

---

# Auctions and Peer Prediction for Scientific Peer Review

---

**Siddarth Srinivasan**

School of Computer Science and Engineering  
University of Washington  
Seattle, WA 98195  
sidsrini@cs.washington.edu

**Jamie Morgenstern**

School of Computer Science and Engineering  
University of Washington  
Seattle, WA 98195  
jamiemmt@cs.washington.edu

## Abstract

Peer reviewed publications are considered the gold standard in certifying and disseminating ideas that a research community considers valuable. However, we identify two major drawbacks of the current system: (1) the overwhelming demand for reviewers due to a large volume of submissions, and (2) the lack of incentives for reviewers to participate and expend the necessary effort to provide high-quality reviews. In this work, we adopt a mechanism-design approach to propose improvements to the peer review process. We present a two-stage mechanism which ties together the paper submission and review processes, simultaneously incentivizing high-quality reviews and high-quality submissions. In the first stage, authors participate in a VCG auction for review slots by submitting their papers along with a bid that represents their expected value for having their paper reviewed. For the second stage, we propose a novel peer prediction-style mechanism (H-DIPP) building on recent work in the information elicitation literature, which incentivizes participating reviewers to provide honest and effortful reviews. The revenue raised by the Stage I auction is used in Stage II to pay reviewers based on the quality of their reviews.

## 1 Introduction

Scientific publishing through peer review is widely regarded as the gold standard for disseminating scientific research. Yet, despite its widespread use and acceptance, it is plagued by a number of well-documented problems: willing reviewers are often in short supply and reviews themselves are often of poor quality, giving rise to noisy acceptance decisions (Cortes and Lawrence, 2021; Shah et al., 2018). Studying the history of peer review in sociology, Merriman (2020) writes that while peer review is ‘a practice intended to assess the merits of pieces of scholarship, its features did not arise because they were believed to be an especially apt means of identifying the best work.’ The issues with peer review have become especially salient in machine learning as well as computing more broadly. Computing relies largely on a conference model for publishing and disseminating research, and while the number of submissions to major conferences has exploded, the pool of qualified reviewers has grown more slowly (Sculley et al., 2018b; Shah, 2019; Stelmakh et al., 2020). As these issues are only likely to persist, a solution addressing the problem of low-quality reviews is needed.

In this work, we adopt a mechanism-design approach to tackle two drawbacks of the current system of conference-based publishing in machine learning: (1) a large number of submissions, leading to poor and often arbitrary criteria for desk rejection, and (2) poor incentives for reviewers that make it challenging to find qualified reviewers willing to provide effortful and honest reviews. Our proposed mechanism tackles (1) by requiring authors to participate in an auction for review slots, and tackles (2) via a prediction market-style mechanism (involving predicting peers’ review scores) that compensates reviewers for honest and effortful reviews. The mechanism involves minimal modification to the current peer review system; it requires authors to submit a bid accompanying every

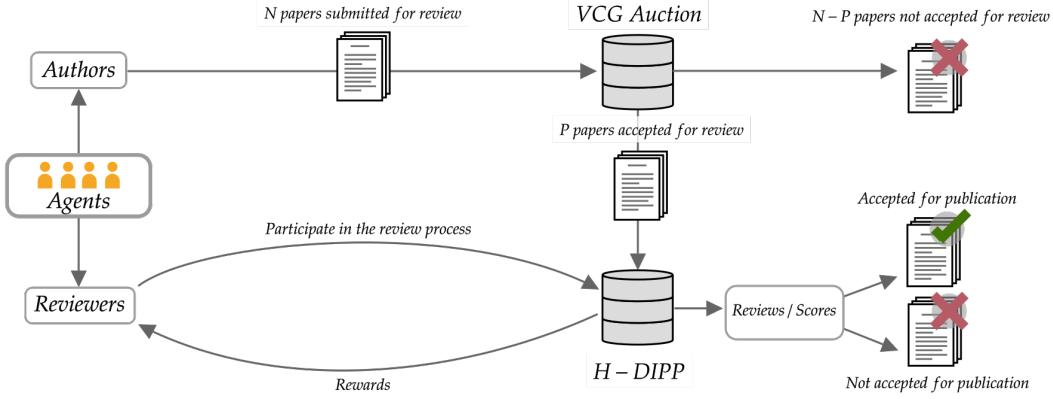


Figure 1: Diagram of proposed mechanism for scientific peer review

paper, and reviewers to participate in a simple prediction game after submitting their reviews. The paper-reviewer matching process and accept/reject decision process may be independently improved or remain unchanged. Our mechanism would adopt a novel currency usable within a consortium of conferences, with the goal of incentivizing authors to provide the much needed service of peer review in order to earn credits exchangeable for having their own papers reviewed. While our proposal is generally applicable to peer review, our perspective is informed by the conference publication model in machine learning. There remain many questions to be answered before our proposed mechanism can be deployed in practice, and we intend for this work to serve as a plausible (but by no means comprehensive) approach to be iterated upon in designing a robust scientific publication process.

Our paper is structured as follows: in section 2, we provide some background on scientific peer review and motivate the problem; in section 3, we present related work; in section 4, we describe our proposed mechanism; in section 5, we discuss some important considerations for a practical implementation of the mechanism, and in Section 6, we conclude with ideas for future work.

**Main Contributions** We make two main contributions: (1) we propose a two-stage peer review mechanism consisting of a submission stage and a reviewing stage, where the revenue raised by the former is used to pay for the latter, and (2) we design the reviewing stage by building on previous work in the information elicitation literature to develop a novel prediction market-style mechanism that incentivizes reviewers to provide honest and effortful reviews.

## 2 Background

Although there are some minor variations, the process of peer review can be summarized as follows: (1) authors submit their work to a conference, (2) some papers receive desk rejection, (3) a paper that is not desk rejected is matched with a small number (3-5) of willing reviewers (who do not know the identity of the authors and vice-versa), (4) the reviewers independently score the paper’s quality for publication based on several criteria such as technical correctness, interest to the community, and quality of exposition (the reviewers may or may not know each others’ identities), (5) a meta-reviewer aggregates the reviewers’ scores and makes an accept/reject decision, incorporating the conference’s targeted acceptance rate into their decision. With this context, we now discuss the often cited primary motivations for and criticisms of peer review below. Merriman (2020) argues that these ‘intellectual’ motivations tend to be post-hoc rationalizations, and the current process of peer review is ‘best understood as the product of continuous efforts to steward editors’ scarce attention while preserving an open submission policy that favors authors’ interests’. Part of our goal then is to develop a peer review mechanism that lives up to the stated ‘intellectual’ motivations for the process.

### 2.1 Motivations for Peer Review

**M1 Screening Results for Validity:** Peer review serves as a community-issued certificate that the approach and results in a publication seem plausible, replicable, and not fraudulent.

- M2 Screening Results for Importance to the Community:** The publication record serves as a central repository of curated results that the field considers important or impactful and wishes to focus on; the conference itself serves to build consensus and host debates on the cutting-edge issues in the field.
- M3 Reputation Building/Signaling:** Publication at selective venues also serves to build reputation and advance careers of authors; the idea is that a selective acceptance process that identifies valid and interesting results also identifies individuals who do good work.
- M4 Feedback and Networking:** Peer review serves as a systematic approach to obtaining feedback on both ongoing and completed work, as well as finding collaborators interested in tackling similar problems.

M1 and M2 are often cited as the primary motivations for peer review, and are arguably the most foundational. While M3 is not often cited explicitly as a motivation, it is well-known that publishing frequently at top-tier venues improves career advancement prospects; peer review clearly serves as a mechanism to identify individuals who should be rewarded. M3 derives its legitimacy from M1 and M2. Meanwhile, M4 is an auxiliary benefit derived from M1 and M2. We argue that the peer review system should be judged on its ability to satisfy M1 and M2 while cognizant of its social functions M3 and M4 so as to avoid unintended consequences.

## 2.2 Criticisms of Peer Review

The general criticism of peer review is that the reviews are often sub-par and the accept/reject decisions are too noisy. Various NeurIPS experiments (Cortes and Lawrence, 2021; Shah et al., 2018; NeurIPS, 2019) provide some empirical basis for this criticism; the 2014 NeurIPS experiment in particular (see Appendix B) found that if the conference was re-run, about a quarter of papers would have their accept/reject decisions swapped, and only about half of accepted papers would be re-accepted. There is also a perception that reviewers are often misled by unsubstantiated claims (Lipton and Steinhardt, 2018) or work that *looks* important but isn't (Vazire, 2017). This may incentivize the submission of premature, lower quality work which if rejected, may be resubmitted at no cost to future conferences with minor tweaks with the goal of getting through a noisy review process (Bengio, 2020). In this context, we discuss some of the explanations for why the current system of peer review within machine learning may be sub-par, although our discussion may apply to peer review in general.

- C1 The Extensive Margin of Reviewing:** There are too many submissions and not enough reviewers, so the current system struggles to match submissions with qualified and interested reviewers. There are no explicit incentives for reviewers, so reviewers participate entirely for pro-social reasons. The shortage of reviewers has been exacerbated by the recent explosion in submissions to top conference venues in machine learning over the past decade. The large number of submissions can also lead to desk-rejection policies that can be somewhat opaque and arbitrary. Poor matching due to a scarcity of reviewers could explain why reviews are often sub-par, so a mechanism that incentivizes reviewers to participate could be valuable.
- C2 The Intensive Margin of Reviewing:** Even when reviewers participate voluntarily, they may *not be willing to exert enough effort* to provide a high-quality assessment of the paper, and there is currently no explicit incentive to do so. In an attempt to redress the shortage of reviewers, some conferences require authors to also participate in review. While this would increase the supply of reviewers, it could potentially crowd-out pro-social motivations and worsen the average quality of reviews without proper incentives. Thus, we may need to provide the appropriate incentives to elicit effortful reviews, especially when requiring participation in the review process.
- C3 Dishonest Reviews:** Some have noted that even when reviewers exert enough effort to understand a paper, they may not report their honest opinion for various reasons. Thus, it is desirable to design the review process to remove obstacles to truthful reviewing.
- C4 Ability of Reviewers:** Another strain of criticism is that peer review is fundamentally flawed and *it cannot work* even when reviewers provide honest and effortful reviews on well-matched papers. There are two versions of this criticism: the first is that reviewing is too subjective for reviewers to agree on results' validity and importance to the community, and the second is that even when reviewers agree, their reviews do not track any meaningful measure of importance to the community. Empirical evidence from the 2014 NeurIPS experiment (where all the

aforementioned incentive issues persist) suggests that reviewers’ scores correlate moderately well with each other as well as with log-citation counts; it may be quite possible to improve upon these results by remedying incentive issues.

Our focus in this paper is to address the first three criticisms head on with a mechanism-design approach. On the paper submission side of the process, our proposed mechanism disincentivizes lower-quality submissions aimed at getting through a noisy review process. On the paper reviewing side, our proposed mechanism incentivizes reviewers to participate in the peer review process, exert effort, and report their assessments truthfully. Our mechanism is agnostic to the reviewer-paper matching process, and we discuss this in Section 4.4. The key challenge we face is the lack of a ground-truth verification for the submitted reviews, which makes it difficult to verify whether reviewers’ assessments are truthful and effortful. We tackle this issue by building on recently proposed mechanisms from the information elicitation without verification (IEWV) literature.

### 3 Related Work

We discuss three bodies of related work: 1) proposals to improve scientific peer review from the machine learning community, 2) general proposals to reform or replace scientific peer review, and 3) the information elicitation without verification literature (Faltings and Radanovic, 2017; Waggoner and Chen, 2013). Readers may also be interested in a broader historical overview of the development of peer review (focused primarily on the social sciences) presented by Merriman (2020).

#### 3.1 Proposals from the Machine Learning Community

A number of proposals to handle the rapid rise in submissions and reviewing load have been put forward by the machine learning community. One idea is to allow some editorial screening, where Area Chairs can reject submissions without full reviews; this would eliminate the need for  $\sim 4\%$  of reviews but some papers may still be screened out unfairly. Another proposal is to introduce submission caps, which limit the number of submissions by a single author, but a cap of 10 submissions would again only eliminate  $\sim 4\%$  of the reviewing load. An analysis of NeurIPS 2019 data (NeurIPS, 2019) concluded that it was unclear how best to rapidly filter papers prior to full review; our proposed submission mechanism uses a VCG auction to determine which papers are accepted for review, where authors signal how ready their work is and the mechanism accepts as many papers for review as it can find quality reviewers for.

Criticizing the slow pace of peer review, LeCun (2012) and Zhao (2012) propose an ‘open-market’, where all papers are uploaded to an online repository as soon as they are ready for viewing. ‘Reviewing entities’ may then identify papers they wish to review, although authors may request reviews as well. However, data from past NeurIPS conferences (Shah et al., 2018; NeurIPS, 2019) suggest that under the current implementation, reviewers do not make many eager bids to review papers, and eager bids are not predictive of eventual acceptance. Thus, mechanisms that rely on reviewers to identify and review papers may need to provide explicit incentives. Identifying the tension between rapid dissemination of ideas and the need to carefully screen results, Bengio (2020) proposes a hybrid arXiv-journal model, where arXiv addresses the former need and journals take on the role of slow, careful screening of results to promote higher-quality publications. In these aforementioned approaches, reviewing incentive issues would still persist (as noted in fields with journal publication models); indeed, the reviewing burden may move from conferences to the open-market system or journals, as authors from less well-known backgrounds seeking community-issued certificates of the quality of their work for reputation-building purposes begin to substitute conference submissions for journal submissions. Our submission mechanism disincentivizes weak submissions that add to the reviewing load while being unlikely to be accepted, and our reviewing mechanism provides direct incentives for honest and effortful reviews.

