is lower than $\lambda(q)$, the bias towards rejection occurs. We thank an anonymous referee for suggesting this generalization.

[11] Hence, our analysis suggests that an increase in the ex ante expected paper quality may yield a disproportionately high increase in the publication probability of the paper. Evidence in line with this argument is presented by Blank (1991), who finds that authors from near top universities have a higher publication probability under single-blind rather than double-blind refereeing procedures.

a signal coinciding with the true state of the world with probability $1 - \varepsilon > \frac{1}{2}$). If the editor obtains a bad signal, the manuscript does not enter the refereeing process. Preselection alters the referee's assessment of the manuscript's quality. The average quality of papers entering the refereeing process increases. In addition, recommending a rejection loses its safe haven character, as the referee is uncertain whether the editor has no information or a positive signal.

Proposition 2 *Consider a preselection rule* (η, ε) *with* $\eta \leq 1$ *and* $\varepsilon \geq 0$. *With a tougher preselection rule* (η', ε'), $\eta' \geq \eta$ *and* $\varepsilon' \leq \varepsilon$ *(at least one inequality strict), the probability that the mediocre referee recommends the manuscript for publication is (weakly) higher.*

The qualification "weakly" applies if under the original preselection rule, the mediocre referee has a strict preference for always recommending to reject the paper or for truthtelling. Then, a marginally tougher preselection rule has no effect on the reporting behavior. In this case, there always exist tougher preselection rules for which the probability that the mediocre referee recommends the manuscript for publication is strictly higher. In particular, there exists a class of preselection rules under which the referee has a strict preference for truthtelling.

Proposition 3 *For any* (p, q), *there is a class of preselection rules* (η, ε) *which induce truthtelling.*

The incentives to recommend a publication become stronger for higher η and lower ε, as the manuscript's quality as perceived by the mediocre referee increases. If the editor applies a too tough preselection rule, a bias towards publication arises. For instance, if the editor always acquires perfect information ($\eta = 1, \varepsilon = 0$), the mediocre referee knows that all manuscripts which he obtains are good. Therefore, he recommends a publication in equilibrium regardless of his own signal.[12]

2.1.3 Decision quality

Here, we discuss quality measures which do not require to further specify the editor's preferences (this will be done in the next section). We continue to assume that the editor is committed to a certain preselection rule, which is correctly anticipated by the referee. First, we observe that the more effective bad papers are filtered out at the preselection stage, the worse the refereeing stage performs as a filter.

Proposition 4 *The tougher the preselection rule* ($\eta \uparrow, \varepsilon \downarrow$, *or both) the (weakly) higher is the probability that a bad paper which is not desk-rejected is recommended for publication.*

[12]Note that a strict preference for always recommending the manuscript for publication can only arise if the editor acquires perfect information, i.e., $\varepsilon = 0$. Suppose instead $\varepsilon > 0$, but $\beta = 0$. Then, $\gamma = 1$ and only the smart type ever recommends a rejection. Thus, the mediocre type gains from deviating to also recommending a rejection. Hence, we must have $\beta > 0$ in equilibrium.

The tougher the preselection stage is, the (weakly) more often the mediocre referee recommends the paper for publication. The more likely it is that the mediocre type recommends the paper for publication, the more often a recommendation for publication stems from the mediocre type, who bases this recommendation on a weaker signal than the smart type. Hence, the probability that a bad paper is recommended for publication increases. As a consequence, summary rejections may have the effect that bad papers are published with a higher probability. Through the impact on the referee's behavior, preselection influences the journal's quality in a non-trivial way.

Corollary 1 *Consider the case $q > \overline{q}(p)$. The probability that an accepted paper is good is non-monotonic in the preselection rule.*

Without preselection, the probability that an accepted paper is good is one, because only smart referees recommend a publication (and only if the paper is in fact good). With the toughest preselection rule ($\eta = 1, \varepsilon = 0$), it is one, because every paper that enters the refereeing process is good and recommended for publication. In between, the probability that an accepted paper is good is smaller than one.

The mistakes which the editor's decision may involve are (i) accepting the paper if it is bad and (ii) rejecting the paper if it is good. Both result from the imperfectness of information on both stages, the preselection stage and the refereeing stage. For any preference specification which satisfies $u_{ag} > u_{rg}$ and $u_{rb} > u_{ab}$, the editor's expected utility ceteris paribus decreases in the probability with which each type of mistake occurs. If both mistakes occur (weakly, at least one strictly) more often under a preselection rule, then the editor's expected utility is unambiguously lower under that rule.

The probability to accept a bad paper is $(1 - \eta(1 - \varepsilon))(1 - q)(p(1 - \beta) + (1 - p)\gamma)$. The probability to reject a good paper is $\eta\varepsilon + (1 - \eta\varepsilon)(1 - q)(p(1 - \gamma) + (1 - p)\beta)$. It is easy to see and intuitively plausible that both mistakes are ceteris paribus less likely the lower ε, i.e., the less likely the editor makes a mistaken judgement at the preselection stage, and the higher q, i.e., the more likely the editor is advised by the smart referee. Moreover, the higher q, the less the mediocre referee's strategic adjustment to the preselection rule matters. As we have seen in Proposition 3, there is a class of preselection rules for which the mediocre referee gives sincere advice. Within this class, choosing a preselection rules with the same η and a lower ε unambiguously reduces both mistakes. Preselection rules with the same ε and a higher η in this class lead to a lower probability to accept a bad paper, but the effect on the probability to reject a good paper is opposite. Whether a tougher preselection rule is beneficial for the editor depends on how he trades off one mistake against the other (and on the cost of information processing at both stages, as we will see in the next section).

We conclude this section with the identification of preselection rules which are unambiguously bad, i.e., under which both types of mistakes are (weakly) more likely than if no preselection rule is applied. Consider again the case $q > \overline{q}(p)$, implying $\gamma = 0$ without preselection.

With a preselection stage inducing $\gamma > 0$, the editor incurs the mistake to accept a bad paper more often. If ε is too high, both mistakes occur more often with a preselection stage than without.

Proposition 5 *Consider $q > \overline{q}(p)$. For a class of preselection rules (η', ε'), the editor's expected utility (neglecting information processing costs) is strictly higher if he does not preselect.*

For any $\varepsilon > 0$, there is a $q(\varepsilon)$ such that if $q > q(\varepsilon)$, the introduction of a preselection stage has no impact on the mediocre referee's reporting behavior. If such a rule is used, the probability to reject a good paper increases compared to the case of no preselection because of the editor's mistaken judgment, whereas the probability to accept a bad paper is unaffected. In the proof of Proposition 5, we show that there is a class of preselection rules for which both types of mistake occur strictly more often.

2.2 The editor's strategy

In this section, we identify a set of parameter constellations such that the previously assumed editor behavior is indeed compatible with equilibrium play. We impose conditions which make sure that in any equilibrium (i) the editor asks for the referee's recommendation if $s_e \in \{\emptyset, g\}$, but not if $s_e = b$, and (ii) desk-rejects the manuscript upon observing a bad signal and follows the referee's advice if he acquires it. The purpose of this section is not to identify the entire set of parameter constellations for which an equilibrium with the above properties exists, but to show that it is non-empty. In particular, for parameters that satisfy the conditions identified in this section, the above prescribed behavior is optimal for the editor independently of his beliefs with respect to the mediocre referee's reporting behavior and independently of his own behavior off the equilibrium path (i.e., what he would have done had he not consulted the referee). As a consequence, the identified conditions are stronger than necessary to guarantee the existence of an equilibrium with the above editor behavior. Nevertheless, the identified set of parameter constellations is rich enough to show in Section 2.3 that inefficient preselection rules as identified in Proposition 5 (but taking into account the cost of information acquisition) may indeed be used in equilibrium. We start our analysis with the editor's publication behavior.

The editor's publication behavior

The editor maximizes his expected payoff taking all the available information into account. He decides to publish the paper if the probability that it is good exceeds a threshold. Under Condition 1, without access to the referee's recommendation, the probability that the paper is good does not pass this threshold.

Condition 1 $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{1-\varepsilon}{\varepsilon}$.

The following condition makes sure that the referee's recommendation to publish the paper convinces the decision maker to do so.

Condition 2 $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} \leq \frac{\varepsilon}{(1-\varepsilon)(1-q)}$.

The interval with the lower bound defined by Condition 1 and the upper bound given by Condition 2 exists if q is high relative to ε.

Remark 1 *Conditions 1 and 2 can be satisfied simultaneously if* $q > \frac{1-2\varepsilon}{(1-\varepsilon)^2}$.

Lemma 5 *Assume that Conditions 1 and 2 are satisfied. In any equilibrium in which the smart referee is sincere, the editor behaves as follows: In case that he does not ask for the referee's recommendation, the editor rejects the paper. He follows the referee's recommendation if he has asked for it.*

If the editor's preferences satisfy Conditions 1 and 2, he is biased towards rejection in the sense that the paper's expected quality has to be sufficiently high to be accepted for publication. In particular, if he observes a positive signal on his own, this alone does not suffice to convince him to publish the paper. The referee's recommendation to accept the paper convinces him to do so.[13] Note that under Condition 1, the editor acquires own information at a preselection stage only if he anticipates to condition the choice whether to ask for the referee's recommendation on this information.[14] If he acquires own information, he does so for the sake of reducing the expected cost of refereeing, which also seems to be the major intention when introducing real-world desk-rejection policies.