Other proposals involve improvements to various procedural aspects of peer review, such as a recruitment and mentorship pipeline for junior academics to alleviate the scarcity of reviewers (Stelmakh et al., 2020), optimizing reviewer assignment to obtain less noisy reviews (Shah, 2019) and mitigate conflicts of interest (Xu et al., 2018; Jecmen et al., 2020), and better aggregation of subjective opinions on criteria scores to accept/reject recommendations (Noothigattu et al., 2021). Such ideas are complementary to our reviewing mechanism; they focus on improving paper-reviewer matching or the acceptance decision process, while we focus on incentivizing honest and effortful reviews after reviewers are assigned and before scores are aggregated into accept/reject decisions.

### 3.2 General Proposals to Reform or Replace Peer Review

One strain of research proposes modest changes to the system, such as allowing reviewers to re-purpose reviews from past submissions (Alberts et al., 2008) to reduce the demand for reviews, or improving the publication culture so authors submit more complete and relevant work to the appropriate venues instead of simply targeting the most prestigious venues. Others (Arns, 2014) propose reducing the reviewing burden by allowing some papers (e.g. null results) to be accepted without peer review. However, these proposals do not systematically tackle the previously discussed core incentive issues.

Fox and Petechy (2010) propose an approach similar to ours, where authors must pay in an artificial currency to have their paper reviewed, and reviewers are paid for their completed reviews. However, reviewers receive flat payment for their participation, and so may not provide honest or effortful reviews. Sculley et al. (2018a) propose pecuniary and non-pecuniary compensation for reviewers based on review quality (as judged by a rubric), but rating reviews would increase the workload on area chairs.

In a different approach, de la Rosa and Szymanski (2007b,a) propose entirely *replacing the existing peer review system* with market-based citation auctions, i.e., where earned citations are treated as currency, every submission comes with a “citation bid”, and the papers with the highest bids are accepted for publication. The citations that the paper goes on to earn are used to recoup the bet, with the idea being that the bid will reveal how impactful (as measured by number of citations) a researcher expects their work to be and the mechanism will publish papers expected to be highly cited. They attempt to use citation count to avoid the fundamental problem of verifying review quality, and citations certainly provide some crowd-sourced opinion on the quality of a paper. However, even as it fixes the shortage of reviewers by getting rid of the need for them in the first place, predicting citations is likely extremely high variance, can lead to distorted citation practices, and poses equity issues if highly-cited authors can publish any paper they want by simply outbidding everyone else regardless of the quality of their work.

Prüfer and Zetland (2010) also propose a citation-based paper submission system for journals, although their proposal does not abolish scientific peer review. Instead, they propose a new currency of ‘academic dollars’ used by *editors* of various journals to bid for submitted papers in an auction; if accepted for review, the dollar value is paid to the authors, referees, and editors of *cited* works. This is essentially a citation betting market for journal editors; a paper’s bid will be commensurate with its future expected citation count. Referees will exert effort to provide better feedback to authors to improve the quality of their work in order to earn dollars from future citations, and authors have a marginally higher incentive to write better to earn citations. While an intriguing proposal, we suspect the task of predicting citations is challenging and the singular focus on optimizing for citations may distort citation practices.

Frijters and Torgler (2019) also propose a modified peer review incorporating market-based ideas, with a focus on mitigating ‘gate-keeping’ by insiders and improving transparency. Their proposal includes allowing authors to bid for reviews of varying quality, and opening up the reviews themselves to be rated thereby allowing for ‘professionalization’ of the process. However, the system is somewhat complicated and the market for rating reviews may not be very liquid, especially in technical fields. Xiao et al. (2018) propose a peer review mechanism to elicit effortful reviews by assigning papers to reviewers based on the quality of the authors’ past reviews, but it computes this review quality by asking authors to rate the reviews they received and assuming these ratings to be honest. This is likely too strong an assumption in our setting where authors may simply report how favourable a review was; this would incentivize flattering and likely dishonest reviews, undermining the mechanism.

### 3.3 Information Elicitation without Verification

The IEWV literature is focused on the problem of eliciting truthful private ‘signals’ when the mechanism designer is unable to verify the responses; this may be because the ground truth is difficult or impossible for the mechanism to access or there is simply no ground truth since the responses are subjective. The key application domains of interest include tasks like crowdsourcing data, product ratings, community sensing, and peer grading (Cai et al., 2015; Radanovic et al., 2016; Dasgupta and Ghosh, 2013; Kong and Schoenebeck, 2018). More recent work is also interested in ensuring that private signals are obtained with effort. The typical setting in this literature consists of risk-neutral, utility-maximizing agents who maintain some prior distribution over ‘signals’ they can receive from a mechanism. When the agents participate in the mechanism, the world reveals to them their private information or ‘signal’, and they subsequently form posterior beliefs over the distribution of signals observed by the other agents. In our setting, this would be analogous to reviewers having some prior

belief over the quality of papers they may see, and when presented with a specific paper, the reviewer forms their opinion of that particular paper and develops a posterior belief over other reviewers' assessments of the paper. The task of the mechanism designer is to elicit agents' private signals when verification is not possible by providing incentives for truthful reports.

Two of the foundational approaches to this problem are the peer prediction framework (Miller et al., 2005) and the Bayesian Truth Serum (BTS) (Prelec, 2004). The fundamental idea is to design payments by reasoning about agents' reports and posteriors over other agents' reports in a way that makes truth-telling a Bayesian Nash equilibrium. However, these mechanisms make strong assumptions that are unlikely to hold in practice; subsequent works have relaxed these assumptions (Witkowski and Parkes, 2012a, 2013, 2012b; Radanovic and Faltings, 2013, 2014). There are also *output agreement* mechanisms (Waggoner and Chen, 2014) with weaker assumptions, but these eschew truthfulness and seek to elicit common knowledge. A directly relevant proposal to our work is by Carvalho et al. (2013), which proposes a peer prediction-like mechanism for the same task of scientific peer review. They assume that reviewers' signals on a given paper are drawn from a multinomial distribution, and model reviewers' prior over the multinomials with the conjugate (Dirichlet) prior so inference is tractable, but this is a somewhat restrictive assumption. They further test their mechanism on an experiment with Amazon Mechanical Turk workers, and find that their mechanism does incentivize more accurate reviews and such reviews also produce higher payments for the participating agents. The truthful mechanism applicable in the most general setting is the recently proposed prediction market-style Differential Peer Prediction (DPP) mechanism (Schoenebeck and Yu, 2020). The truth-telling equilibrium maximizes agent welfare, and the mechanism is applicable to a small group of heterogeneous agents. Our proposed reviewing mechanism is a direct extension of this work.

These aforementioned works do not incorporate any model of agent effort, and instead assume that nature or some black box simply reveals the private signal to the agents. The early literature incorporating effort models (Dasgupta and Ghosh, 2013; Witkowski et al., 2014; Radanovic et al., 2016; Liu and Chen, 2016) is limited to the case of binary effort: either agents exert low/no effort, or they exert high/full effort. A binary effort model may be appropriate for quick tasks like those on crowdsourcing platforms, but reviewing technical papers is a time-intensive process, so a more sophisticated effort model is needed. Cai et al. (2015) propose a continuous effort model for data elicitation, where exerting effort gives agents a lower variance signal. However, in our setting, lower effort reviews may not only be noisier, they may also be systematically biased (Kong and Schoenebeck, 2018). Gao et al. (2016) caution that in practice (and as highly relevant in our case), agents have a variety of low-effort signals such as paper length, topic, etc., potentially leading to low-effort coordination into uninformative equilibria. Indeed, Gao et al. (2014) present empirical information showing that agents can successfully coordinate into uninformative equilibria, especially if these equilibria pay more than truthful equilibria.

With these considerations, Kong and Schoenebeck (2018) propose a Hierarchical Mutual Information Paradigm (HMIP) that constructs a hierarchy of criteria that capture a range of cheap to expensive signals, and pays agents the *information gained* from more expensive signals. This work builds on the Mutual Information Paradigm (Kong and Schoenebeck, 2019) which shows that mechanisms incentivize truth-telling when the payment is based on the mutual information between agents' responses, and this is what peer prediction and BTS are attempting to do. We integrate the very general differential peer prediction (DPP) mechanism into the hierarchical mutual information paradigm (HMIP) to develop a strictly truthful prediction market-style mechanism that incentivizes effortful reviews, with the desired equilibria paying more than uninformative equilibria.

## 4 A Mechanism for Scientific Peer Review

We now propose a two-stage mechanism for scientific peer review that incentivizes authors to submit papers they believe are likely to be accepted, and reviewers to participate, exert effort, and honestly review papers. Stage I deals with the paper submission process, where authors bid for review slots, and Stage II deals with the paper review process, where reviewers are paid based on the quality of their reviews. The paper-reviewer matching process and accept/reject decision process may remain unchanged. We envision the mechanism using its own currency (henceforth referred to simply as 'credits') that can be utilized within a consortium of conferences.

**Summary** In the Stage I submission mechanism, authors participate in a VCG auction for review slots, where they submit their papers  $\omega_i$  along with an accompanying bid  $b_i$  (in mechanism-specific

credits). The winners of the auction pay the highest losing bid and have their papers *accepted for review*. In the Stage II reviewing mechanism (which we call H-DIPP), every paper  $\omega_i$  accepted for review is assigned 3 reviewers (denoted  $A$ ,  $B$ , and  $C$ ). Reviewers proceed per the usual peer review process, supplying reviews and scoring the paper on a common, fixed set of  $T$  criteria (e.g. novelty, correctness, interest to the community, quality of writing). The reviewers will then also participate in a simple prediction market-like game to report *predictions* (as probability distributions) of the scores of one other peer, twice for each criterion; reviewers first predict the peer's review score on a given criterion, are shown the third peer's score on that criterion, and are asked to update their prediction. This proceeds sequentially for each criterion in a pre-determined order. Reviewers' predictions of a peer's review scores are rewarded based on their accuracy (with the log proper scoring rule), and reviewers' review scores on each criterion are rewarded based on how much better their peer's second prediction of that review score was compared to their first. We show that reviewers maximize their expected utility under such a payment scheme when they provide honest and effortful reviews. We emphasize that the prediction game involves predicting other reviewers' review scores, and not something like the paper's expected impact. The revenue raised from Stage I of the auction can be used to pay reviewers for their reviews in Stage II. Area chairs then determine which papers are *accepted for publication* per the usual process. Authors who submitted bids must ensure that they have (or can earn) enough credits to cover their bids in Stage I, or risk having their reviews and potential acceptance withheld. In general, authors face the budget constraint of submitting bids no larger than their store of credits they are willing and able to earn through reviewing. Authors who provide the valuable and scarce service of reviewing will be able to earn and bid more credits, providing a positive incentive to participate in the review process.

**Remark 1.** *We note that there are two simple ways to incentivize authors to earn sufficient credits in Stage II to cover their bid from Stage I. The first is to refuse to accept papers without sufficient reviewing credit; the second is to model the conference submission process as a repeated game where an author can either save credits or be indebted credits from round to round.*

#### 4.1 Setup

**Notation** Suppose we have  $N$  papers submitted to the conference peer review mechanism denoted  $\omega_i$  (for  $i \in [N]$ ) with  $K$  participating agents. We use the index  $i$  when referring to papers in general and the index  $k$  in referring to agents in general. For ease of exposition, we treat a submitted paper as the product of a single agent (we relax this in Section 5). We represent the papers submitted by agent  $k$  as  $\{\omega_{i_s}^k\}$  and the papers reviewed by an agent  $k$  as  $\{\omega_{i_r}^k\}$  with the natural condition that  $\{\omega_{i_s}^k\} \cap \{\omega_{i_r}^k\} = \emptyset$ . Author  $k$ 's Stage I bid accompanying paper  $\omega_i$  is  $b_i$ . Reviewer  $k$ 's scores of paper  $\omega_i$  on the  $T$  criteria are written as  $(\mathbf{x}_k^{[T]})_i$ , their first predictions of peer  $k'$ 's reported scores as  $(\hat{\mathbf{p}}_{k'[T] \leftarrow k})_i$ , and their second predictions of peer  $k'$ 's reported scores as  $(\hat{\mathbf{p}}_{k'[T] \leftarrow k}^+)_i$ . When reviewer  $k$  exerts enough effort to receive informed private signals up to criterion  $t_k$ , the cost of the effort is  $e_k(t_k)$ . A complete table of notation is presented in Appendix A.

**Mechanism Design Background** We define an agent  $k$ 's *strategy* as a map from their private information to a distribution over reports elicited by a mechanism. An agent's *Stage I strategy* consists of their bids for all their submitted papers, denoted  $s_k^{(1)} = (b_i)_{i \in \{\omega_{i_s}^k\}}$ . An agent's *Stage II strategy* consists of the effort they exerted to arrive at an informed review score on  $t_k$  criteria, reported scores on all  $T$  criteria and their two predictions of a peer  $k'$ 's score on every criterion for every paper they review, denoted  $s_k^{(2)} = \left( \left\{ e_k(t_k), \hat{\mathbf{x}}_k^{[T]}, \hat{\mathbf{p}}_{k'[T] \leftarrow k}, \hat{\mathbf{p}}_{k'[T] \leftarrow k}^+ \right\} \right)_{i \in \{\omega_{i_r}^k\}}$ . The complete strategy for an agent participating in the mechanism is simply  $s_k = s_k^{(1)} \times s_k^{(2)}$ . A *strategy profile* is a tuple of every agent's strategy  $\mathbf{s} = (s_1, s_2, \dots, s_K)$ . We represent the strategy profile of all agents excluding agent  $k$  as  $\mathbf{s}_{-k}$ .

**Definition 1** (Dominant Strategy). *A strategy  $s_k$  is dominant if the strategy profile  $(s_k, \mathbf{s}_{-k})$  maximizes agent  $k$ 's payoff for any strategy profile  $\mathbf{s}_{-k}$ .*

**Definition 2** ((Strict) Bayesian Nash Equilibrium (BNE)). *Suppose player  $k$  observes private signal  $x_k$ . A strategy profile  $\mathbf{s} = (s_1(x_1), s_2(x_2), \dots, s_K(x_K))$  is a Bayesian Nash Equilibrium if no agent  $k$  can increase their expected payoff by deviating from strategy  $s_k(\cdot)$  when the other agents play  $\mathbf{s}_{-k}$ , for any private signal  $x_k$ . The BNE is strict if deviating from  $s_k$  decreases agent  $k$ 's expected payoff. If the strategy profile of the BNE consists of identical strategies for all agents, the BNE is symmetric.*

**Utility Function** We specify the utility function  $U_k^{(1)}$  of an agent  $k$  participating in the submission mechanism (Stage I) as follows:

$$U_k^{(1)} = \begin{cases} \sum_{i \in \{\omega_{i_s}^k\}} (v_k \delta_i - c_i) & \text{accepted for review} \\ 0 & \text{reject without review} \end{cases} \quad (1)$$

where  $\delta_i$  is a random variable indicating whether or not the paper is accepted *for publication*,  $v_k$  is the agent's value in credits for a paper accepted for publication and  $c_i$  is the price charged by the Stage I mechanism for a review slot. We can similarly specify the utility function  $U_k^{(2)}$  of an agent  $k$  participating in the H-DIPP reviewing mechanism (second stage) as follows:

$$U_k^{(2)} = \sum_{i \in \{\omega_{i_r}^k\}} r_k \left( \left[ s_k^{(2)} \right]_i, \left[ \mathbf{s}_{-k}^{(2)} \right]_i \right) - (e_k(t_k))_i \quad (2)$$

where  $r_k$  is the reward credits paid to reviewer  $k$  by the H-DIPP mechanism for a given strategy  $[s_k^{(2)}]_i$  on paper  $\omega_i$ . The total utility of an agent participating in the mechanism is  $U_k = U_k^{(1)} + U_k^{(2)}$ .

We will analyze the incentives and strategy profiles of the two stages separately. This would be appropriate if the value of a review slot and payments to reviewers had no bearing on each other. We discuss when this may fail and what can be done about it in Section 4.4. With this setup, we now present the two stages of the mechanism.

## 4.2 Stage I: Paper Submission

Our Stage I mechanism uses a VCG auction to allocate review slots and charge winners the threshold bid necessary for their paper to be reviewed. We assume that the number of review slots  $P$  is fixed and announced by the conference before paper submission deadline, based on how many reviews it expects to be able to pay for (perhaps with some predictive model using historical data, with some leeway to run a deficit). We defer a discussion of how the number of review slots  $P$  may be determined dynamically in Section 4.4, along with the associated considerations and challenges. The auction asks authors to submit a bid  $b_i$  along with every paper  $\omega_i$  indicating how much they are willing to pay for a review slot for that paper. Authors can observe the quality of their paper  $\omega_i$  and arrive at a belief that the paper will be accepted for publication with probability  $\eta_i$  conditional on review under the equilibrium of the submission mechanism<sup>1</sup> (*bids* have no bearing on the chance of acceptance for publication). Hence, denoting reviewer  $k$ 's value of a paper accepted for publication as  $v_k$ , the expected value of a review slot is  $\mathbb{E}[v_k \delta_i] = v_k \eta_i$ . We also have the following natural conditions: (1) the review slots are identical<sup>2</sup>, (2) each paper may only receive a single review slot, and (3) package bids are not allowed, i.e., authors cannot make a single bid for multiple review slots for multiple papers, but must instead make a separate bid for each slot for each paper. These conditions make it simple and practical to identify the winners of the auction when a single author can submit multiple papers for review.

After collecting bids, the VCG auction allocates the  $P$  review slots to the  $P$  papers with the highest bids, denoted  $\{\omega_1^*, \dots, \omega_P^*\}$ , with associated bids  $\{b_p^*\}_{p \in [P]}$ . Denoting the paper with the  $P+1$ -highest bid  $b_{P+1}^*$  as  $\omega_{P+1}^*$ , the mechanism charges all authors of the papers with winning bids a price of  $c_i = \sum_{\ell=1, \ell \neq i}^{P+1} b_\ell^* - \sum_{\ell=1, \ell \neq i}^P b_\ell^* = b_{P+1}^*$ . Thus, all papers accepted for review are charged the highest losing bid<sup>3</sup>, and accepting these papers maximizes author welfare for a given number of review slots  $P$ .

<sup>1</sup>In other words, authors should reason about their paper's acceptance probability in a pool with a particular combination of  $P$  of the 'the best papers' and papers from authors who can bid higher since they review more. This would be the equilibrium of the submission mechanism; if the pool had some other combination of papers, a stronger paper or credit-rich author would be willing to bid higher to displace a paper in the pool with their paper. We are assuming that by attending conferences and participating in review, authors have some notion of how competitive their paper will be in relative terms.

<sup>2</sup>Even though from the reviewing side, the review slots are heterogeneous (papers will differ in terms of cost of finding well-matched reviewers), we treat the slots as identical from the author's side in the interest of fairness. This could be relaxed by splitting the conference into 'tracks' within which this treatment is more justifiable.

<sup>3</sup>This is correct when viewing papers-as-bidders, but a potential concern when viewing authors-as-bidders is that authors who submit multiple papers may try to bid strategically. Strictly speaking, VCG would require charging author  $k$  who gets  $p_k$  slots the sum of the  $p_k$  highest losing bids excluding their own. In practice, this is unlikely to matter since the number of review slots vastly outnumbers an author's submissions.

Truthful bidding is a *weakly dominant strategy* for VCG, so we expect that rational actors will place bids equal to their value for their paper being reviewed. The natural assumption of identical review slots and prohibition of package bids precludes the need to solve a computationally challenging knapsack problem to identify winners. Thus, the *Stage I strategy profile* that maximizes aggregate author welfare under a given number of review slots  $P$  is  $\mathbf{s}^{(1)} = \left( (v_k \eta_i)_{i \in \{\omega_{i_s}^k\}} \right)_{k=1, \dots, K}$ . The mechanism allocates review slots to the highest bidding papers, which will generally be papers that are likely to be accepted for publication and/or have authors who generate a lot of value as reviewers. Thus, the papers accepted for review should also ideally be of higher quality if the review process screens well for quality. This should also eliminate the need for desk rejection.

### 4.3 Stage II: Peer Review

We now present the second stage of the mechanism which focuses on incentivizing honest and effortful reviews of the  $P$  papers accepted by the VCG auction. The context here is as follows: the bid requirement for the submission process in Stage I creates demand for credits, and since this can only be earned by reviewing, it incentivizes researchers to join the reviewer pool. While this is desirable, we must structure the payment for reviews carefully. A flat fee to reviewers could result in reviewers putting in minimal effort and getting paid anyway; if credits could be earned with little effort, bids could be made arbitrarily high, defeating the point of the mechanism. Instead, we need to structure the payment such that honest and effortful reviews are paid more than low-effort or false reviews. This is a challenging problem in the absence of ground truth verification; and our proposed H-DIPP mechanism is a novel solution to this problem. Note that the H-DIPP mechanism stage describes only the reviewing process for a single paper; we are agnostic about the paper-reviewer matching process, and meta-reviewers can make final accept/reject decisions as usual.

Thus, our two main considerations are incentivizing reviewers to exert effort, and to report their true reviews. We achieve both by adapting the Target Differential Peer Prediction (TDPP) (Schoenebeck and Yu, 2020) to operate within the Hierarchical Mutual Information Paradigm (HMIP) (Kong and Schoenebeck, 2018). The Differential Peer Prediction mechanism is a recently proposed prediction market-style mechanism applicable to our small-group, heterogeneous reviewer setting and determines payments for reviews of each paper independently. They present their mechanism for a group of 3 agents, ideal for our peer review setting. The mechanism works as follows: first, it assigns the 3 agents the roles of *expert*, *target*, and *source*. Then, it asks the *expert* for a probability distribution representing their prediction of the *target*'s signal. Finally, it reveals the *source*'s signal and asks the *expert* for an updated prediction of the *target*'s signal. Schoenebeck and Yu (2020) show that the total expected payment to agents is the mutual information of the source and target's signals conditioned on the expert's signals, and this yields a strong truthfulness guarantee (Kong and Schoenebeck, 2019).

However, truthfulness is not enough; reporting low-effort signals can still be truthful. Additionally, since these truthful signals are paid their mutual information, when low-effort signals correlate more than high-effort signals, it becomes natural to coordinate into less informative equilibria. HMIP proposes a framework to get around this issue by constructing a hierarchy of criteria that range from requiring low-effort to high-effort signals, and paying agents the *gain* in information on higher level criteria. Single-HMIM is an instantiation of the HMIP framework to incentivize effortful reviews on a single task, but uninformative equilibria can pay more than truthful equilibria and we seek to overcome this.

We combine the insights from TDPP and HMIP to propose a Hierarchical-Differential Peer Prediction (H-DIPP) mechanism. We present some fundamentals of the Mutual Information Paradigm (Kong and Schoenebeck, 2019) in Appendix C, and detail the relationship between our H-DIPP mechanism and the related single-HMIM and TDPP mechanisms in Appendix D. Proofs are presented in Appendix E. We now give a more detailed discussion of our reviewing mechanism H-DIPP (Mechanism 1).

#### 4.3.1 Model Preliminaries

We begin by assuming that a given paper  $\omega_i$  will be scored on a set of  $T$  *criteria* by exactly 3 reviewers<sup>4</sup>  $A, B, C$ . Each reviewer rates the paper  $\omega_i$  on criterion  $t$  on a  $D_t$ -point scale, i.e., reviewers' reported scores can take on any value in a  $\Sigma_t = \{1, 2, \dots, D_t\}$ . The same H-DIPP mechanism will be used determine payments for every paper's reviewers, so we drop the index  $i$  referring to a specific paper going for

---

<sup>4</sup>This can be extended to any number of reviewers by cyclically assigning the roles of expert, target, and source.

ward. Now, suppose we have a sample space  $\Omega$  of all the possible realizations of papers ('world states') a reviewer may see. A reviewer  $k$ 's *private review score* on criterion  $t$  is the random variable  $X_k^t : \Omega \rightarrow \Sigma_t$  describing their honest private assessment (or 'signal') of the paper on that criterion, *when they exert the necessary effort to arrive at an informed opinion*. A reviewer's joint distribution of effort-informed review scores for all criteria is represented by the multivariate random variable  $\mathbf{X}_k^{[T]} = (X_k^1, X_k^2, \dots, X_k^T) \in \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_T$ , and the full joint distribution of effort-informed review scores over all criteria for all reviewers is  $\mathbb{P}(\mathbf{X}_A^{[T]}, \mathbf{X}_B^{[T]}, \mathbf{X}_C^{[T]})$ . We assume that this joint distribution is common knowledge and has full support, i.e., every assignment of outcomes has non-zero probability. Unlike HMIP, differential peer prediction (and our approach) *does not* require the assumption that the common prior be symmetric, and is applicable for heterogeneous agents. We state this assumption formally below:

**Assumption 1** (Common Prior with Full Support). *Reviewers' effort-informed signals on all criteria  $\{\mathbf{x}_A^{[T]}, \mathbf{x}_B^{[T]}, \mathbf{x}_C^{[T]}\}$  are drawn from the common prior  $\mathbb{P}(\mathbf{X}_A^{[T]}, \mathbf{X}_B^{[T]}, \mathbf{X}_C^{[T]})$ , and the common prior is common knowledge. The common prior also has full support.*

**Criteria Hierarchy** Next, we construct the *hierarchy of criteria* that each reviewer will rate the paper on. We seek to construct a universal hierarchy for all reviewers by imposing an ordering such that reviewers require less effort to come to an informed private assessment on 'lower' level criteria than 'higher' level criteria. Naturally, when a reviewer exerts enough effort to come to an informed private assessment on some higher level criterion, they will also have arrived at an informed private assessment on all lower level criteria, so the total effort required to arrive at informed private assessments increases up the hierarchy. Thus, we say  $t_i > t_j$  if scoring the paper on criterion  $t_i$  requires strictly more effort than providing a score on criterion  $t_j$ . When a reviewer has exerted enough effort to arrive at an informed review score  $x_k^t$  on criterion  $t$ , we say that reviewer has *completed* criteria  $[t_k] = \{1, \dots, t_k\}$ . The idea is to explicitly construct lower level criteria to collect the cheap signals and pay agents more for the information *added* by valuable higher level criteria. We can define the positive, increasing function  $e_k : [T] \rightarrow \mathbb{R}^+$  specifying reviewer  $k$ 's *total effort cost* (in the mechanism-specific credits) to obtain effort-informed signals up to criterion  $t$ , where  $e_k(t) > e_k(t')$  if  $t > t'$ . While we aim to construct a single universal criteria hierarchy where the effort cost is monotonically increasing for all reviewers, the cost function  $e_k$  itself may be different for different reviewers; for instance, veteran peer reviewers may require less effort to score papers on every criterion. We give a sample plausible criteria hierarchy below.