The editor's information acquisition behavior

Next, we make sure that it is optimal for the editor to ask for the referee's advice if he has observed a good signal or no information, and not to ask for the referee's advice if he has observed a bad signal. Assume in the following that Conditions 1 and 2 are satisfied.

Without acquiring the referee's advice, the editor rejects the paper with probability one. Asking for the referee's advice is equivalent to delegating the decision to the referee. If the manuscript is in fact good, asking for the referee's recommendation and then following the advice has the effect that the manuscript is published at least with probability q rather than

[13] In particular, under the stated conditions, even if the editor has observed a bad signal himself, the referee's recommendation outweighs this information (and that is true even if the editor has the most pessimistic beliefs regarding the reporting behavior of the mediocre referee). That requires, of course, that the editor's signal is relatively weak compared to q, the (minimum) probability that the referee's recommendation to accept the paper is correct.

[14] Note also that, if the editor was allowed to acquire information after the referee has submitted his report, he would not want to use this option (assuming that he can acquire a signal at most once).

with probability zero.[15] Hence, refereeing yields a redistribution of probability mass from the event of rejecting a good paper to the event of publishing it, which is good. Accordingly, if the manuscript is in fact bad, delegating the decision to the referee has the effect that the manuscript is rejected with a probability in between q and 1 instead of with probability one. Hence, refereeing may also yield a redistribution of probability mass from the event of rejecting to the event of publishing a bad paper, which is bad.[16] The overall effect of refereeing (net of the cost) on the editor's payoff is positive, as it is optimal for him to make the publication choice contingent on the information. Hence, if refereeing is costless, the editor always asks for the referee's advice. The value of refereeing depends on the editor's information.

The extent of the redistribution of probability mass induced by refereeing is independent of the editor's information. However, the positive effect on the probability to publish the paper if it is good occurs with a higher probability if the editor has not yet observed an opposing signal, and, accordingly, the negative effect occurs with a lower probability. Hence, given the publication behavior, refereeing is unambiguously more valuable if the editor does not possess information than if he observed a bad signal (and it is even more valuable if the editor observed a good signal).

Hence, for any $(\hat{\gamma}, \hat{\beta})$, there exist intermediate values for c such that it is optimal for the editor to ask for the referee's recommendation if $s_e \in \{g, \emptyset\}$, and to desk-reject the manuscript if $s_e = b$. Moreover, if Condition 3 below holds in addition to Conditions 1 and 2, c assumes a value that ensures that the editor behaves as described above in any equilibrium (i.e., for all $\hat{\gamma}, \hat{\beta}$).

Condition 3 (A) $p < 1 - \varepsilon$, (B) $\frac{u_{rb} - u_{ab}}{u_{ag} - u_{rg}} < \frac{\frac{1}{2} - \varepsilon q}{\frac{1}{2} - \frac{1}{2} q}$, (C) $q\varepsilon(u_{ag} - u_{rg}) \leq c \leq \frac{1}{2}(u_{ag} - u_{rg}) - \frac{1}{2}(1 - q)(u_{rb} - u_{ab})$.

Proposition 6 *Let Conditions 1 and 2 be satisfied.*
(i) There is a parameter range for the cost of refereeing c such that an equilibrium exists in which the editor desk-rejects the manuscript if he observes a negative signal, and asks for the referee's recommendation if he observes a positive signal or no information. He follows the referee's advice.
(ii) In addition, let Condition 3 be satisfied. In any equilibrium, the editor desk-rejects the manuscript if he observes a negative signal, and asks for the referee's recommendation if he observes a positive signal or no information. He follows the referee's advice.

Remark 2 $q > \frac{\frac{1}{2} - \varepsilon}{\frac{1}{2}(1 - \varepsilon) - \varepsilon^2}$ *guarantees that Condition 1 and Condition 3B can be satisfied simultaneously and implies Remark 1.*

[15]The probability of accepting a good paper and rejecting a bad paper, respectively, depends on the mediocre referee's reporting strategy. As the smart referee gives truthful advice, both are at least q.

[16]This effect occurs if and only if the mediocre referee sometimes recommends to publish the paper, i.e., $\gamma > 0$.

Given that the editor uses the information acquired at the preselection stage to condition his further information acquisition choice on it, the information has a value. If his cost of information acquisition ξ does not exceed this value, it is optimal to acquire the information. Conditions 1 and 2 and the assumption that ξ is sufficiently small are sufficient for the existence of an equilibrium in which the editor adopts a desk-rejection policy.

2.3 Equilibria and efficiency

In this section, we show that applying an inefficient preselection rule, in the sense that the editor's expected payoff is higher if he commits not to preselect, can be the unique equilibrium outcome. We assume that the referee cannot observe the editor's own information acquisition choice. First, we complete our equilibrium analysis by studying the set of equilibria that can arise for a given cost of information acquisition ξ, maintaining the assumption that Conditions 1-3 are satisfied. The set of equilibria is identified in the following proposition.

Proposition 7 *Let Conditions* $1-3$ *be satisfied.*
(i) If the editor's information acquisition cost ξ *is sufficiently small, then there exists a unique equilibrium in which the editor always acquires information.*
(ii) If the editor's information acquisition cost ξ *is sufficiently large, then there exists a unique equilibrium in which the editor never acquires information.*
(iii) For intermediate values of the editor's information acquisition cost ξ*, there exists an equilibrium in which the editor sometimes acquires information. This equilibrium coexists with two other equilibria: one in which the editor always acquires information, and one in which the editor never acquires information.*

Figure 2 illustrates the editor's expected utility when acquiring information, EU_{IA}, and when not acquiring information, EU_{NIA}, for $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$.[17] On the horizontal axis, we display the editor's belief that the mediocre referee recommends a publication, ranging from zero ($\gamma = 0, \beta = 1$) to one ($\gamma = 1, \beta = 0$).[18] Denote with $\gamma(\eta = 1, \varepsilon), \beta(\eta = 1, \varepsilon)$ the referee's equilibrium strategy induced by ε and $\eta = 1$. If the editor's information acquisition cost ξ is low, taking as given the referee's reporting behavior, the editor gains from acquiring information, because this allows him to sometimes save the refereeing cost c. As a downside, he sometimes forgoes the opportunity to change his attitude towards the manuscript. As the referee anticipates the editor's information acquisition, the mediocre type recommends a publication of the manuscript more often. Therewith, the referee's recommendation for publication becomes less valuable. This in turn increases the attractiveness of preselection for the editor, because now its downside

[17] Note that this is compatible with Conditions 1-3 if q is high enough.

[18] Note that it may not be possible to induce every reporting strategy. If $q < \overline{q}(p)$, $\gamma > 0$ for all η. Likewise, $\gamma = 1, \beta = 0$ can be induced only if $\varepsilon = 0$. If $\varepsilon > 0$, then $\beta > 0$.

Figure 2: The editor's expected utility when acquiring information (EU_{IA}) and when not acquiring information (EU_{NIA}) for $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$.

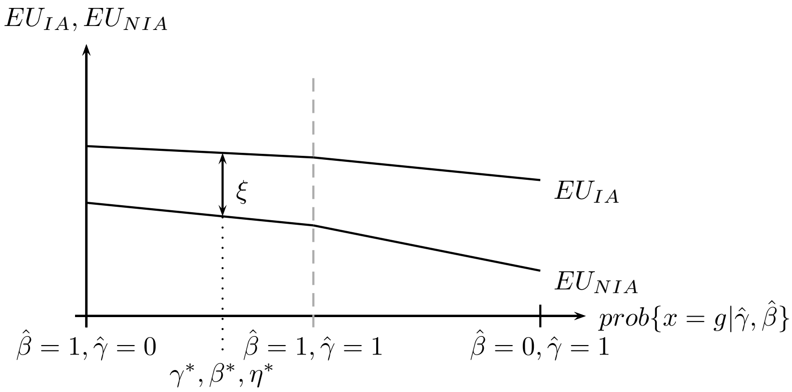

is less severe as a forgone attitude change would be based on less precise information. Thus, the incentive to acquire information is reinforced by the referee's strategic adjustment. Hence, for low values of ξ, the equilibrium in which the editor always acquires information is unique. Although the value of preselection increases the more often the mediocre referee recommends a publication, the editor's ex ante expected utility decreases if $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$.[19] Hence, preselection has the negative side-effect to change the referee's behavior in an unfavorable way. For intermediate values of ξ, there exists an equilibrium with $0 < \eta < 1$. Note, however, that this equilibrium is unstable. For high ξ, the editor never acquires information in equilibrium.