In the current peer review setting at machine learning conferences, reviewers are already required to report scores on a variety of criteria, so we simply require a particular structuring of existing criteria. As an example, one possible hierarchy of criteria could be the following (where responses fall on a 5-point scale):

- **Criterion 1:** (presentation) "How clear is the writing/exposition?"
- **Criterion 2:** (completeness) "How easy would it be to reproduce major results?"
- **Criterion 3:** (correctness) "How accurate are the technical claims and methodology?"
- **Criterion 4:** (contribution) "How interesting and valuable of a contribution is the work? "
- **Criterion 5:** (overall quality) "In which quintile of submitted papers would this work fall?"

We see that for example, answering the level 4 question on contribution requires a level of effort that would also yield informed opinions on the presentation, completeness, and correctness; the reverse isn't obviously true. In general, the criteria hierarchy should be structured so higher level criteria require more effort.

**Strategic and Uninformed Reports** Our mechanism will elicit reviewers' private review scores  $X_k^t$  on all criteria, and we denote the (potentially strategic) *reported review score* as  $\hat{X}_k^t$ . With  $\theta_k$  as the transition matrix describing reviewer  $k$ 's reporting strategy, the joint distribution of true signals and reported signals under a signal strategy profile  $(\theta_A, \theta_B, \theta_C)$  is  $\mathbb{Q}(\hat{\mathbf{X}}_A^{[T]}, \hat{\mathbf{X}}_B^{[T]}, \hat{\mathbf{X}}_C^{[T]}, \mathbf{X}_A^{[T]}, \mathbf{X}_B^{[T]}, \mathbf{X}_C^{[T]}) = \theta_A(\hat{\mathbf{X}}_A^{[T]} | \mathbf{X}_A^{[T]})\theta_B(\hat{\mathbf{X}}_B^{[T]} | \mathbf{X}_B^{[T]})\theta_C(\hat{\mathbf{X}}_C^{[T]} | \mathbf{X}_C^{[T]})\mathbb{P}(\mathbf{X}_A^{[T]}, \mathbf{X}_B^{[T]}, \mathbf{X}_C^{[T]})$ . When a reviewer plays a truthful reporting strategy (i.e.,  $\theta_k$  is identity), we denote it as  $\tau_k$ .

Now, if a reviewer only completes up to  $t$  criteria, they will not arrive at an informed private score on criteria at higher levels. However, their signals on lower level criteria do provide some information,

allowing them to make reasonable guesses of the scores on higher level criteria<sup>5</sup>. Indeed, we assume that on incomplete higher level criteria, a reviewer's default private signal is their 'best guess' of the signal that other agents would expect of them given the information they already have, and this may depend on how much effort the other agents exert. This is given by some function  $g_k^{t'|t} : \Sigma_1 \times \dots \times \Sigma_t \rightarrow \Sigma_{t'}$  where  $x_k^{t'} = g_k^{t'|t}(\mathbf{x}_k^{[t]})$  represents reviewer  $k$ 's best guess of their private assessment on criterion  $t'$  when they have completed only criteria at level  $t$  (where  $t' > t$ ). Note that  $x_k^t$  thus either refers to an effort-informed signal drawn from the common prior or a best-guess, depending on the reviewer's effort level. We use the notation  $\mathbf{x}_k^{[t_k]}$  if we wish to emphasize that we are only referring to the effort-informed signals on criteria  $[t_k]$  (when  $t_k < T$ ).

**Other Assumptions** We now list some other assumptions we make on the information structure of the reviewers. First, we adopt a conditional independence assumption from Kong and Schoenebeck (2018), which states that a reviewer  $k$ 's private score on some criterion  $t$  contains all the information needed to predict their peer  $k$ 's score on that criterion, given that peer's scores on lower level criteria. We write  $I(\cdot)$  to represent the mutual information (see Appendix C for more details).

**Assumption 2** (Conditional Independence). *We make the following conditional independence assumption on the common prior:*

$$I(\mathbf{X}_{k'}^{[T]}, X_k^t | \mathbf{X}_k^{[t-1]}) = I(X_{k'}^t, X_k^t | \mathbf{X}_k^{[t-1]}) \quad (3)$$

or equivalently,

$$I(\mathbf{X}_{k'}^{[T]} \setminus X_{k'}^t, X_k^t | \mathbf{X}_k^{[t-1]}, X_{k'}^t) = 0 \quad (4)$$

Finally, we make an assumption on the informativeness of the signals a reviewer might observe. When a reviewer  $k$  arrives at review scores  $\mathbf{x}_k^{[t]}$  on  $[t]$  criteria, they also arrive at a posterior belief  $\mathbb{P}(\mathbf{X}_{-k}^{[T]} | \mathbf{x}_k^{[t]})$  about the review scores of other reviewers. Informally, we assume there is some signal a reviewer could observe, which when paired with two different signals from a second reviewer, would induce two different beliefs about the signals of the third reviewer.

**Assumption 3** (Second Order Stochastic Relevance). *Suppose that for any distinct signals  $x_B^t, \tilde{x}_B^t \in \Sigma_t$  on any criterion  $t$ , there is  $x_A^t \in \Sigma_t$  such that:*

$$\mathbb{P}(X_C^t | x_A^t, x_B^t) \neq \mathbb{P}(X_C^t | x_A^t, \tilde{x}_B^t)$$

We assume the above holds for any permutation of the reviewers  $\{A, B, C\}$ , so the common prior  $\mathbb{P}$  is second order stochastic relevant.

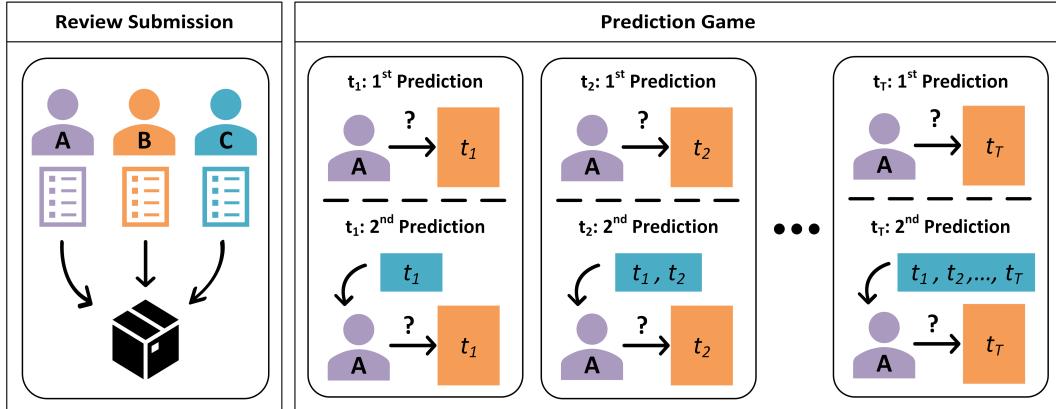


Figure 2: An illustration of the H-DIPP mechanism from reviewer  $A$ 's perspective, when reviewer  $B$  is their assigned target and reviewer  $C$  is their assigned source.

<sup>5</sup>This is an important difference between our approach and the original HMIP, where reporting on incomplete criteria is optional. See Appendix D for more details.

### 4.3.2 Mechanism and Payments

Having detailed our assumptions, we now describe exactly how the mechanism collects reports and computes the payment. The mechanism is summarized as Mechanism 1 and illustrated in Figure 2.

**Reports** First, each reviewer reports their review scores to the mechanism per the usual process. Reviewers arrive at private review scores  $\mathbf{x}_k^{[t_k]}$  on  $t_k$  criteria by exerting  $e_k(t_k)$  effort, and report this to the mechanism (potentially strategically) as  $\hat{\mathbf{x}}_k^{[T]} \sim \theta_k(\cdot | \mathbf{x}_k^{[t_k]})$ .

Once all the reviews are in, every reviewer (independently) plays a simple prediction game. We refer to the reviewer playing the prediction game as an ‘expert’, and every ‘expert’ will be assigned a ‘source’ and a ‘target’. We use the following assignments of (expert, target, source) roles:  $(A, B, C)$ ,  $(B, C, A)$ , and  $(C, A, B)$ , i.e., when  $A$  plays the game as ‘expert’,  $B$  is the assigned ‘target’ and  $C$  is the assigned ‘source’, and so on. During the prediction game, the *expert* must make two probabilistic predictions of their *target*’s reported review score on each criterion; the expert’s first prediction on a criterion is their best guess of their target’s review score on that criterion, and the expert’s second prediction is their updated best guess after the mechanism reveals their *source*’s review score on that criterion. The criteria are revealed in a pre-determined order specified by the criteria hierarchy. As an example, consider the case where reviewer  $A$  plays the prediction game after all the reviews are in (with target  $B$  and source  $C$ ). The mechanism will first elicit the expert  $A$ ’s prediction  $\hat{\mathbf{p}}_{B^1 \leftarrow A} = P(\hat{X}_B^1 | \mathbf{x}_A^{[t_A]})$  of the target  $B$ ’s score on criterion 1, and then reveals the source  $C$ ’s signal  $\hat{x}_C^1$  on that criterion and asks for an updated prediction  $\hat{\mathbf{p}}_{B^1 \leftarrow A}^+ = P(\hat{X}_B^1 | \mathbf{x}_A^{[t_A]}, \hat{x}_C^1)$ . Next on criterion 2, the mechanism asks for the expert  $A$ ’s prediction  $\hat{\mathbf{p}}_{B^2 \leftarrow A} = P(\hat{X}_B^2 | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^1)$  of the target  $B$ ’s score on criterion 2, and then reveals the source  $C$ ’s signal  $\hat{x}_C^2$  on that criterion and asks for an updated prediction  $\hat{\mathbf{p}}_{B^2 \leftarrow A}^+ = P(\hat{X}_B^2 | \mathbf{x}_A^{[t_A]}, \hat{x}_C^2)$ . This proceeds in order until the final  $T$ th criterion. The mechanism thus requests  $3T$  reports: reviewers’  $T$  review scores, their first predictions of their assigned target’s scores on  $T$  criteria, and updated second predictions of that target peer’s scores on  $T$  criteria given the source peer’s scores. A reviewer’s strategy on the reviewing task is  $s_k = (e_k(t_k), \hat{\mathbf{x}}_k^{[T]}, \hat{\mathbf{p}}_{k^t \leftarrow k}, \hat{\mathbf{p}}_{k^t \leftarrow k}^+)$ .

---

#### Mechanism 1: H-DIPP

---

**Input:**  $\omega_i$  (paper under review),  $\{t_1, t_2, \dots, t_T\}$  (hierarchical criteria),  $\{A, B, C\}$  (set of 3 reviewers),  $\{\alpha_A^{[T]}, \beta_A^{[T]}, \alpha_B^{[T]}, \beta_B^{[T]}, \alpha_C^{[T]}, \beta_C^{[T]}\}$  (mechanism hyperparameters criteria)

**Output:**  $\{r_A, r_B, r_C\}$  (payments for the reviewers)

- 1 Elicit reviewers’ scores on each criterion  $\{\hat{\mathbf{x}}_k^{[T]}\}_{k \in \{A, B, C\}}$  per the usual peer review process
  - 2 Assign (expert, target, source) roles as follows:  $(A, B, C)$ ,  $(B, C, A)$ ,  $(C, A, B)$ .
  - 3 **for each criterion  $t$  ascending the criteria hierarchy do**
    - 4 Elicit  $\{\hat{\mathbf{p}}_{B^t \leftarrow A}, \hat{\mathbf{p}}_{C^t \leftarrow B}, \hat{\mathbf{p}}_{A^t \leftarrow C}\}$ , i.e., reviewers  $A, B, C$  predictions of the reports of reviewers  $B, C, A$  respectively on criterion  $t$ .
    - 5 Reveal source scores on criterion  $t$ , i.e., reveal reviewer  $A$ ’s scores  $\hat{x}_A^t$  to reviewer  $B$ , reviewer  $B$ ’s score  $\hat{x}_B^t$  to reviewer  $C$ , and reviewer  $C$ ’s score  $\hat{x}_C^t$  to reviewer  $A$ .
    - 6 Elicit  $\{\hat{\mathbf{p}}_{B^t \leftarrow A}^+, \hat{\mathbf{p}}_{C^t \leftarrow B}^+, \hat{\mathbf{p}}_{A^t \leftarrow C}^+\}$ , i.e., reviewers  $A, B, C$  *updated* predictions of the reports of reviewers  $B, C, A$  respectively on criterion  $t$ .
    - 7  $r_{A^t}^{\text{expert}} \leftarrow S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B^t \leftarrow A}) + S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{A^t \leftarrow C}^+); r_{A^t}^{\text{target}} \leftarrow S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A^t \leftarrow C}^+) - S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A^t \leftarrow C})$
    - 8  $r_{B^t}^{\text{expert}} \leftarrow S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C^t \leftarrow B}) + S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{B^t \leftarrow A}^+); r_{B^t}^{\text{target}} \leftarrow S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B^t \leftarrow A}^+) - S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B^t \leftarrow A})$
    - 9  $r_{C^t}^{\text{expert}} \leftarrow S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A^t \leftarrow C}) + S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{C^t \leftarrow B}^+); r_{C^t}^{\text{target}} \leftarrow S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C^t \leftarrow B}^+) - S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C^t \leftarrow B})$
  - 10  $r_A \leftarrow \sum_{t=1}^T \alpha_{A^t}^t r_{A^t}^{\text{expert}} + \beta_{A^t}^t r_{A^t}^{\text{target}}$
  - 11  $r_B \leftarrow \sum_{t=1}^T \alpha_{B^t}^t r_{B^t}^{\text{expert}} + \beta_{B^t}^t r_{B^t}^{\text{target}}$
  - 12  $r_C \leftarrow \sum_{t=1}^T \alpha_{C^t}^t r_{C^t}^{\text{expert}} + \beta_{C^t}^t r_{C^t}^{\text{target}}$
- 

**Payments** The mechanism will pay each reviewer for their roles as expert and target (source payment is 0). Reviewers receive ‘expert payment’ based on their performance in the prediction game, and

receive ‘target payment’ based on a peer’s performance in the prediction game. The goal is for the expert payment to incentivize truthful predictions, and the ‘target payment’ to incentivize honest reports of private review scores. With reviewer-specific mechanism hyperparameters  $\{\alpha_A^{[T]}, \beta_A^{[T]}\}$ , the criterion- $t$  specific payment to reviewer  $A$  using the log scoring rule  $S_{\log}$  is:

$$r_{At} = \alpha_A^t \left( \underbrace{S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B^{t \leftarrow A}}) + S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B^{t \leftarrow A}}^+)}_{\text{expert payment}} \right) + \beta_A^t \left( \underbrace{S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A^{t \leftarrow C}}^+) - S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A^{t \leftarrow C}})}_{\text{target payment}} \right) \quad (5)$$

A reviewer’s *expert payments* depend on how good their predictions were during the prediction game (incentivizing truthful predictions). A reviewer’s *target payments* depend how much better a peer expert’s second prediction of their reported review score was over their first; the idea is that their peer’s second prediction can update reliably towards the target reviewer’s true review score (thus incentivizing truthful reviews). To compute reviewer  $A$ ’s total payment, we simply sum the individual criterion-specific payments  $r_A = \sum_t r_{At}$ . Payments for the other reviewers  $B, C$  are determined analogously.

The mechanism hyper-parameters are used to appropriately weight the criterion-specific payments to cover the marginal effort costs of completing higher level criteria, and will describe how to set them in relation to reviewers’ common prior and effort costs in the next subsection. The task of gleaning information about the common prior and reviewers’ effort costs is important, but we leave this to future work.

#### 4.3.3 Equilibria

We now analyze the equilibria of our mechanism<sup>6</sup>. First, we give the strict truthfulness of our mechanism when all reviewers complete a certain number of criteria.

**Theorem 1** (Strict Truthfulness). *Suppose reviewers  $A, B, C$  complete  $[t_A], [t_B]$ , and  $[t_C]$  criteria respectively. Then, on shared complete criteria  $[t'] = [\min(t_A, t_B, t_C)]$ , the H-DIPP mechanism is strictly truthful, i.e., all reviewers honestly reporting their true predictions and private review scores on shared complete criteria  $[t']$  is a strict Bayesian Nash equilibrium.*

However, mere truthfulness is not enough; the mechanism must also incentivize reviewers to exert the necessary effort to complete all criteria. After all, we care about eliciting reviewers’ *effort-informed signals* on all criteria. To achieve this, we must reason about reviewers’ marginal payments and costs of effort. Going forward, we use  $R$  to denote expected payment. We break up the total marginal expected payment  $\Delta R_k$  into the marginal expected expert and target payments  $\Delta R_k^{\text{expert}}$  and  $\Delta R_k^{\text{target}}$ , which in turn are each broken up into marginal payments  $\Delta R_{k^t}^{\text{expert}}$  and  $\Delta R_{k^t}^{\text{target}}$  on each criterion. We give the following lemma on reviewers’ marginal expected payments for completing an additional criterion and show that it is always non-negative (when peers have completed all criteria and are reporting truthfully).

**Lemma 1** (Non-negative Marginal Payments). *Suppose reviewers  $B, C$  complete all criteria and truthfully report their private review scores and predictions. Then, reviewer  $A$ ’s marginal expected payment for completing criterion  $t_A$  and honestly reporting private review scores and predictions is  $\Delta R_A(t_A) := R_A(t_A) - R_A(t_A - 1) = \Delta R_A^{\text{expert}}(t_A) + \Delta R_A^{\text{target}}(t_A)$ , where:*

1.  $\Delta R_A^{\text{expert}}(t_A) := R_A^{\text{expert}}(t_A) - R_A^{\text{expert}}(t_A - 1) = \sum_{t=1}^T \alpha_A^t \Delta R_{At}^{\text{expert}}(t_A)$
2.  $\Delta R_A^{\text{target}}(t_A) := R_A^{\text{target}}(t_A) - R_A^{\text{target}}(t_A - 1) = \sum_{t=1}^T \beta_A^t \Delta R_{At}^{\text{target}}(t_A)$
3.  $\Delta R_{At}^{\text{expert}}(t_A) := I(X_B^t; X_A^{t_A} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[t]}) + I(X_B^t; X_A^{t_A} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[t-1]})$
4.  $\Delta R_{At}^{\text{target}}(t_A) := \begin{cases} 0, & t < t_A \\ I(X_B^{t_A}; X_A^{t_A} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) - I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{X}_A^{[t_A-1]}) | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}), & t = t_A \\ \mathbb{E} \left[ KL \left( \mathbb{P}(X_B^t | g_A^{t|t_A}(\mathbf{X}_A^{[t_A]}), \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t-1]}) || \mathbb{P}(X_B^t | g_A^{t|t_A-1}(\mathbf{X}_A^{[t_A-1]}), \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t-1]}) \right) \right], & t > t_A \end{cases}$

Furthermore, the marginal expected expert and target payments on each criterion are always non-negative, i.e.,  $\Delta R_{At}^{\text{expert}}(t_A), \Delta R_{At}^{\text{target}}(t_A) \geq 0$  for all  $t, t_A \in [T]$  and the inequality is strict when  $t = t_A$ . Similar results hold for reviewers  $B$  and  $C$  when peers complete all criteria and report truthfully.

<sup>6</sup>We ignore permutation equilibria, where the labels for signals are permuted. Such equilibria exist in our mechanism, but they are not of serious concern; see Kong and Schoenebeck (2019) for details.

With this, we can now specify how to set the mechanism hyperparameters to incentivize a reviewer to complete all criteria, such that all reviewers completing all criteria and reporting truthfully is a strict BNE (equivalent to the *potent* property of hyperparameters in Kong and Schoenebeck (2018)).

**Theorem 2** (Optimal Mechanism Hyperparameters). *Suppose mechanism hyperparameters  $\{\alpha_A^{[T]}, \beta_A^{[T]}, \alpha_B^{[T]}, \beta_B^{[T]}, \alpha_C^{[T]}, \beta_C^{[T]}\}$  are solutions to the following linear program:*

$$\begin{aligned}
& \min \sum_{k \in \{A, B, C\}} \sum_{t=1}^T \alpha_k^t R_{k^t}^{expert}(T) + \beta_k^t R_{k^t}^{target}(T) \\
& \text{s.t. } \sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{expert}(t_A) + \beta_A^t \Delta R_{A^t}^{target}(t_A) > \Delta e_A(t_A) \text{ for } 1 \leq t_A \leq T \\
& \quad \sum_{t=1}^T \alpha_B^t \Delta R_{B^t}^{expert}(t_B) + \beta_B^t \Delta R_{B^t}^{target}(t_B) > \Delta e_B(t_B) \text{ for } 1 \leq t_B \leq T \\
& \quad \sum_{t=1}^T \alpha_C^t \Delta R_{C^t}^{expert}(t_C) + \beta_C^t \Delta R_{C^t}^{target}(t_C) > \Delta e_C(t_C) \text{ for } 1 \leq t_C \leq T \\
& \quad \alpha_k^t, \beta_k^t \geq 0
\end{aligned} \tag{6}$$

Then, for any reviewer  $k$ , if their peers complete all criteria and report predictions and review scores truthfully, reviewer  $k$  strictly maximizes their expected utility by also completing all criteria and reporting predictions and review scores truthfully. In other words, completing all criteria and reporting truthfully is a strict Bayesian-Nash equilibrium. We refer to this equilibrium as the *fully-informative equilibrium*.

Having shown that our desired strategy profile (all reviewers completing all criteria and reporting truthfully) is a strict BNE, we now discuss how this compares to other equilibria and strategy profiles. A common challenge is that truthful equilibria pay less than effortless uninformative equilibria (like in single-HMIM), so the truthful equilibria are not realized in practice. However, this is not the case with H-DIPP; we will show that an uninformative equilibrium pays nothing in expectation and the fully informative equilibrium pays strictly more. We begin with the following lemma on the total payment to reviewers in any equilibrium:

**Lemma 2** (Aggregate Reviewer Payments). *Let  $s_k = (e_k(t_k), \hat{\mathbf{x}}_k^{[T]}, \hat{\mathbf{p}}_{k' \leftarrow k}^+, \hat{\mathbf{p}}_{k' \leftarrow k}^-)$  be reviewer  $k$ 's strategy, where  $k'$  is reviewer  $k$ 's target. If  $\mathbf{s} = (s_A, s_B, s_C)$  is an equilibrium strategy profile, reviewer  $A$ 's expected payment is:*

$$\mathbb{E}[r_A] = \sum_{t=1}^T \alpha_A^t \left( -\mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]}\right)\right)\right] - \mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]}\right)\right)\right] \right) + \beta_A^t I\left(\hat{X}_B^t; \hat{X}_A^t | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]}\right) \tag{7}$$

(and expected payments for peers  $B, C$  computed analogously). Reviewers' total expected payment is:

$$\begin{aligned}
\mathbb{E}[r_A + r_B + r_C] &= \sum_{t=1}^T \alpha_A^t \left( -\mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]}\right)\right)\right] - \mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]}\right)\right)\right] \right) + \beta_A^t I\left(\hat{X}_B^t; \hat{X}_A^t | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]}\right) \\
&\quad + \alpha_B^t \left( -\mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t]}\right)\right)\right] - \mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t-1]}\right)\right)\right] \right) + \beta_B^t I\left(\hat{X}_C^t; \hat{X}_B^t | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{X}}_C^{[t-1]}\right) \\
&\quad + \alpha_C^t \left( -\mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t]}\right)\right)\right] - \mathbb{E}\left[H\left(\mathbb{Q}\left(\hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t-1]}\right)\right)\right] \right) + \beta_C^t I\left(\hat{X}_A^t; \hat{X}_C^t | \mathbf{X}_B^{[t_B]}, \hat{\mathbf{X}}_A^{[t-1]}\right)
\end{aligned} \tag{8}$$

We make several observations. First, unlike the ideal-HMIP proposal (and like the concrete single-HMIM proposal) in Kong and Schoenebeck (2018), we cannot claim that truthful reporting of private signals is a dominant strategy, since a reviewer's target payment for reporting private scores will depend on their peer's expert prediction strategy. However, it is the case that when we compare equilibria fixing (truthful or untruthful) peer strategies, a reviewer earns higher target rewards in the equilibria where they report truthfully, compared to ones where they strategically manipulate the score, since  $I\left(\hat{X}_B^t; \hat{X}_A^t | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]}\right) \leq I\left(\hat{X}_B^t; X_A^t | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]}\right)$  for all  $t \in [T]$  (and similarly for other reviewers) due to the data processing inequality (Kong and Schoenebeck, 2019).

Second, unlike TDPP (which does not account for effort), we cannot guarantee the strong truthfulness of the fully informative truth-telling equilibrium, i.e., we cannot guarantee that the fully informative equilibrium has higher aggregate utility than any other equilibrium. It may be possible for a reviewer to strategically manipulate their private signals so that when used as a source for some expert's prediction, the target receives higher payment; e.g., it may be possible that reviewer  $B$ 's strategic reports provides a higher target payment for reviewer  $A$  with  $I(X_B^t; \hat{X}_A^t | \mathbf{X}_C^{[t_C]}, \mathbf{X}_B^{[t-1]}) \leq I(\hat{X}_B^t; \hat{X}_A^t | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]})$ . Reviewer  $B$ 's strategic reports could then increase reviewer  $A$ 's target payment by more than it hurts their own target payment, thereby raising aggregate utility (costs are the same). However, reviewers do not receive payments for their role as source, and further must accept lower target payment for such strategic manipulation, so reviewers are disincentivized from pursuing such equilibria unilaterally. Nevertheless, we cannot rule out the possibility of viable collusive equilibria.

Thirdly, a reviewer's expected expert payment is always non-positive (since the Shannon entropy is non-negative), and a reviewer's expected target payment is always non-negative. We also must ensure that a reviewer's expected payment is greater than their cost in the fully informative equilibrium, so that it is individually rational for reviewers to participate in the mechanism. We can achieve this by scaling the hyperparameters such that the target payment pays more than the expert payment subtracts. This also grants that the fully-informative equilibrium pays more than an uninformative equilibrium where no reviewers exert effort. We state this formally below.

**Theorem 3** (Uninformative Equilibria and Individual Rationality).

- (a) *In the highest paying uninformative equilibrium (where all reviewers exert zero effort), every reviewer's expected payment (and hence aggregate reviewer welfare) is zero.*
- (b) *Suppose that in addition to the constraints specified in Theorem 2, the mechanism hyperparameters also satisfy  $R_k(T) - e_k(T) > 0 \Leftrightarrow \sum_{t=1}^T \alpha_k^t R_{k^t}^{\text{expert}}(T) + \beta_k^t R_{k^t}^{\text{target}}(T) > e_k(T)$  for  $k \in \{A, B, C\}$ , i.e., reviewers' expected utility in the fully-informative equilibrium is positive. Then, it is individually rational to participate in the mechanism and complete all criteria and report truthfully, given peers are doing the same. Additionally, the fully-informative equilibrium has higher individual and aggregate utility than the uninformed equilibrium.*

Finally, we note that while the mechanism can hide reviewers' identities from each other to provide some measure of protection against collusion, a formal analysis of H-DIPP's robustness to collusion would be valuable direction for future work.