Having identified the set of equilibria, we continue with a discussion of the efficiency of equilibrium outcomes. Our efficiency criterion is the editor's ex ante expected utility.[20] Our assumptions with respect to the editor's preferences imply that he is biased towards rejecting the paper. Whether or not he acquires own information, his ex ante expected utility is unambiguously falling in the probability with which the referee recommends a publication upon observing a bad signal ($\beta \downarrow$). With respect to the effect of γ on the editor's payoff, we can distinguish two parameter ranges for the editor's preferences, as summarized in Table 2 (and depicted in Figure 2 for the second range).

If the editor's bias towards rejection is moderate, i.e., if $\frac{p}{1-p} < \frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} < \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$, and if his own signal is not too strong in the sense that preselection induces $\beta(\eta = 1, \varepsilon) = 1$, then the referee's strategic adjustment to higher preselection levels is beneficial for the editor. If his information acquisition cost ξ is low, then the unique equilibrium outcome is efficient. Likewise,

[19] If the editor is less biased towards rejection, he gains from an increase in γ, see Table 2.

[20] If the referee obtains no transfer (i.e., if the refereeing cost c arises due to searching) and incurs no cost of refereeing, then maximizing the editor's expected utility is equivalent to maximizing ex ante expected welfare. The referee's ex ante expected utility is the ex ante expected editor's belief that he is smart, which is always equal to the prior q.

Preference parameter range		Marginal effects of an increase in γ
$\frac{p}{1-p} < \frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} < \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$	:	$EU_{IA}\uparrow, EU_{NIA}\downarrow, EU_{IA}-EU_{NIA}\uparrow$
$\frac{p(1-\varepsilon)}{(1-p)\varepsilon} < \frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}}$	:	$EU_{IA}\downarrow, EU_{NIA}\downarrow, EU_{IA}-EU_{NIA}\uparrow$

Table 2: Effects of an increase in γ on the editor's expected utility EU_{IA}, EU_{NIA}

if ξ is so high such that the equilibrium without preselection is unique, the equilibrium outcome is efficient.

If $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$, the referee's strategic adjustment of his reporting strategy is detrimental for the editor (see Figure 2), and may even dominate the positive effect of reducing the expected refereeing cost. The editor would like the referee to believe that there are no desk-rejections. However, if ξ is low, the referee anticipates that the editor preselects with probability one. The editor would be better off if he could commit not to use the option of preselection.

Proposition 8 *Consider ξ small enough such that $\eta = 1$ in the unique equilibrium. There is a parameter range for the editor's preferences, for q' and an associated $c(q')$ such that if $q = q'$ and $c < c(q')$ the editor attains a strictly higher expected utility than his equilibrium payoff if he is able to commit to $\eta = 0$.*

If two imperfect positive signals (the editor's and the mediocre referee's) do not suffice to convince the editor to publish the paper, then he prefers the mediocre referee to reject the paper independently of his private information. That way, the paper is published only if it is recommended for publication by the smart referee, i.e., if it is good. Preselection allows the editor to save the cost of refereeing in the case that it is relatively unlikely that the paper will be recommended for publication. However, preselection induces the mediocre referee to recommend a publication with a positive probability. If the effect on the decision quality is strong relative to the cost advantage of preselection, the editor is better off if he can commit not to preselect.

2.4 Endogenous participation

In many applications, including the refereeing process, the expert may refrain from providing advice. We therefore point out effects of the preselection rule on the referee's participation behavior. For this analysis, we take the preselection rule as exogenously given and derive the equilibrium composition of the pool of participating referees. For low participation costs, there is no effect of the preselection rule on participation behavior. Hence, our previous analysis allows for endogenous participation under the presumption that the participation cost is low.

We will discuss alternative cost and timing structures. To begin with, we assume that prior to receiving his signal about the manuscript's quality, the referee has the possibility to decline the editor's invitation to review the manuscript. Denote by k the cost that the referee faces when agreeing to review the paper. The editor takes the referee's participation strategy into account when updating his belief about the referee's type. We denote with q^{part} the probability that a participating referee is smart. There always exists an equilibrium in which the referee never participates, because the out-of-equilibrium-event of observing participation is associated with adverse out-of-equilibrium-beliefs. We will ignore this type of equilibrium in the following. We continue to restrict attention to equilibria in which the smart type gives sincere advice. In equilibrium, there is participation only if the smart type participates. Otherwise, a mediocre referee would reveal his mediocrity and incur a cost, and would be strictly better off pooling with the smart type by not participating. Moreover, in equilibrium, the smart type participates with a higher probability than the mediocre type. As a consequence, the probability that a participating referee is smart is weakly higher than the ex ante probability that he is smart.

Lemma 6 $q^{part} \geq q$.

Lemma 7 describes the equilibrium participation behavior depending on the participation cost.

Lemma 7 *Consider a preselection rule* (η, ε).
(i) The mediocre referee participates with a positive probability if and only if the smart referee always participates.
(ii) There is an equilibrium in which the smart referee always participates if $k < 1$.
(iii) There is an equilibrium in which the smart referee does not participate if $k \geq 1 - q$.
(iv) There is an equilibrium in which the smart referee participates with a probability in between zero and one if $1 - q < k < 1$.
(v) Parts (ii)-(iv) describe the entire set of equilibria for $k < 1$.
(vi) For $k > 1$*, there is a unique equilibrium in which the referee never participates.*

If $k \in [1 - q, 1)$, we face multiple equilibria. We focus on the equilibrium in which the mediocre referee participates with a positive probability, because this is the only one in which the preselection rule has an impact on the referee's behavior. There is no direct effect of the preselection rule on the smart type's participation behavior.[21] The smart type's expected payoff is affected by the preselection rule only if the mediocre type participates. If the mediocre type participates, then a change of the preselection rule may change his participation and/or reporting behavior. However, in spite of an impact on the smart type's expected payoff, there is no incentive for the smart type to change his participation or reporting behavior. Hence, in the following we focus on the mediocre type's participation behavior.

The mediocre type's incentive to participate is driven by the opportunity to pool with the smart type. Without preselection, recommending a rejection is a safe haven alternative,

[21] An indirect effect may arise from the possibility that the preselection rule may be responsible for equilibrium selection.

which ceteris paribus becomes less attractive if the editor preselects. Hence, if the referee's recommendation is biased towards rejection, preselection is detrimental for him and reduces his participation. If the preselection rule induces a bias towards accepting the paper, participation becomes more attractive the tougher the preselection rule is, that is the higher the chances are that this advice is correct. The tougher the preselection rule, the better the mediocre type can predict, and hence imitate, the smart referee's reporting behavior. We state this observation formally for particularly illustrative parameter constellations, for which the preselection rule does not impact on the referee's reporting strategy.[22]

Proposition 9 *(i) Consider $\varepsilon > 0$. If $q > \frac{p(1-\varepsilon)}{\varepsilon^2(1-p)+p(1-\varepsilon)}$, then increasing the preselection level reduces the mediocre referee's participation and has no effect on reporting behavior.*
(ii) There are ε' sufficiently low and $\eta' < 1$ sufficiently high such that if $\varepsilon \leq \varepsilon'$ and $\eta \geq \eta'$, then increasing the preselection level increases the mediocre referee's participation.

Under the conditions stated in the Proposition 9(i), conditional on receiving the referee's advice, the editor is strictly better off with a higher level of preselection. The probability that a recommendation to publish the paper is correct is unaffected, and the probability that a recommendation to reject the paper is correct is higher. Note that the increase in the probability of receiving no advice does not hurt the editor, because it is only uninformative advice which occurs with a lower probability. Hence, conditional on entering the refereeing stage, the editor is better off with a tougher preselection stage. The editor has to trade off these effects against a mistaken judgment at the preselection stage.

In the following, we consider an alternative timing: Suppose that the referee can decide again whether he wants to participate after having observed the signal s_r (but already having sunk k). The picture does not change. Because any report yields a strictly positive expected payoff and non-participation yields zero utility, the referee does not change his mind after the realization of his signal.

If k (or another cost) arises at the reporting stage (say it is costly to read the paper, but also to write a report), then if the reporting cost is high enough such that full participation is not optimal, participation is contingent on the signal. The mediocre referee participates less often with the signal that is less favorable for his reporting strategy. For instance, if his report is biased towards rejecting the paper, participation is more valuable if his signal indicates bad paper quality. Hence, he participates less often having observed a good signal than having observed a bad signal. Note that if participation is conditional on the signal, then the decision not to participate contains information that the editor can use for updating his assessment of the paper's quality. The direction of the effect of the preselection level on

[22]The referee can react to a change of the preselection rule at the participation stage and/or at the reporting stage. Increasing the preselection level ceteris paribus reduces (increases) the mediocre referee's expected utility if the probability that his advice is correct decreases (increases).

participation behavior depends on its impact on the mediocre referee's possibility to pool with the smart type. Hence the arguments above apply, except that the adoption of participation behavior to a new preselection rule is more sophisticated. The qualitative effect on the pool of participating referees is the same. Due to the possibility to condition choices on more information, there are additional effects on the editor's information as sketched below. Consider again the case that q and ε are sufficiently high such that $\gamma = 0$ for any preselection level and any participation behavior. Suppose that the cost for writing a report is such that under the presumption of full participation, the mediocre referee endowed with a good signal is just indifferent between participation and non-participation. Increasing the preselection level strictly decreases his expected utility from participation. He will react with reducing participation upon the observation of a good signal. The editor enjoys an increased report quality as he faces the smart type more often. In addition, the mediocre referee's report becomes "more truthful" in the sense that conditional on reporting, the probability that the report contains the truth increases. Finally, the editor is able to retrieve the mediocre referee's good signal from the fact that he does not participate. These insights suggest that if the participation cost arises predominantly from writing a report, then it may be beneficial to allow the referee to condition his participation choice on information, i.e., to grant access to the paper prior to agreeing to review. Moreover, raising the reporting cost (e.g., by demanding a detailed statement) deters the mediocre type's participation and therewith allows a better usage of his information.