#### 4.4 Demand and Supply

In Section 4.2, we assumed the VCG auction had a fixed number of review slots  $P$  to be auctioned. Here, we present the considerations and challenges in designing a mechanism that sets  $P$  dynamically *after* paper submissions based on the supply of reviewers, but a comprehensive solution is a task for future work. The foremost considerations in setting  $P$  based on the availability of reviewers are:

1. **Individual Rationality and Truthfulness:** Authors and reviewers should not expect to lose credits by participating the mechanism, and authors should not be able to game the number of review slots by bidding something other than their expected value for having their paper reviewed.
2. **Balanced Budget:** We aim to set  $P$  based on the number of reviews that can be paid for under the H-DIPP payment scheme using revenue from the auction. Recall that the cost to the mechanism to have one paper reviewed by a *given reviewer triplet* in the fully-informative equilibrium (if this is indeed the observed equilibrium) is given by Lemma 2, where hyperparameters have been set using Theorem 2.
3. **Maximize  $P$ :** We would like to accept as many papers for review as feasible.
4. **Fair Expertise-based Matching:** Reviewers must be well-qualified to review their assigned papers and have the appropriate domain expertise. Additionally, all papers accepted for review should have equal consideration when being assigned reviewers; the bid accompanying a paper should have no bearing on the quality of reviewer matches<sup>7</sup>. Thus, although the review

<sup>7</sup> Aside from the motivation of fairness, if higher bidding papers received priority in matching reviewers, reviewers could infer a paper's bid from the quality of their match, giving them a low-effort signal of paper quality.

slots are identical from authors' perspective (for fairness), the mechanism must match them to reviewers appropriately. We assume access to a black-box algorithm that can score a reviewer's match with a paper, and assign reviewers to papers based on this score (e.g. Toronto Paper Matching System (Charlin and Zemel, 2013)).

The final consideration comes from the observation that the mechanism should not blindly seek to minimize its total cost over all papers in an attempt to squeeze in as many review acceptances  $P$  as possible. The cost to the mechanism for reviewing a single paper depends both on the hyperparameters (optimally set via Theorem 2) as well as the mutual information in reviewers' signals, so total costs can be made abnormally low by badly matching reviewers and paying them little since their signals will have naturally low mutual information. Heuristically, we expect that well-matched expert reviewers hit the sweet spot of having low effort costs to review a paper and more mutually informative review scores with expert peers, so the mechanism pays a fair price for such expert reviews.

Nevertheless, we face a trade-off where accepting more papers for review can decrease the average quality of the paper-reviewer matches. Formally, let  $S_{i_k}^*$  denote the match score for the  $k$ -th reviewer of paper  $\omega_i$  and  $R_{i_k}^*$  denote the expected payment to the  $k$ -th reviewer of paper  $\omega_i$  under the optimal assignment of reviewers (as determined by the black-box algorithm) for some given number of papers  $P$  accepted for review. Then, a naive mechanism designer's objective is to pick the number of review slots to be the largest  $P$  with an acceptable average reviewer match score that also satisfies the balanced budget constraint (where slots are priced as in the VCG auction):

$$\begin{aligned} & \underset{P}{\operatorname{argmax}} \left( \frac{1}{P} \sum_{i=1}^P S_{i_A}^* + S_{i_B}^* + S_{i_C}^* \right) + \lambda P \\ & \text{s.t. } b_{P+1}^* \cdot P \geq \sum_{i=1}^P R_{i_A}^* + R_{i_B}^* + R_{i_C}^* \end{aligned} \tag{9}$$

The parameter  $\lambda$  controls the tradeoff between match quality and number of papers accepted for review, and is a design decision to be made by individual conferences. The objective can be modified if desired; substituting the mean reviewer match score for the median or an additional term penalizing variance in reviewer match scores are alternative choices. In a naive algorithm, we can find the optimal number of review slots by first computing all the valid number of review slots under the balanced budget constraint, and picking the largest one with an acceptable average reviewer match score (summarized in Algorithm 2).

---

**Algorithm 2:** Naive Algorithm to Determine Number of Review Slots  $P$

---

**Input:**  $N$  papers submitted to the conference, a black-box algorithm

matching reviewers to papers based on expertise (e.g. Toronto Paper Matching System)

**Output:** Optimal Number of Review Slots  $P$

- 1 **for**  $n = 1 : N$  **do**
  - 2     Obtain optimal matching of reviewers to  $n$  papers from black-box algorithm
  - 3     If accepting  $n$  papers satisfies the budget constraint  $b_{P+1}^* \cdot P \geq \sum_{i=1}^P R_{i_A}^* + R_{i_B}^* + R_{i_C}^*$ ,  
        add  $n$  to the list of allowed number of review slots  $\mathcal{P} \leftarrow \mathcal{P} \cup \{n\}$
  - 4 Select the largest  $P \in \mathcal{P}$  with an acceptable average reviewer
- match score, i.e., largest  $P$  that maximizes the objective  $\left( \frac{1}{P} \sum_{i=1}^P S_{i_A}^* + S_{i_B}^* + S_{i_C}^* \right) + \lambda P$
- 

However, such a mechanism faces several challenges. Firstly, it is not obvious that it is truthful; authors may be able to affect the number of review slots by changing their bids. Secondly, we need to identify the mechanism hyperparameters by solving the linear program in Theorem 2 and computing the expected payment in the fully informative equilibrium (Lemma 2); this requires estimating or eliciting information about the common prior  $\mathbb{P}$  and the reviewer triplet's effort cost functions  $e_k(\cdot)$ . Furthermore, in practice we may need a single set of mechanism hyperparameters  $\{\alpha^{[T]}, \beta^{[T]}\}$  that determine payments to all reviewers. Lastly, when the number of review slots  $P$  is fixed, authors' bids will not affect their utilities in the reviewing stage of the same conference since the hyperparameters will have been pre-determined; this allows us to analyze the two stages of the mechanism independently. However, endogenizing  $P$  would mean the setting of hyperparameters (and hence payments to

reviewers) is sensitive to the revenue raised (and hence bids made by authors). This coupling of the two stages complicates the analysis. These challenges are important directions for future work.

## 5 Discussion

Here, we discuss some important considerations for a practical implementation of our mechanism.

**Collaboration** We assumed for ease of exposition that every paper has a single author and this author must participate in the review process to earn enough credits to cover their bids. In practice, papers are largely multi-author efforts, and not every co-author is equally qualified to participate in the review process. The natural relaxation of the single-author assumption is to formalize the role of a ‘sponsor’ who puts up the credits needed for the bid, and will typically be a senior author who has banked enough credits by reviewing. However, non-author researchers may also supply the credits for a paper’s accompanying bid and be acknowledged in their capacity as sponsor, mitigating incentives to add credit-rich non-author sponsors as authors in exchange for a high bid. Additionally, we may wish to grant a lump-sum of credits to first-time authors to give them the right to submit some number of papers while they build the necessary expertise for review. We emphasize that bids do not have any bearing on a paper’s acceptance for publication, so sponsors will still be sensitive about their bids as a weak paper accepted for review will likely be rejected and the spent credits will be wasted.

**Common Prior Assumption** We have assumed that reviewers’ review scores are drawn from a common prior that is common knowledge, and we allow for asymmetric common priors so reviewers may be heterogeneous. This assumption is important to obtain the equilibrium results, but the extent to which it holds in practice, particularly with anonymous reviewers, needs to be answered empirically. This assumption can be made more realistic by making examples or summary statistics of a reviewer’s past reviews available for reference (so any reviewer can gauge how their opinion correlates with their peers’).

**Post-processing Review Scores** The H-DIPP mechanism elicits honest and effortful opinions from *heterogeneous* reviewers, with the goal of reducing the noise in estimates of a paper’s quality. However, this also means that reviewers’ honest and effortful opinions may be miscalibrated with respect to the rating scale (Shah et al., 2018). Thus, some post-processing of these scores (Wang and Shah, 2018; Cortes and Lawrence, 2021) to correct reviewer’s individual biases may produce a better picture of a paper’s quality.

**Conflicts of Interest** We note that our results on incentivizing effortful and truthful reviews hold under the assumption that there are no other incentives in play. However, if the matched reviewers have a conflict of interest, this would weaken the incentive to report honest and effortful reviews. Thus, related work on optimizing paper-reviewer matches (Xu et al., 2018; Jecmen et al., 2020) to avoid conflicts of interest is complementary and important to the success of our mechanism.

**Behavioural and Social Considerations** Our focus in this paper has been to develop a mechanism for paper submission and reviewing with desirable theoretical properties. However, our fundamental assumptions about human behaviour may or may not hold in practice. A primary concern is that introducing payments for reviewers could crowd-out pro-social motivations (Squazzoni et al., 2013; Gneezy and Rustichini, 2000). Empirical evidence for when such prediction market-style mechanisms succeed (Debmalya et al., 2020) and fail (Gao et al., 2014) can help refine our mechanism and similar ideas.

**A Secondary Market** We have assumed that our mechanism operates with its own currency, but we could further allow the sale and purchase of credits on a secondary market. This has the potential to allow individuals to earn real-world currency as reviewers, and could even encourage specialization in the review process: researchers who excel at reviewing in H-DIPP mechanism could allocate more of their time providing high quality reviews and be appropriately compensated for their work.

## 6 Conclusion and Future Work

We tackle incentive issues in scientific peer review by proposing a two-stage mechanism that allocates review slots to papers via a VCG auction and uses the revenue raised to pay participating reviewers in our

novel H-DIPP mechanism. The H-DIPP mechanism aims to incentivize reviewers to obtain and report a clearer signal of a paper’s quality in order to reduce the noise in the accept/reject decisions. We showed that the H-DIPP mechanism is strictly truthful, and described how to set mechanism hyperparameters so that exerting enough effort to score the paper honestly on all criteria is a strict Bayesian Nash equilibrium. We also showed that such a fully informative equilibrium pays more in expectation than any uninformative equilibrium, and described how the mechanism hyperparameters can be further tweaked so participating in the H-DIPP reviewing mechanism is individually rational. Finally, we discussed the trade-off between the number of available review slots and quality of paper-reviewer matches. However, there remain important questions to be answered, and we present some of them below:

1. **Criteria Hierarchy:** What is the best construction of the criteria hierarchy?
2. **Estimating Hyperparameters and Costs:** Setting mechanism hyperparameters and the number of review slots  $P$  requires information about reviewers’ marginal payments and costs, so identifying ways to learn (Liu and Chen, 2016) or elicit this information would be valuable.
3. **Endogenizing the Number of Review Slots  $P$ :** How can we allow  $P$  to be set based on the availability of reviewers, while guaranteeing an individually rational, truthful, and balanced budget mechanism that accepts as many well-matched papers as possible?
4. **Variance in Payments:** Is the variance in payments so large that honest and effortful reviews get negative payments with unacceptable regularity? If so, can the addition of a few more reviewers reduce this variance?
5. **Improving Equilibria:** Can we leverage information across reviewing tasks (Dasgupta and Ghosh, 2013; Agarwal et al., 2017) in an efficient way to further increase total payoffs in the truth-telling equilibrium? How does the utility of untruthful equilibria compare to the fully-informative equilibrium? How robust is the mechanism to collusion?
6. **Truthful, but Useful?:** H-DIPP is strictly *truthful*, but whether or not this is useful is a different empirical question. If H-DIPP-incentivized review scores vary widely even with post-processing (i.e., truthful reviews end up being quite idiosyncratically subjective), or do not track other measures of paper quality, it may not be possible to extract a meaningful signal (criticism C4). Running small scale experiments with real money could shed light on whether the benefits of our proposed mechanism warrant the overhead.

Such issues merit further in-depth discussion and analysis. Our proposed mechanism is intended to serve as a theoretically well-grounded solution to incentive issues in scientific peer review. Ultimately, experimental validation will be the final arbiter of the plausibility of our approach.

## Acknowledgments and Disclosure of Funding

We thank Sandesh Adhikary for assistance with figures in the paper.

## References

- Agarwal, A., Mandal, D., Parkes, D. C., and Shah, N. (2017). Peer prediction with heterogeneous users. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 81–98.
- Alberts, B., Hanson, B., and Kelner, K. L. (2008). Reviewing peer review.
- Arns, M. (2014). Open access is tiring out peer reviewers. *Nature*, 515(7528):467–467.
- Bengio, Y. (2020). Time to rethink the publication process in machine learning.
- Cai, Y., Daskalakis, C., and Papadimitriou, C. (2015). Optimum statistical estimation with strategic data sources. In *Conference on Learning Theory*, pages 280–296.
- Carvalho, A., Dimitrov, S., and Larson, K. (2013). Inducing honest reporting without observing outcomes: An application to the peer-review process. *arXiv preprint arXiv:1309.3197*.
- Charlin, L. and Zemel, R. (2013). The toronto paper matching system: an automated paper-reviewer assignment system.
- Cortes, C. and Lawrence, N. D. (2021). Inconsistency in conference peer review: Revisiting the 2014 neurips experiment. *arXiv preprint arXiv:2109.09774*.
- Dasgupta, A. and Ghosh, A. (2013). Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330.
- de la Rosa, J. L. and Szymanski, B. K. (2007a). Citation auctions as a method to improve selection of scientific papers. In *2007 2nd International Conference on Digital Information Management*, volume 1, pages 479–486. IEEE.
- de la Rosa, J. L. and Szymanski, B. K. (2007b). Selecting scientific papers for publication via citation auctions. *IEEE Intelligent Systems*, 22(6):16–20.
- Debmalya, M., Goran, R., and David, P. (2020). The effectiveness of peer prediction in long-term forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2160–2167.
- Faltings, B. and Radanovic, G. (2017). Game theory for data science: Eliciting truthful information. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 11(2):1–151.
- Fox, J. and Petchev, O. L. (2010). Pubcreds: fixing the peer review process by “privatizing” the reviewer commons. *The Bulletin of the Ecological Society of America*, 91(3):325–333.
- Frank, R. H. (2016). *Success and Luck*. Princeton University Press.
- Frijters, P. and Torgler, B. (2019). Improving the peer review process: a proposed market system. *Scientometrics*, 119(2):1285–1288.
- Gao, A., Wright, J. R., and Leyton-Brown, K. (2016). Incentivizing evaluation via limited access to ground truth: Peer-prediction makes things worse. *arXiv preprint arXiv:1606.07042*.
- Gao, X. A., Mao, A., Chen, Y., and Adams, R. P. (2014). Trick or treat: putting peer prediction to the test. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 507–524.
- Gneezy, U. and Rustichini, A. (2000). Pay enough or don’t pay at all. *The Quarterly journal of economics*, 115(3):791–810.
- Jecmen, S., Zhang, H., Liu, R., Shah, N., Conitzer, V., and Fang, F. (2020). Mitigating manipulation in peer review via randomized reviewer assignments. *Advances in Neural Information Processing Systems*, 33:12533–12545.
- Kong, Y. and Schoenebeck, G. (2018). Eliciting expertise without verification. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 195–212.

- Kong, Y. and Schoenebeck, G. (2019). An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. *ACM Transactions on Economics and Computation (TEAC)*, 7(1):1–33.
- LeCun, Y. (2012). A new publishing model in computer science.
- Lipton, Z. C. and Steinhardt, J. (2018). Troubling trends in machine learning scholarship. *arXiv preprint arXiv:1807.03341*.
- Liu, Y. and Chen, Y. (2016). Learning to incentivize: Eliciting effort via output agreement. *arXiv preprint arXiv:1604.04928*.
- Merriman, B. (2020). Peer review as an evolving response to organizational constraint: Evidence from sociology journals, 1952–2018. *The American Sociologist*, pages 1–26.
- Miller, N., Resnick, P., and Zeckhauser, R. (2005). Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373.
- NeurIPS (2019). What we learned from neurips 2019 data. *Medium*.
- Noothigattu, R., Shah, N., and Procaccia, A. (2021). Loss functions, axioms, and peer review. *Journal of Artificial Intelligence Research*, 70:1481–1515.
- Prelec, D. (2004). A bayesian truth serum for subjective data. *science*, 306(5695):462–466.
- Prüfer, J. and Zetland, D. (2010). An auction market for journal articles. *Public Choice*, 145(3-4):379–403.
- Radanovic, G. and Faltings, B. (2013). A robust bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI'13)*, number CONF, pages 833–839.
- Radanovic, G. and Faltings, B. (2014). Incentives for truthful information elicitation of continuous signals. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28.
- Radanovic, G., Faltings, B., and Jurca, R. (2016). Incentives for effort in crowdsourcing using the peer truth serum. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(4):1–28.
- Schoenebeck, G. and Yu, F.-Y. (2020). Two strongly truthful mechanisms for three heterogeneous agents answering one question. In *International Conference on Web and Internet Economics*, pages 119–132. Springer.
- Sculley, D., Snoek, J., and Wiltschko, A. (2018a). Avoiding a tragedy of the commons in the peer review process. *arXiv preprint arXiv:1901.06246*.
- Sculley, D., Snoek, J., Wiltschko, A., and Rahimi, A. (2018b). Winner’s curse? on pace, progress, and empirical rigor.
- Shah, N. B. (2019). Principled methods to improve peer review. *Retrieved May, 27:2020*.
- Shah, N. B., Tabibian, B., Muandet, K., Guyon, I., and Von Luxburg, U. (2018). Design and analysis of the nips 2016 review process. *Journal of machine learning research*.
- Squazzoni, F., Bravo, G., and Takács, K. (2013). Does incentive provision increase the quality of peer review? an experimental study. *Research Policy*, 42(1):287–294.
- Stelmakh, I., Shah, N. B., Singh, A., and Daumé III, H. (2020). A novice-reviewer experiment to address scarcity of qualified reviewers in large conferences. *arXiv preprint arXiv:2011.15050*.
- Vazire, S. (2017). Quality uncertainty erodes trust in science. *Collabra: Psychology*, 3(1).
- Waggoner, B. and Chen, Y. (2013). Information elicitation sans verification. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content (SC13)*.
- Waggoner, B. and Chen, Y. (2014). Output agreement mechanisms and common knowledge. In *HCOMP*.

- Wang, J. and Shah, N. B. (2018). Your 2 is my 1, your 3 is my 9: Handling arbitrary miscalibrations in ratings. *arXiv preprint arXiv:1806.05085*.
- Witkowski, J., Bachrach, Y., Key, P., and Parkes, D. C. (2014). Dwelling on the negative: Incentivizing effort in peer prediction.
- Witkowski, J. and Parkes, D. C. (2012a). Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981.
- Witkowski, J. and Parkes, D. C. (2012b). A robust bayesian truth serum for small populations.
- Witkowski, J. and Parkes, D. C. (2013). Learning the prior in minimal peer prediction. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce*, volume 14.
- Xiao, Y., Dörfler, F., and Van Der Schaar, M. (2018). Incentive design in peer review: Rating and repeated endogenous matching. *IEEE Transactions on Network Science and Engineering*, 6(4):898–908.
- Xu, Y., Zhao, H., Shi, X., Zhang, J., and Shah, N. B. (2018). On strategyproof conference peer review. *arXiv preprint arXiv:1806.06266*.
- Zhao, S. (2012). Two proposals for a new publishing and reviewing system in computer science.

## A Table of Notation

Symbol	Description
$N$	Number of papers submitted to the mechanism
$K$	Number of agents participating in the mechanism
$T$	Number of criteria reviewers must evaluate a paper on
$i$	Index for a paper
$k$	Index for a reviewer
$t$	Index for a criterion
$\omega_i$	The $i$ -th paper submitted to the mechanism
$s_k^{(\cdot)}$	Agent $k$ 's Stage I/II strategy
$s_k$	Agent $k$ 's full Stage I and Stage II strategy
$\mathbf{s}$	Full strategy profile of all agents on Stage I and Stage II
$\mathbf{s}_{-k}$	Strategy profile of all agents except agent $k$
$U_k^{(\cdot)}$	Agent $k$ 's Stage I/II utility
$U_k$	Agent $k$ 's total utility
$b_i$	The bid accompanying paper $\omega_i$ for a review slot
$c_i$	The price charged for accepting paper $\omega_i$ for review
$b_p^*$	The $p$ -th highest bid
$\omega_p^*$	The paper with the $p$ -th highest bid
$\delta_i$	Boolean denoting whether paper $\omega_i$ is accepted <i>for publication</i> after review
$\eta_i$	Probability that paper $\omega_i$ is accepted <i>for publication</i> conditional on review
$v_k$	Agent $k$ 's value for a paper accepted for publication
$P$	The number of papers accepted for review
$A, B, C$	The reviewer triplet assigned to a given a paper
$\Sigma_t$	Set of allowed review scores on criterion $t$
$e_k(t)$	Reviewer $k$ 's cost for completing up to $t$ criteria
$\mathbb{P}$	A reviewer triplet's common prior over effort-informed scores
$\mathbb{Q}$	A reviewer triplet's joint distribution over reported scores and effort-informed scores
$\mathbf{X}_k^{[t]}$	Reviewer $k$ 's true review scores on criteria 1,..., $t$
$g_k^{t' t}(\mathbf{x}_k^{[t]})$	Reviewer $k$ 's 'best guess' of their signal on criterion $t'$ given signals $\mathbf{x}_k^{[t]}$ (where $t' > t$ )
$\hat{\mathbf{X}}_k^{[T]}$	Reviewer $k$ 's (potentially strategically) reported review scores on criteria 1,..., $T$
$\theta_k$	Transition matrix representing reviewer $k$ 's signal reporting strategy
$\hat{\mathbf{p}}_{k' \leftarrow k}^{[T]}$	Probability distribution of reviewer $k$ 's first prediction of $k'$ 's scores on criteria 1,..., $T$
$\hat{\mathbf{p}}_{k' \leftarrow k}^{+}$	Probability distribution of reviewer $k$ 's second prediction of $k'$ 's scores on criteria 1,..., $T$
$r_k$	Reviewer $k$ 's payment for reviewing
$R_k(t)$	Reviewer $k$ 's expected payment for completing $t$ criteria when peers complete $T$ criteria and report truthfully
$R_k^{\text{role}}(t)$	Reviewer $k$ 's expected expert/target payment for completing $t$ criteria when peers complete $T$ criteria and report truthfully
$R_{k^{t'}}^{\text{role}}(t)$	Reviewer $k$ 's expected expert/target payment <i>on criterion <math>t'</math></i> for completing $t$ criteria when peers complete $T$ criteria and report truthfully
$\Delta R_k(t)$	Reviewer $k$ 's marginal expected payment for completing criterion $t$ when peers complete $T$ criteria and report truthfully
$\Delta R_k^{\text{role}}(t)$	Reviewer $k$ 's marginal expected expert/target payment for completing criterion $t$ when peers complete $T$ criteria and report truthfully
$\Delta R_{k^{t'}}^{\text{role}}(t)$	Reviewer $k$ 's marginal expected expert/target payment <i>on criterion <math>t'</math></i> for completing criterion $t$ when peers complete $T$ criteria and report truthfully
$S_{\log}(\cdot, \cdot)$	Log proper scoring rule
$I(\cdot)$	Mutual information

## B The 2014 NeurIPS Experiment

In 2014, program chairs of the NeurIPS conference sought to study the consistency of the peer review process (Cortes and Lawrence, 2021), and set up an experiment where 10% of submitted papers underwent review by two independent committees. We describe the major findings below:

- **Inconsistency of Reviews:** About 26% of papers received different outcomes from the two committees (compared to 37.5% expected from random acceptance).
- **Accept Precision:** Of the accepted papers, only about 50% of papers would have survived re-review by an independent committee (compared to 25% expected from random acceptance).
- **Reject Precision:** Of the rejected papers, about 80% would have also been rejected in a re-review by an independent committee (compared to 75% expected from random rejection).

These findings showed that the *accept/reject* decisions were moderately better than pure chance, although still a fair bit noisy. Some of this is unavoidable in any highly selective mechanism seeking to allocate a limited number of spots to a large number of candidates on the basis of some noisily observed measure of ‘quality’ (Frank, 2016). When a mechanism seeks to accept some small percentage of top quality candidates, even a small amount of noise in observing the quality of a candidate can add a great deal of randomness to which candidates make the cut. Yet it will be the case that the bulk of the randomness will be faced by the candidates clustered around the cutoff (‘messy middle’); candidates well above the cutoff will likely be consistently accepted, and candidates well below the cutoff will likely be consistently rejected. Thus, the noise in observations of paper quality is an important factor affecting the randomness faced by authors in a highly selective acceptance process. Shah et al. (2018) find that about 45% of NeurIPS 2015 papers and 30% of NeurIPS 2016 papers are in ‘messy middle’ where ‘inter-reviewer agreements are near random’. Back in the 2014 NeurIPS experiment, an analysis of the independent committees’ ratings of paper quality, finds the following:

- **Correlation of Quality Scores:** The correlation between calibrated ‘quality’ scores assigned by the two independent committees (adjusted for a reviewers’ systematic biases) was 0.546.
- **Citation Impact:** For accepted papers, there was little correlation (0.051) between reviewers’ assessments of paper quality and papers’ log-citation count (likely due to the restricted range effect). For rejected papers, there was a positive correlation (0.22) between reviewers’ assessments of paper quality and papers’ log citation count.

The finding that about half the variance in scores is intrinsic to the paper (‘objective’) and half is idiosyncratic to the reviewer (‘subjective’) (Cortes and Lawrence, 2021) provides more information on the noise in observations. This subjectivity can arise from more readily identifiable sources like differences in effort or expertise, or from truly subjective differences of perspective. Naturally, if reviewers exert different levels of effort or have differing levels of expertise, their reviews may be systematically different from each other and would weaken the ‘objective’ proportion of review scores. The correlation between quality scores and impact also suggests that the objective proportion of review scores has some better-than-random ability to predict the long-term impact of papers as measured by citations (relevant for criticism C4). In sum, the above results on accept/reject decisions as well as reviewers’ quality scores paint a picture of a peer review system that performs noticeably better than random, but with a lot of room for improvement.

These results provide a baseline against which to judge the performance of our proposed mechanism. The goal then is for our proposed mechanism to incentivize reviewers to coordinate into a high-effort equilibrium where the ‘subjectivity’ of reviewers’ scores is lower, the correlation between independent reviewers’ quality scores is higher, and the messy middle is smaller. The hypothesis is that the provision of incentives for effort *during the review process* may help reviewers better observe the quality of papers, and reduce the variance of decisions under the current *selective acceptance process*.

Finally, we note that conference selectiveness serves the dual purposes of focusing the community’s attention on major results (motivation M2), as well as building reputation for authors who do high-quality work (motivation M3). While Cortes and Lawrence (2021) argues against using conference publications as a proxy for research quality or career advancement, decision-makers looking to make hiring decisions or determine what research to focus on will implicitly or explicitly be seeking relevant signals; in the absence of an alternative signal, we argue that it is important to ensure a peer-reviewed publication gives an appropriate signal of quality.