Lastly, suppose reading the paper is costly but it is possible to formulate a report without incurring a cost. Then, k is effectively an information acquisition cost. In an equilibrium in which the smart type participates, non-participation is dominated by participation and giving uninformed advice. As pointed out in Suurmond et al. (2004), acquiring information gives the smart type the opportunity to separate from the mediocre type. He will hence use this opportunity if it is not too costly. The mediocre type acquires information only if it has a value for him, i.e., if he plans to report truthfully.[23] In this case, the preselection rule impacts only on the information acquisition behavior. Again, the direction of the effect depends on the level of preselection and the associated reporting behavior. If for a given preselection rule, the mediocre referee does not acquire information and recommends a rejection, then increasing the preselection level increases the incentive to acquire information.[24] However, increasing the preselection level too much reduces the incentive to acquire information as the referee's inference that the manuscript is good becomes so strong that he always recommends a publication.

[23] In Suurmond et al. (2004), the mediocre ("dumb") type cannot acquire useful information.

[24] The expected utility from recommending a publication given $s_r = g$ increases, whereas the expected utility from recommending a rejection decreases. The value of truthtelling relative to the value of unconditionally recommending a rejection increases. If the value of unconditionally recommending a rejection stays above the value of unconditionally recommending a publication, then the value of information increases.

3 Discussion

Preselection is a prominent mechanism in decision procedures. Often, preselection rules are applied in order to save information acquisition costs at later stages.

Cherkashin et al. (2009) quantify the effect of a hypothetical introduction of a preselection stage at the Journal of International Economics. They consider desk-rejecting papers solely based on observable information, like the rank of the author's university. They find that if the 40% "worst" papers are desk-rejected, then only 8% are wrongfully rejected compared to the actual journal decision, using a refereeing process. They conclude that preselecting even without looking at the paper does not have a significant effect on the decision quality.

A preselection rule that is exclusively based on public information as applied in the analysis of Cherkashin et al. (2009) does not affect the assessment of the quality of a paper which enters the refereeing stage and should not influence the referee's behavior. However, if additional private information is acquired at the preselection stage, the referees' reaction to a change of the preselection policy may have an additional effect on the decision quality that should be taken into consideration when implementing a desk-rejection rule. As pointed out in the present paper, a desk-rejection policy based on the editor's expertise may increase the probability that a paper is wrongfully accepted.

Apart from the academic publishing process, our analysis applies to many other decision problems. The crucial features of our model are (i) that information-processing agents care about their reputation for expertise and (ii) that some decisions are more easily identified as being wrong than others. In such a decision environment, agents may shy away from revealing information which is likely to induce a decision that hurts their reputation if it turns out to be wrong. The introduction of a preselection stage (or of an additional information processing layer) may affect the agent's reporting behavior at later stages and may also have consequences for participation and information acquisition behavior.

APPENDIX

Proof of Lemma 1

Suppose $\gamma < 1, \beta < 1$ in equilibrium.
$\gamma < 1$ requires:[25]

$$prob\{\omega = g|s_r = g\}\pi_{gg}(\gamma,\beta) \leq \pi_{b\emptyset}(\gamma,\beta) - prob\{s_e = g|s_r = g\}\left(\pi_{b\emptyset}(\gamma,\beta) - \pi_{bg}(\gamma,\beta)\right). \quad (1)$$

$\beta < 1$ requires:

$$prob\{\omega = g|s_r = b\}\pi_{gg}(\gamma,\beta) \geq \pi_{b\emptyset}(\gamma,\beta) - prob\{s_e = g|s_r = b\}\left(\pi_{b\emptyset}(\gamma,\beta) - \pi_{bg}(\gamma,\beta)\right). \quad (2)$$

Because $prob\{\omega = g|s_r = g\} > prob\{\omega = g|s_r = b\}$, the LHS of (1) is greater than the LHS of (2), and because $prob\{s_e = g|s_r = g\} > prob\{s_e = g|s_r = b\}$ and $\pi_{b\emptyset} > \pi_{bg}$, the RHS of (1) is smaller than the RHS of (2). Hence, both conditions cannot be satisfied simultaneously.

Proof of Lemma 2

Denote with $prob\{x = g|\hat{\gamma},\hat{\beta}\}$ the probability that the editor assigns to the event that the mediocre referee recommends a publication. Note that $prob\{x = g|\hat{\gamma} = 0, \hat{\beta} = 1\} = 0$ and $prob\{x = g|\hat{\gamma} = 1, \hat{\beta} = 0\} = 1$. In Figures 3 and 4, $prob\{x = g|\hat{\gamma},\hat{\beta}\}$ increases gradually from 0 to 1, first by fixing $\beta = 1$ and gradually increasing γ, and then fixing $\gamma = 1$ and gradually decreasing β.

Suppose $s_r = g$. The mediocre referee's expected utility from recommending publication, $prob\{\omega = g|s_r = g\}\pi_{gg} = prob\{\omega = g|s_r = g\}\frac{q}{q+(1-q)(\hat{\gamma}p+(1-\hat{\beta})(1-p))}$, is strictly decreasing in $\hat{\gamma}$, increasing in $\hat{\beta}$, and hence decreasing in $prob\{x = g|\hat{\gamma},\hat{\beta}\}$. Similarly, the expected utility from recommending a rejection is strictly increasing in $prob\{x = g|\hat{\gamma},\hat{\beta}\}$. Hence, there exists at most one equilibrium with $0 < \gamma < 1$. Such an equilibrium exists iff $\mathbb{E}[\pi|x = g, s_r = g, \hat{\gamma}, \hat{\beta}]$ and $\mathbb{E}[\pi|x = b, s_r = g, \hat{\gamma}, \hat{\beta}]$, intersect in the first region, i.e. to the left of $(\hat{\beta} = 1, \hat{\gamma} = 1)$. Figure 3 illustrates an example of this case.
Analogously, there exists at most one equilibrium with $0 < \beta < 1$. An equilibrium with $\beta \in (0,1)$ exists iff $\mathbb{E}[\pi|x = g, s_r = b, \hat{\gamma}, \hat{\beta}]$ and $\mathbb{E}[\pi|x = b, s_r = b, \hat{\gamma}, \hat{\beta}]$ intersect in the second region, i.e. to the right of $(\hat{\beta} = 1, \hat{\gamma} = 1)$. We need to show that equilibria with $0 < \gamma < 1$ and equilibria with $0 < \beta < 1$ do not coexist in order to complete the proof.

Suppose there exists an equilibrium with $\gamma < 1$. The expected utility from recommending publication is higher in case $s_r = g$ than in case $s_r = b$, and the expected utility from recommending a rejection is higher in case $s_r = b$ than in case $s_r = g$ (if $\eta = 0$, as in Figure 4, they are the same). Hence, $\mathbb{E}[\pi|x = g, s_r = b, \hat{\gamma}, \hat{\beta}]$ and $\mathbb{E}[\pi|x = b, s_r = b, \hat{\gamma}, \hat{\beta}]$ intersect to the left of

[25] The referee's assessment $prob\{.\}$ is conditional on being asked a report (and hence conditional on (η, ε)). We suppress this information for the sake of better readability.

Figure 3: Expected utility when recommending publication (dashed), recommending rejection (dotted), upon observing $s_r = g$ given beliefs $(\hat{\gamma}, \hat{\beta})$, $p = 0.55$, $q = 1/2$, $\eta = 0$.

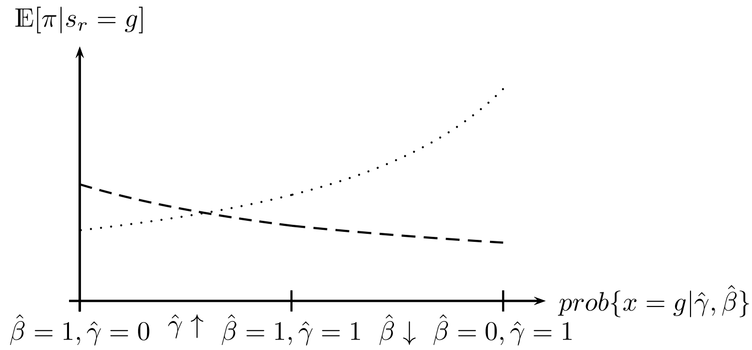

Figure 4: Expected utility when recommending publication upon observing $s_r = g$ (black, dashed), $s_r = b$ (gray, dashed), recommending rejection, upon observing $s_r = g$ (black, dotted), $s_r = b$ (gray, solid) given beliefs $(\hat{\gamma}, \hat{\beta})$, $p = 0.55$, $q = 1/2$, $\eta = 0$.

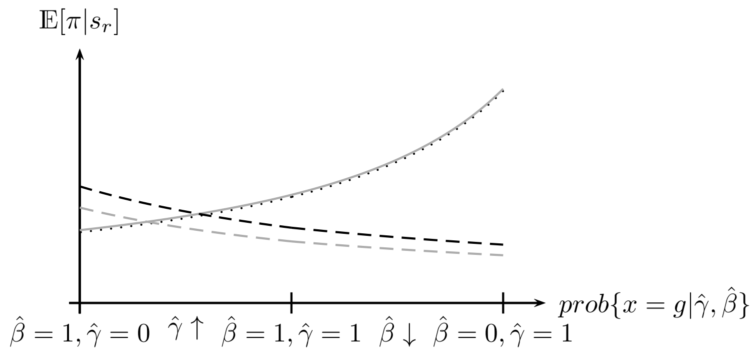

the intersection between $\mathbb{E}[\pi|x = g, s_r = g, \hat{\gamma}, \hat{\beta}]$ and $\mathbb{E}[\pi|x = b, s_r = g, \hat{\gamma}, \hat{\beta}]$, i.e. in the region where $\beta = 1$. It follows that there exists no equilibrium with $\beta < 1$. Figure 4 illustrates the argument.

Suppose there exists an equilibrium with $\beta < 1$. As $\mathbb{E}[\pi|x = g, s_r = g, \hat{\gamma}, \hat{\beta}]$ and $\mathbb{E}[\pi|x = b, s_r = g, \hat{\gamma}, \hat{\beta}]$ intersect to the right of the intersection between $\mathbb{E}[\pi|x = g, s_r = b, \hat{\gamma}, \hat{\beta}]$ and $\mathbb{E}[\pi|x = b, s_r = b, \hat{\gamma}, \hat{\beta}]$, i.e., in the region where $\gamma = 1$. It follows that there exists no equilibrium with $\gamma < 1$.

Proof of Lemma 3

The referee faces the best incentives to report $x = g$ if the editor assumes $\hat{\gamma} = 0, \hat{\beta} = 1$. In this case, $\pi_{gg} = 1$. Having observed $s_r = g$, reporting $x = g$ yields the expected utility p. If he has observed $s_r = b$, his expected utility from recommending $x = g$ is even lower, $1 - p$. Reporting $x = b$ yields utility $\frac{q}{q+2(1-q)}$. Hence, the agent has an incentive to report $x = g$ only if

$$p \geq \frac{q}{q + 2(1-q)} \Leftrightarrow q \leq \frac{2p}{1+p}.$$

Thus, if $q > \frac{2p}{1+p} := \overline{q}(p)$, the referee always reports $x = b$.

Proof of Lemma 4

In any putative equilibrium with $\hat{\gamma} = 1$, the editor's belief that the referee is smart when observing $x = b$, $\pi_{b\emptyset}$, is at least q. When observing $x = g$ and $\omega = g$, the editor's belief that the agent is smart, i.e. π_{gg}, is at most $\frac{q}{1-(1-p)(1-q)}$. $\pi_{b\emptyset}$ is obtained with probability one in the case of a rejection, and π_{gg} is obtained with probability $p < 1$ in the case of a publication. It is easy to verify that $q > \frac{pq}{1-(1-p)(1-q)}$. Hence, there is no equilibrium in which the mediocre referee always reveals a good signal truthfully.

Proof of Proposition 2

Consider the preselection rules (η, ε) and (η', ε'), $\eta' \geq \eta$ and $\varepsilon' \leq \varepsilon$ and denote with (γ, β) and (γ', β') the induced equilibrium reporting behavior, respectively.
$prob\{\omega = g|s_r = g, (\eta, \varepsilon)\} = \frac{p(1-\eta\varepsilon)}{p(1-\eta\varepsilon)+(1-p)(1-\eta(1-\varepsilon))}$ is strictly increasing in η and strictly decreasing in ε. Hence, c.p., the expected payoff from recommending a publication upon the observation of a good signal, $prob\{\omega = g|s_r = g, (\eta, \varepsilon)\}\pi_{gg}$, increases in η and decreases in ε. (Note that the π's depend only on the editor's beliefs, but not on the preselection rule.) Similarly, $prob\{s_e = g|s_r = g, (\eta, \varepsilon)\}$ is strictly increasing in η and strictly decreasing in ε. Hence, c.p., the expected payoff from recommending a rejection upon the observation of a good signal, $\pi_{b\emptyset}(\hat{\gamma} = \gamma, \hat{\beta} = \beta) - prob\{s_e = g|s_r = g, (\eta, \varepsilon)\}(\pi_{b\emptyset}(\gamma \hat{=} \gamma, \hat{\beta} = \beta) - \pi_{bg}(\hat{\gamma} = \gamma, \hat{\beta}) = \beta)$, is strictly decreasing in η and strictly increasing in ε.
Hence, the incentive to recommend a publication upon the observation of a good signal is better the tougher the preselection rule. Three cases can emerge. (i) The incentive is (still) not strong enough. This is the case if $prob\{\omega = g|s_r = g, (\eta', \varepsilon')\}\pi_{gg}(\hat{\gamma} = 0, \hat{\beta} = 1) < \pi_{b\emptyset}(\hat{\gamma} = 0, \hat{\beta} = 1) - prob\{s_e = g|s_r = g, (\eta', \varepsilon')\}(\pi_{b\emptyset}(\hat{\gamma} = 0, \hat{\beta} = 1) - \pi_{bg}(\hat{\gamma} = 0, \hat{\beta} = 1))$ and we have $\gamma' = \gamma = 0$. (ii) The increase in the incentive to recommend a publication translates into an increase in the probability to do so, i.e., $\gamma' > \gamma$. (iii) The incentive was already strong for for preselection rule (η, ε), i.e., $prob\{\omega = g|s_r = g, (\eta, \varepsilon)\}\pi_{gg}(\hat{\gamma} = 1, \hat{\beta} = 1) > \pi_{b\emptyset}(\hat{\gamma} = 1, \hat{\beta} = 1) - prob\{s_e = g|s_r = g, (\eta, \varepsilon)\}(\pi_{b\emptyset}(\hat{\gamma} = 1, \hat{\beta} = 1) - \pi_{bg}(\hat{\gamma} = 1, \hat{\beta} = 1))$. In this case, γ is already maximal and we have $\gamma' = \gamma = 1$.
Analogously, the incentive to recommend a publication upon the observation of a bad signal increases. Hence, we either have $\beta' = \beta = 1$, $\beta' < \beta$ or $\beta' = \beta = 0$.

Proof of Proposition 3

Necessary and sufficient for the existence of an equilibrium in which $\gamma = \beta = 1$ are the following conditions:

$$prob\{\omega = g|s_r = g\}\pi_{gg}(\hat{\gamma} = 1, \hat{\beta} = 1) \quad (3)$$

$$\geq$$

$$\pi_{b\emptyset}(\hat{\gamma} = 1, \hat{\beta} = 1) - prob\{s_e = g|s_r = g\}(\pi_{b\emptyset}(\hat{\gamma} = 1, \hat{\beta} = 1) - \pi_{bg}(\hat{\gamma} = 1, \hat{\beta} = 1)), \quad (4)$$

and

$$prob\{\omega = g|s_r = b\}\pi_{gg}(\hat{\gamma} = 1, \hat{\beta} = 1) \quad (5)$$

$$\leq$$

$$\pi_{b\emptyset}(\hat{\gamma} = 1, \hat{\beta} = 1) - prob\{s_e = g|s_r = b\}(\pi_{b\emptyset}(\hat{\gamma} = 1, \hat{\beta} = 1) - \pi_{bg}(\hat{\gamma} = 1, \hat{\beta} = 1)). \quad (6)$$

From Lemma 4, it follows that for $\eta = 0$, (3) $<$ (4) and (5) $<$ (6). For $\eta = 1, \varepsilon = 0$, (3) $>$ (4) and (5) $>$ (6). For any $\varepsilon < 1/2$, (3) and (5) are monotonically increasing in η, and (4) and (6) are monotonically decreasing in η. There exists an ε' such that if $\varepsilon < \varepsilon'$, there exists $\eta'(\varepsilon)$ such that (3) $\geq$ (4) $\Leftrightarrow \eta \geq \eta'(\varepsilon)$. Moreover, (3) $>$ (5) and (6) $>$ (4) such that (3) $=$ (4) implies (6) $>$ (5). Hence, there exists a range $[\eta'(\varepsilon), \eta''(\varepsilon)]$ such that if $\eta \in [\eta'(\varepsilon), \eta''(\varepsilon)]$, truthful revelation is an equilibrium.

Proof of Proposition 4

Conditional on entering the refereeing process, the probability that a bad paper is recommended for publication is $(1-q)(p(1-\beta)+(1-p)\gamma)$, which is increasing in γ and decreasing in β. From Proposition 2 we know that γ is weakly higher and β weakly lower for tougher preselection rules.

Proof of Proposition 5

It is easy to see that the introduction of a preselection stage yields a (weakly) higher probability that a bad manuscript is accepted. In addition, we identify preselection rules that induce sincere advice, for which the probability that a good manuscript is rejected (either at the preselection stage or after the refereeing) is strictly higher if the preselection rule is used than without preselection. The proof proceeds in four steps. In the first three steps, we focus on preselection rules with $\eta = 1$. Step (i) is to identify a range for ε such that $prob\{x = b|\omega = g\}$ is higher in case of preselection and truthtelling by the mediocre agent (where, with slight abuse of notation, $x = b$ refers to both, the recommendation of a rejection by the referee and the case of a desk-rejection) than in the case without preselection where the mediocre type always recommends a rejection. In Step (ii) we identify the range for ε such that truthtelling is an equilibrium. Step (iii) is to show that the intersection of the previously determined ranges is non-empty. Step

(iv) is to recognize that for some ε in the identified range the conditions for truthtelling and decision quality are not binding. Hence, because of continuity, for $\eta \neq 1$, but close to one, the above argument continues to hold. A formal proof of Step (iv) is omitted.

Step (i) Without preselection, $prob\{x = b|\omega = g\} = 1 - q$. With preselection and truthtelling, $prob\{x = b|\omega = g\} = \varepsilon + (1-\varepsilon)(1-q)(1-p)$. We have $p < \frac{q}{q+2(1-q)}$, hence $\varepsilon + (1-\varepsilon)(1-q)(1-p) > \varepsilon + (1-\varepsilon)(1-q)\frac{2(1-q)}{q+2(1-q)}$. Preselection yields a lower decision quality if

$$\varepsilon + (1-\varepsilon)(1-q)\frac{2(1-q)}{q+2(1-q)} > 1-q \tag{7}$$

$$\Leftrightarrow \varepsilon > \frac{1}{\frac{1}{1-q}+2}. \tag{8}$$

Step (ii) Recommending publication after observing a positive signal is consistent with equilibrium if

$$\frac{p(1-\varepsilon)}{p(1-\varepsilon)+(1-p)\varepsilon} \cdot \frac{q}{1-(1-p)(1-q)} \geq \frac{q\varepsilon}{q\varepsilon + (1-q)((1-\varepsilon)(1-p)+\varepsilon p)} \tag{9}$$

$$\Leftrightarrow \varepsilon \leq \frac{1}{\sqrt{\frac{q}{p(1-q)}+1}+1}. \tag{10}$$

Recommending rejection after observing a negative signal is consistent with equilibrium if

$$\frac{q\varepsilon}{q\varepsilon + (1-q)((1-\varepsilon)(1-p)+\varepsilon p} \geq \frac{(1-p)(1-\varepsilon)}{(1-p)(1-\varepsilon)+p\varepsilon} \cdot \frac{q}{1-(1-p)(1-q)} \tag{11}$$

$$\Leftrightarrow \varepsilon \geq \frac{1}{\frac{\sqrt{\frac{p}{1-q}-p(1-p)}}{1-p}+1}. \tag{12}$$

It is easy to check that $\frac{1}{\sqrt{\frac{q}{p(1-q)}+1}+1} > \frac{1}{\frac{\sqrt{\frac{p}{1-q}-p(1-p)}}{1-p}+1}$.

Step (iii) It is easy to verify that $\frac{1}{\frac{1}{1-q}+2} < \frac{1}{\sqrt{\frac{q}{p(1-q)}+1}+1}$.

Proof of Lemma 5

The editor maximizes his expected payoff taking all the available information into account. He accepts the paper for publication if and only if:

$$\begin{aligned} & prob\{\omega = g|\text{ editor's info}\}u_{ag} + (1 - prob\{\omega = g|\text{ editor's info}\})u_{ab} \\ \geq\ & prob\{\omega = g|\text{ editor's info}\}u_{rg} + (1 - prob\{\omega = g|\text{ editor's info}\})u_{rb} \\ \Leftrightarrow\ & \frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} \leq \frac{prob\{\omega = g|\text{ editor's info}\}}{1 - prob\{\omega = g|\text{ editor's info}\}}. \end{aligned} \tag{13}$$

Under Condition 1, $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{1-\varepsilon}{\varepsilon}$, the editor prefers rejecting to accepting the paper upon the observation of a good signal, $s_e = g$. Note that Condition 1 implies that the editor also rejects the paper if he has no or bad information, or if the referee recommends to reject the paper.

Upon receiving the referee's recommendation to accept the paper, the editor's assessment of the probability that the paper is of good quality is lowest if his own signal is bad ($s_e = b$) and he believes that the mediocre type always recommends to accept the paper ($\hat{\gamma} = 1, \hat{\beta} = 0$). In this case, we have $prob\{s = g|x = g, s_e = b, \hat{\gamma} = 1, \hat{\beta} = 0\} = \frac{\varepsilon}{1-q+\varepsilon q}$. The editor follows the referee's recommendation to accept the paper for any s_e and any $\hat{\gamma}, \hat{\beta}$, if he does so for $s_e = b$ and $\hat{\gamma} = 1, \hat{\beta} = 0$, i.e. if Condition 2, $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} \leq \frac{\varepsilon}{(1-\varepsilon)(1-q)}$, is satisfied.

Proof of Proposition 6

If the editor does not ask for the referee's recommendation, his expected payoff is:

$$prob\{\omega = g|s_e\}u_{rg} + prob\{\omega = b|s_e\}u_{rb}. \tag{14}$$

If he asks for the referee's advice, the editor's payoff (neglecting the refereeing cost) is:

$$\begin{aligned} &prob\{\omega = g|s_e\}(prob\{x = g|\omega = g\}u_{ag} + (1 - prob\{x = g|\omega = g\})u_{rg}) \\ &+prob\{\omega = b|s_e\}(prob\{x = g|\omega = b\}u_{ab} + (1 - prob\{x = g|\omega = b\})u_{rb}). \end{aligned} \tag{15}$$

Subtracting (14) from (15) yields the value of refereeing:

$$prob\{\omega = g|s_e\}prob\{x = g|\omega = g\}\left(u_{ag} - u_{rg}\right) \tag{16}$$

$$- \quad prob\{\omega = b|s_e\}\left(1 - prob\{x = b|\omega = b\}\right)\left(u_{rb} - u_{ab}\right). \tag{17}$$

In the case that the editor's signal is bad, $s_e = b$, the value of refereeing is:

$$\varepsilon prob\{x = g|\omega = g\}\left(u_{ag} - u_{rg}\right) - (1 - \varepsilon)\left(1 - prob\{x = b|\omega = b\}\right)\left(u_{rb} - u_{ab}\right). \tag{18}$$

If the editor does not have a signal, $s_e = \emptyset$, the value of refereeing is:

$$\tfrac{1}{2}prob\{x = g|\omega = g\}\left(u_{ag} - u_{rg}\right) - \tfrac{1}{2}\left(1 - prob\{x = b|\omega = b\}\right)\left(u_{rb} - u_{ab}\right). \tag{19}$$

If the editor has a good signal, $s_e = g$, the value of refereeing amounts to:

$$(1 - \varepsilon)\, prob\{x = g|\omega = g\}\left(u_{ag} - u_{rg}\right) - \varepsilon\left(1 - prob\{x = b|\omega = b\}\right)\left(u_{rb} - u_{ab}\right). \tag{20}$$

It is easy to see that (18) < (19) < (20) for any γ, β. If $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{\varepsilon p}{(1-\varepsilon)(1-p)}$, which is implied by Condition 1 under Condition 3A, (18) is decreasing in γ. Under Condition 1, both (18) and (19) are decreasing in $-\beta$. If $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} < \frac{\frac{1}{2}-\varepsilon q}{\frac{1}{2}-\frac{1}{2}q}$, i.e. if Condition 3B holds, (18) assumes a lower value at $\gamma = 0, \beta = 1$ than (19) at $\gamma = 1, \beta = 0$, and hence that the smallest value for (19), $\frac{1}{2}(u_{ag} - u_{rg}) - \frac{1}{2}(1 - q)(u_{rb} - u_{ab})$, is larger than the largest value for (18), $q\varepsilon(u_{ag} - u_{rg})$.

The last step is to verify that the editor indeed (sometimes) acquires a signal, i.e. $\eta > 0$. Notice that the editor uses the signal as his further information acquisition decision depends on its realization, hence the information has a value.

Proof of Proposition 7

(i) The editor uses the information for further decision making. Hence, it has a value. If the cost of information acquisition ξ is sufficiently small, the editor always acquires information in equilibrium.

(ii) The value of the editor's information is bounded. If ξ exceeds that value, the editor never acquires information in equilibrium.

(iii) The editor's expected utility before he acquires information (neglecting the cost for this information and taking as given his later information acquisition and publication behavior), is given by:

$$\begin{aligned} EU_{IA} = \quad & prob\{s_e = g\} \cdot \begin{pmatrix} prob\{\omega = g | s_e = g\} \left(u_{rg} + prob\{x = g | \omega = g\}(u_{ag} - u_{rg})\right) \\ +prob\{\omega = b | s_e = g\} \left(u_{ab} + prob\{x = b | \omega = b\}(u_{rb} - u_{ab})\right) \\ -c \end{pmatrix} \\ + \quad & prob\{s_e = b\} \cdot \left(prob\{\omega = b | s_e = b\} u_{rb} + prob\{\omega = g | s_e = b\} u_{rg}\right) \\ = \quad & 1/2 \cdot \begin{pmatrix} (1-\varepsilon)\left(u_{rg} + (q + (1-q)(\gamma p + (1-\beta)(1-p)))(u_{ag} - u_{rg})\right) \\ +\varepsilon \left(u_{ab} + (q + (1-q)((1-\gamma)(1-p) + \beta p))(u_{rb} - u_{ab})\right) \\ -c \end{pmatrix} \\ + \quad & 1/2 \cdot ((1-\varepsilon) u_{rb} + \varepsilon u_{rg}). \end{aligned}$$

If he does not acquire own information, the editor's expected utility (again taking as given his later information acquisition and publication behavior), is given by:

$$\begin{aligned} EU_{NIA} = \quad & 1/2 \left(u_{rg} + (q + (1-q)(\gamma p + (1-\beta)(1-p)))(u_{ag} - u_{rg})\right) \\ + \quad & 1/2 \left(u_{ab} + (q + (1-q)((1-\gamma)(1-p) + \beta p))(u_{rb} - u_{ab})\right) \\ - \quad & c. \end{aligned}$$

We have:

$$\frac{\partial EU_{IA}}{\partial \gamma} = 1/2(1-q)\left((1-\varepsilon)p(u_{ag} - u_{rg}) - \varepsilon(1-p)(u_{rb} - u_{ab})\right) \tag{21}$$

$$\frac{\partial EU_{NIA}}{\partial \gamma} = 1/2(1-q)\left(p(u_{ag} - u_{rg}) - (1-p)(u_{rb} - u_{ab})\right) \tag{22}$$

$$\frac{\partial EU_{IA}}{\partial \beta} = 1/2(1-q)\left(-(1-\varepsilon)(1-p)(u_{ag} - u_{rg}) + \varepsilon p(u_{rb} - u_{ab})\right) \tag{23}$$

$$\frac{\partial EU_{NIA}}{\partial \beta} = 1/2(1-q)\left(-(1-p)(u_{ag} - u_{rg}) + p(u_{rb} - u_{ab})\right) \tag{24}$$

Condition 3A implies that (21) > (22), and Condition 1 implies that (24) > (23). Hence, the difference between the editor's expected utility when acquiring information and when not increases in the probability with which the mediocre type recommends a publication. Hence, for any reporting strategy, there is exactly one value for ξ for which the editor is indifferent between the two options. If it is possible to induce the associated reporting strategy with a mixed information acquisition strategy, then there exists an equilibrium in which the editor uses a mixed information acquisition strategy. To see the coexistence with the equilibria in which the editor always/never acquires information, note again that $EU_{IA} - EU_{NIA}$ increases in γ and $-\beta$. Hence, if the referee thinks that the editor always/never acquires information, he recommends a publication more/less often, such that the editor strictly prefers to acquire/not to acquire information (see Figure 2).

Proof of Proposition 8

Let Conditions 1-3 be satisfied, let $\frac{u_{rb}-u_{ab}}{u_{ag}-u_{rg}} > \frac{p(1-\varepsilon)}{(1-p)\varepsilon}$, let $q > \frac{2p}{1+p}$ and, for the ease of exposition, let the editor's preferences be given by:

	$\omega = g$	$\omega = b$
accept	1	$(1-\alpha)$
reject	0	1

Consider a putative equilibrium with $\eta = 1, \gamma \in (0,1], \beta = 1$. The equilibrium γ^* follows from the following equation:

$$\frac{p(1-\varepsilon)}{p(1-\varepsilon)+(1-p)\varepsilon}\frac{q}{q+(1-q)p\gamma} = \frac{q\varepsilon}{q\varepsilon+(1-q)(1-(\varepsilon+(1-2\varepsilon)p)\gamma)}.$$

Rearranging yields:

$$\gamma^* = \frac{p(1-\varepsilon)(1-q)-\varepsilon^2 q(1-p)}{(1-q)p(\varepsilon+p(1-2\varepsilon))}.$$

We have

$$\gamma^* > 0 \quad \Leftrightarrow \quad q < \frac{p(1-\varepsilon)}{\varepsilon^2(1-p)+p(1-\varepsilon)}$$

$$\gamma^* < 1 \quad \Leftrightarrow \quad q > \frac{p(1-2\varepsilon)}{\varepsilon^2+p(1-2\varepsilon)}.$$

Consider $q \geq \frac{p(1-2\varepsilon)}{\varepsilon^2+p(1-2\varepsilon)}$, verifying $\gamma^* \leq 1, \beta^* = 1$ in any equilibrium.

The editor's expected equilibrium payoff for $\eta = 1$ is:

$$1/2 \cdot \begin{pmatrix} (1-\varepsilon)\left(u_{rg} + (q+(1-q)(\gamma p+(1-\beta)(1-p)))(u_{ag}-u_{rg})\right) \\ +\varepsilon\left(u_{ab}+(q+(1-q)((1-\gamma)(1-p)+\beta p))(u_{rb}-u_{ab})\right) \\ -c \end{pmatrix}$$
$$+ \quad 1/2 \cdot (1-\varepsilon)u_{rb} + \varepsilon u_{rg})$$

With $\beta = 1$, we have

$$1/2 \cdot \begin{pmatrix} (1-\varepsilon)\left(q+(1-q)\gamma p\right) \\ +\varepsilon\left((1-\alpha)+(q+(1-q)((1-\gamma)(1-p)+p))\alpha\right) \\ -c \end{pmatrix}$$

$$
\begin{aligned}
& + \quad 1/2 \cdot (1-\varepsilon). \\
= & \quad 1/2 \cdot \Big(\ \gamma(1-q)\left((1-\varepsilon)p - \varepsilon(1-p)\alpha\right) + (1-\varepsilon)q + 1 - c \ \Big)
\end{aligned}
$$

The editor's ex ante expected utility when committing to $\eta = 0$ is:

$$
\begin{aligned}
= & \quad 1/2\left(u_{rg} + (q + (1-q)(\gamma p + (1-\beta)(1-p)))(u_{ag} - u_{rg})\right) \\
+ & \quad 1/2\left(u_{ab} + (q + (1-q)((1-\gamma)(1-p) + \beta p))(u_{rb} - u_{ab})\right) \\
- & \quad c. \\
= & \quad 1/2(q + (1-q)(\gamma p + (1-\beta)(1-p))) \\
+ & \quad 1/2\left(1 - \alpha + (q + (1-q)((1-\gamma)(1-p) + \beta p))\alpha\right) \\
- & \quad c.
\end{aligned}
$$

With $\beta = 1$, this yields

$$
\begin{aligned}
& \quad 1/2(q + (1-q)\gamma p) + 1/2\left(1 - \alpha + (1 - \gamma(1-q)(1-p))\alpha\right) - c. \\
= & \quad 1/2\left(\gamma(1-q)\left(p - (1-p)\alpha\right) + q + 1\right) - c.
\end{aligned}
$$

$EU_{IA}(\gamma^*) < EU_{NIA}(\gamma = 0)$

$$
\begin{aligned}
\Leftrightarrow & \quad 1/2 \cdot \Big(\ \gamma^*(1-q)\left((1-\varepsilon)p - \varepsilon(1-p)\alpha\right) + (1-\varepsilon)q + 1 - c \ \Big) < 1/2\,(q+1) - c \\
\Leftrightarrow & \quad \gamma^*(1-q)\left(\varepsilon(1-p)\alpha - (1-\varepsilon)p\right) > c - \varepsilon q
\end{aligned}
$$

Under Condition 3C, the RHS is positive. The LHS is positive as $\alpha > \frac{(1-\varepsilon)p}{\varepsilon(1-p)}$. Hence, $EU_{IA}(\gamma^*) < EU_{NIA}(\gamma = 0)$ if

$$
\gamma^* > \frac{c - q\varepsilon}{(1-q)((\varepsilon(1-p)\alpha - (1-\varepsilon)p))}.
$$

For each q, we find $c \geq q\varepsilon$ such that the above inequality is satisfied.

Proof of Lemma 6

Suppose, to the contrary, $q^{part} < q$. Bayes' Law dictates that $q^{npart} > q$, where q^{npart} is the probability that the editor assigns to the referee being smart in the case that he does not participate. The mediocre type can only imperfectly imitate the smart type's reporting behavior. Hence, conditional on participating, $\mathbb{E}[\pi|smart] \geq q^{part} \geq \mathbb{E}[\pi|mediocre]$. $q^{npart} > q > q^{part} \geq \mathbb{E}[\pi|mediocre]$. Hence, the mediocre type strictly prefers not to participate as he obtains a higher expected payoff and incurs no cost. But then, $q^{part} = 1 > q$, a contradiction.

Proof of Lemma 7

(i) The mediocre referee has a strict incentive not to participate in the cases that the smart referee (a) is indifferent between participation and non-participation, or (b) strictly prefers non-participation. Hence, the mediocre type participates only if the smart type participates.

In an equilibrium in which the smart type participates with probability one, the editor believes that a referee who does not participate is smart with probability zero. Payoff realizations in case of participation are bounded below by zero. Participation yields a strictly positive payoff with a strictly positive probability. Suppose that the mediocre type participates with a probability $\rho > 0$. Participation yields $\mathbb{E}[\pi|mediocre] - k$, with $\mathbb{E}[\pi|mediocre] \to 1$ if the editor's belief that the mediocre type participates goes to zero. Hence, for any $k < 1$, but sufficiently high, there is a ρ that guarantees the mediocre referee's indifference between participation and non-participation. For low k, the mediocre referee strictly prefers to participate. The exact bound on k below which the mediocre referee always participates depends on the preselection rule.

(ii) If the editor believes that the smart referee always participates, non-participation yields zero utility. The smart referee has an incentive to deviate only if $\mathbb{E}[\pi|smart] < k$, with $\mathbb{E}[\pi|smart] = 1$ if the probability that the editor assigns to the mediocre referee's participation is zero. For $k < 1$, $\mathbb{E}[\pi|smart] < k$ only if the mediocre referee participates with a strictly positive probability. This is compatible with equilibrium play only if $\mathbb{E}[\pi|mediocre] \geq k$. As $\mathbb{E}[\pi|smart] \geq \mathbb{E}[\pi|mediocre]$, this implies that $\mathbb{E}[\pi|smart] \geq k$. Hence, there is no incentive for the smart referee to deviate.

(iii) Suppose that the referee never participates. In this case, non-participation yields utility q. Even if the editor thinks that every participating referee is smart, it does not pay to deviate to participation if $k \geq 1 - q$.

(iv) Suppose that the smart type participates with a probability ρ strictly between zero and one and the mediocre type participates with probability zero. We have $q^{part} = 1$. The smart referee is indifferent between participation and non-participation if $1 - k = q(1 - \rho)$. It is possible to find a $\rho \in (0, 1)$ such that the indifference condition is satisfied if $1 - q < k < 1$. As the smart referee's obtains a higher expected utility form participation than the mediocre referee, the mediocre referee indeed participates with probability zero.

(v) If only the smart referee participates, his expected utility from participation is constant in the editor's belief about the participation level, whereas expected utility from non-participation strictly decreases. The mediocre referee participates if and only if the smart type participates with probability one (see part (i) of the Lemma). The smart referee's payoff exhibits a discontinuity at the point where he participates with probability one, but stays above the expected utility that he derives from non-participation for $k < 1$ (see part (ii) of the Lemma). Hence, there are at most three equilibrium participation levels for the smart referee. For any participation behavior of the smart referee, there is a unique optimal participation behavior for the mediocre referee.

(vi) If $k > 1$, then even if the editor believes that every participating referee is smart, there is no incentive to participate.

Proof of Proposition 9

(i) Take a preselection rule (η, ε) as given and suppose that $q > \frac{p(1-\varepsilon)}{\varepsilon^2(1-p)+p(1-\varepsilon)}$, which guarantees that $\gamma = 0$ for full participation for any preselection level.[26] Lemma 6 implies that $\gamma = 0$ for any participation behavior of the mediocre type. Suppose that for $\eta = 0$, we have $q < q^{part} < 1$ in equilibrium. Consider the effect of increasing the preselection level. Suppose first that the participation level does not change. Because the reporting behavior does not change either, the editor's beliefs stay the same, but now the belief π_{bg} occurs with a strictly higher probability and $\pi_{b\emptyset} > \pi_{bg}$ realizes with a lower probability. Hence, the mediocre referee's expected utility from participation decreases. Reducing participation, the π's conditional on participation increase, restoring indifference between participation and non-participation in the new equilibrium at a lower participation level.

(ii) Suppose that the preselection rule in place induces $\beta < 1$ and the mediocre referee participates with a probability $0 < \rho < 1$. The higher η, the higher is the probability that the manuscript is good conditional on being asked a report. As the mediocre referee's report is biased towards publication, ceteris paribus, his expected utility from participation increases. In order to keep the argument simple assume that $\varepsilon = 0$ and that $\eta > 1 - (1-p)q$, which ensures that $\beta = 0$ for the initial level of preselection as well as for any higher level.[27] Facing a higher expected utility from participation, the mediocre referee increases his participation level. The editor's beliefs adjust to restore indifference in the new equilibrium with more participation.

[26] See the proof of Proposition 8.

[27] The probability that the manuscript is good conditional on being asked a report and having observed a bad signal is $\frac{1-p}{1-p\eta}$. The probability that the editor did not observe a signal is $\frac{1-\eta}{1-p\eta}$. When the mediocre type always recommends a publication, $\pi_{gg} = q, \pi_{b\emptyset} = 1$, and, as π_{bg} occurs only off the equilibrium path, it is set to zero. Always recommending a publication is optimal for the mediocre referee if $\frac{1-p}{1-p\eta}q > \frac{1-\eta}{1-p\eta}$, which is equivalent to $\eta > 1 - (1-p)q$. If $\varepsilon \neq 0$, then $\beta > 0$ in equilibrium (see Footnote 12) and needs to be adjusted to the new preselection rule in order to restore indifference at the reporting stage.

References

[1] AVERY, C. and CHEVALIER, J. (1999): "Herding over the career" *Economic Letters*, 63, 327–333.

[2] BENABOU, R. and LAROQUE, G. (1992): "Using privileged information to manipulate markets: insiders, gurus, and credibility", *Quarterly Journal of Economics*, 107, No. 3, 921–958.

[3] BLANK, R. M. (1991): "The effects of double-blind versus single-blind reviewing: experimental evidence from the American Economic Review", *American Economic Review*, 81, No. 5, 1041–1067.

[4] CHEN, Y. (2010): "Career concerns and excessive risk taking", *working paper*

[5] CHERKASHIN, I., DEMIDOVA, S., IMAI, S. and KRISHNA, K. (2009): "The inside scoop: acceptance and rejection at the Journal of International Economics", *Journal of International Economics*, 77, 120–132.

[6] EFFINGER, M. R. and POLBORN, M. K. (2001): "Herding and anti-herding: A model of reputational di!erentiation", *European Economic Review* 45, 385–403

[7] HOLMSTRÖM, B. (1999): "Managerial incentive problems: A dynamic perspective", *Review of Economic Studies*, 66, 169–182.

[8] LABAND, D. N. (1990): "Is there value-added from the review process in economics?: Preliminary evidence from authors", *Quarterly Journal of Economics*, 105, No. 2, 341–352.

[9] LAGERLÖF, J. and FRISELL, L. (2007): "A model of reputation in cheap talk", *Scandinavian Journal of Economics*, 109, No. 1, 49–70.

[10] LEAVER, C. (2009): "Bureaucratic minimal squawk behavior: theory and evidence from regulatory agencies", forthcoming *American Economic Review* 99, No.3, 572–607.

[11] LEVY, G. (2004): "Anti-herding and strategic consultation", *European Economic Review* 48, 503-525.

[12] LEVY, G. (2005): "Careerist judges and the appeals process", *RAND Journal of Economics*, 36, No. 2, 275–297.

[13] MILBOURN, T. T., SHOCKLEY, R. L. and THAKOR, A. V. (2001): "Managerial career concerns and investments in information", *RAND Journal of Economics*, 32, No. 2, 334–351.

[14] MORRIS, S. (2001): "Political correctness", *Journal of Political Economy*, 109, No. 2, 231–265.

[15] OTTAVIANI, M. and SØRENSEN, P (2006): "Professional advice", *Journal of Economic Theory*, 126, 120–142.

[16] OTTAVIANI, M. and SØRENSEN, P (2001): "Information aggregation in debate: who should speak first?", *Journal of Public Economics*, 81, 393-421.

[17] PRAT, A. (2005): "The wrong kind of transparency", *American Economic Review*, 95, No. 3, 862–877

[18] SCHARFSTEIN, D. and STEIN, J. (1990): "Herd behavior and investment", *American Economic Review*, 80, No. 3, 465–479

[19] SOBEL, J. (1985): "A theory of credibility", *Review of Economic Studies*, 52, 557–573.

[20] SUURMOND, G., VISSER, B. and SWANK, O. H. (2004): "On the bad reputation of reputational concerns", *Journal of Public Economics*, 88, 2817–2838.