## C Key elements of the Mutual Information Paradigm (Kong and Schoenebeck, 2019)

A scoring rule  $S: \Sigma \times \Delta_\Sigma \rightarrow \mathbb{R}$  is a function that accepts some outcome  $\sigma$  (from some signal space  $\Sigma$ ) and a probability distribution  $\delta$  (representing a forecast over that signal space  $\Delta_\Sigma$ ) and scores the forecast with a real number. The scoring rule is *proper* if whenever the outcome  $\sigma$  is drawn from a distribution  $\delta \in \Delta_\Sigma$ , the expected score  $\mathbb{E}_\sigma[S(\sigma, \cdot)]$  is maximized when the forecast is  $\delta$ . The scoring rule is *strictly proper* if this maximum is unique. Our mechanism will use the log scoring rule  $S_{\log}(\sigma, \delta) = \ln \delta(\sigma)$ , which is a strictly proper scoring rule. Next, recall that the Shannon Mutual Information of two random variables is defined as the KL divergence of the product of their marginal distributions from the joint distribution:  $I(X;Y) = D_{\text{KL}}(P(X,Y) || P(X) \otimes P(Y))$ . The conditional Shannon mutual information is given as  $I(X;Y|Z) = \sum_z P(z) D_{\text{KL}}(P(X,Y|z) || P(X|z) \otimes P(Y|z)) = \sum_z P(z) I(X;Y|z)$ . Kong and Schoenebeck (2019) show an interesting connection between conditional Shannon mutual information and the log scoring rule:

$$I(X;Y|Z) = \mathbb{E}_{X,Y,Z}[S_{\log}(Y, P(Y|X,Z)) - S_{\log}(Y, P(Y|Z))] \quad (10)$$

Finally, we present the important data processing inequality, typically useful in proving that deviating from truth-telling decreases payment. The inequality states that whenever a signal  $Y$  is manipulated in a way independent of another signal  $X$ , the manipulated signal can only have lower mutual information with  $X$ .

**Lemma 3** (Strict Data Processing Inequality (Schoenebeck and Yu, 2020)). *If  $(X, Y)$  on a finite space  $\mathcal{X} \times \mathcal{Y}$  is stochastic relevant (i.e., for any distinct  $x, x' \in \mathcal{X}$ ,  $P(Y|x) \neq P(Y|x')$ ) and has full support, then for any random function  $\theta: \mathcal{Y} \rightarrow \mathcal{Y}$  where the randomness is independent of  $(X, Y)$ , we have that  $I(X; \theta(Y)) \leq I(X; Y)$ . We have equality if and only if  $\theta$  is an injective map.*

**Lemma 4** (Strict Conditional Data Processing Inequality (Schoenebeck and Yu, 2020)). *If  $(W, X, Y)$  on a finite space  $\mathcal{W} \times \mathcal{X} \times \mathcal{Y}$  is second-order stochastic relevant (i.e., for any distinct  $x, x' \in \mathcal{X}$ , there is some  $w \in \mathcal{W}$  such that  $P(Y|w, x) \neq P(Y|w, x')$ ) and has full support, then for any random function  $\theta: \mathcal{Y} \rightarrow \mathcal{Y}$  where the randomness is independent of  $(W, X, Y)$ , we have that  $I(X; \theta(Y)|W) \leq I(X; Y|W)$ . We have equality if and only if  $\theta$  is an injective map.*

## D Relationship to HMIM and TDPP

### D.1 Relationship to HMIM

Here we discuss how our mechanism differs from the single-HMIM proposal under the HMIP framework (Kong and Schoenebeck, 2018).

- **Compulsory Reports:** Unlike single-HMIM, reviewers must provide scores for all criteria, not just the ones they have completed. This is currently the norm in peer review, and we would like to design the mechanism around this expectation. We model reviewers' alternative to effortful reports as 'best guess' reports, as opposed to HMIM where a reviewer who does not complete a criterion does not report any signal and simply forgoes any payment on it. Realistically, a reviewer's signals on completed criteria may enable them to make reasonable guesses on higher level criteria, and we account for this in computing reviewers' marginal payments and setting hyperparameters.
- **Partial Vs Total Order:** Kong and Schoenebeck (2018) only impose a partial order on their criteria hierarchy to keep it general; we impose a total order for ease of exposition.
- **Different Payment Scheme:** While we adopt the hierarchical criteria structure from HMIP, the actual payments we provide are adopted from TDPP. However, at equilibrium, a reviewer's expected target payment is almost the mutual information payment scheme proposed in HMIP, under our conditional independence assumption (Assumption 2). The difference is that while HMIP (under the Conditional Independence assumption) would compute the mutual information of the target's and source's signals on criterion  $t$  conditioned on the source's signals on lower level criteria, our mechanism also conditions this on the expert's private information.

- **Strategy and Equilibria:** Our mechanism, like single-HMIM, is strictly truthful. The ‘idealized’ HMIP has the desirable property that truthful reporting is a dominant strategy, but the practical single-HMIM algorithm does not have this property. Truthful reports are not a dominant strategy in our H-DIPP mechanism either, since target payments depend on the prediction strategy of experts. However, it is the case that a reviewer is paid greater target payments in truth-telling equilibria, and the use of proper scoring rules incentivizes experts to predict truthfully if targets deviate to truth-telling.

## D.2 Relationship to TDPP

Here we discuss how our mechanism differs from the TDPP proposal (Schoenebeck and Yu, 2020).

- **Playing All Roles:** TDPP assigns each agent a single role (expert, target, or source). In our mechanism, every agent plays all three roles.
- **Effort Model:** TDPP assumes that signals are revealed to agents at no cost. We assume that the effort level determines the signals that are revealed to agents, so we adopt the hierarchical approach of HMIP to incentivize the effort needed to respond to higher level criteria.
- **Knowledge of Common Prior:** TDPP and other ‘detail-free’ mechanisms assume no knowledge of the common prior. While we do not strictly need to know a reviewer triplet’s common-prior to run the mechanism, we would like to know certain properties (the marginal payments from Lemma 1) of the common prior (as in HMIP) so that we can set the mechanism hyperparameters to incentivize the desired fully informative equilibrium.
- **Multi-criteria Predictions and No Source Payment:** TDPP only elicits source and target responses on a single question, while H-DIPP elicits responses to a hierarchy of criteria. H-DIPP reveals the source’s signal after each criterion (as opposed to revealing it all at once, or revealing target signals, etc.) as this brings the expected target payment close to the HMIP recommendation. However, this does mean that unlike TDPP, we do not provide any source payment for the following reason. The TDPP expected source payment was based on the expert  $k$ ’s first prediction of target  $k$ ’s signals on a criterion  $t$  (made without the source’s information). This term did not change the source’s incentives in the original TDPP (as their actions could not affect the expert’s first prediction), but it was a convenient accounting trick to make the ex ante agent welfare equal to the conditional mutual information of the source and target’s reports. However, we cannot incorporate such a source payment in H-DIPP: the expert makes predictions with the source’s signals on lower level criteria, so *the source can influence the expert’s first prediction on every criterion  $t > 1$* . Thus, including such a source payment could lead the source to strategically misreport so that the expert’s first predictions are poor. Naturally, leaving out source payments preserves source’s incentives to report truthfully, but the ex ante reviewer welfare is no longer the conditional mutual information. This is why there are ‘extra’ entropy terms in Lemma 2.
- **Strategy and Equilibria:** Like TDPP, our H-DIPP mechanism is strictly truthful. However, TDPP has the additional very desirable property of being *strongly truthful*, i.e., the truth-telling equilibrium pays more than all other equilibria, so agents have no reason to coordinate into other equilibria. Owing to the multi-round structure of H-DIPP, we cannot quite make the same guarantee; it *may* be possible to have equilibria where reviewers strategically misreport to increase a peer’s target payment by more than the loss of their target payment thereby increasing aggregate payment, although there is no individual incentive to do so. It may be possible for all reviewers to do this so their individual and aggregate payments are higher than the fully informative equilibrium, although accepting the increased target payment bestowed by a source’s strategic misreports and *not reciprocating* by reporting truthfully would lead to a higher-paying equilibrium for the target. Investigating such equilibria is an important direction for future work. Nevertheless, we can make the weaker claim that the fully informative equilibrium pays more than an effortless uninformative equilibrium (Theorem 3).

## E Proofs

**Theorem 1** (Strictly Truthful) *Suppose reviewers  $A, B, C$  complete  $[t_A], [t_B]$ , and  $[t_C]$  criteria respectively. Then, on shared complete criteria  $[t'] = [\min(t_A, t_B, t_C)]$ , the H-DIPP mechanism is*

strictly truthful, i.e., all reviewers honestly reporting their true predictions and private review scores on shared complete criteria  $[t']$  is a strict Bayesian Nash equilibrium.

*Proof of Theorem 1.* We adapt the proof of strict truthfulness of TDPP by Schoenebeck and Yu (2020) for H-DIPP. We show that agents reporting their honest predictions and scores on shared complete criteria is a strict Bayesian Nash equilibrium; importantly, payments for each criterion are independent and can be strategized about independently, hence we can reason strictly about strategies on the shared complete criteria  $\mathbf{x}_k^{[t']} \subseteq \mathbf{x}_k^{[t_k]}$ . Reviewer  $A$ 's private signals on incomplete criteria  $t > t_A$  are  $\mathbf{x}_A^t = g_A^{t|t_A}(\mathbf{x}_A^{[t_A]})$  for some best guess function  $g$ , and similarly for reviewers  $B$  and  $C$ . Since all reviewers play the role of expert, target, and source, we only analyze reviewer  $A$ 's incentives when peers  $B$  and  $C$  report truthfully; these incentives are symmetrical for all reviewers. Additionally, we hold the number of completed criteria and hence effort cost  $e_k(t_k)$  constant, and look to optimize payoff under this condition.

Suppose reviewers  $B$  and  $C$  report their private signals truthfully, i.e.,  $\hat{x}_B^t = x_B^t$  and  $\hat{x}_C^t = x_C^t$ , and reviewer  $C$  as expert truthfully predicts reviewer  $A$ 's score as the posterior from updating the common prior  $\hat{\mathbf{p}}_{A^t \leftarrow C} = \mathbb{P}(X_A^t | \mathbf{x}_C^{[t']}, \mathbf{x}_B^{[t-1]})$ , and  $\hat{\mathbf{p}}_{A^t \leftarrow C}^+ = \mathbb{P}(X_A^t | \mathbf{x}_C^{[t']}, \mathbf{x}_B^{[t]})$ .

Now, reviewer  $A$ 's expected payment (taken with respect to the joint distribution  $\mathbb{Q}(\hat{\mathbf{X}}_A^{[T]}, \mathbf{X}_A^{[t']}, \mathbf{X}_B^{[t']}, \mathbf{X}_C^{[t']})$  of true signals and  $A$ 's potentially strategic reports) on the shared complete criteria can be written as:

$$\mathbb{E}_{\mathbb{Q}}[r_{A^t}] = \alpha_A^t \mathbb{E} \left[ \underbrace{S_{\log}(x_B^t, \hat{\mathbf{p}}_{B^t \leftarrow A}) + S_{\log}(x_B^t, \hat{\mathbf{p}}_{B^t \leftarrow A}^+)}_{\text{expert payment}} \right] + \beta_A^t \mathbb{E} \left[ \underbrace{S_{\log}(\hat{x}_A^t, \mathbb{P}(X_A^t | \mathbf{x}_C^{[t']}, \mathbf{x}_B^{[t]})) - S_{\log}(\hat{x}_A^t, \mathbb{P}(X_A^t | \mathbf{x}_C^{[t']}, \mathbf{x}_B^{[t-1]}))}_{\text{target payment}} \right] \quad (11)$$

Now, reviewer  $A$ 's payment on a given criterion  $t$  is independent of their reported signals and predictions on other criteria, so the sum of payments on complete criteria can be maximized by maximizing the expert payment and target payment on each criterion:

*Expert Payment* The expected expert payment consists of reviewer  $A$ 's predictions of their target  $B$ 's signal on each criterion  $t$  first with the source  $C$ 's signals on criteria less than  $t$  and then with  $C$ 's signal on criterion  $t$  as well. Predictions on each criterion are scored independently, so a strategic prediction on one criterion cannot impact future payments on other criteria. Since we use a strictly proper scoring rule to score both predictions on each criterion, reviewer  $A$ 's expected payoff for each criterion is strictly maximized when they report their true posterior belief over their target  $B$ 's signals under the common prior, i.e.,  $\hat{\mathbf{p}}_{B^t \leftarrow A}^+ = \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t]})$  and  $\hat{\mathbf{p}}_{B^t \leftarrow A} = \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t-1]})$ .

*Target Payment* Now, consider the reviewer  $A$ 's expected target payment on criteria  $t \leq t'$ . We observe that we can model the strategies for responses to different criteria independently since the target payments are also independent for different criteria. Let  $\hat{x}_A^t = \theta_{A^t}(\mathbf{x}_A^{[t_A]})$  be reviewer  $A$ 's *deterministic best response* report (on criterion  $t$ ) to reviewer  $B$ 's honest signal reports  $\hat{x}_B^{[T]}$  and reviewer  $C$ 's honest (posterior) predictions  $\hat{\mathbf{p}}_{A^t \leftarrow C}^+ = \mathbb{P}(X_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t]})$ ,  $\hat{\mathbf{p}}_{A^t \leftarrow C} = \mathbb{P}(X_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]})$ .

For  $t \leq t'$ , reviewer  $A$ 's expected target payment for playing this strategy  $\hat{x}_A^t = \theta_{A^t}(\mathbf{x}_A^{[t']})$  is:

$$\begin{aligned} u_t(\theta) &:= \mathbb{E} \left[ S_{\log}(\hat{x}_A^t, \mathbb{P}(X_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t]})) - S_{\log}(\hat{x}_A^t, \mathbb{P}(X_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]})) \right] \\ &= \mathbb{E}_{\mathbb{Q}} \left[ \log \frac{\mathbb{P}(\hat{x}_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t]})}{\mathbb{P}(\hat{x}_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]})} \right] \\ &= \mathbb{E}_{\mathbb{Q}} \left[ \log \frac{\mathbb{P}(\hat{x}_A^t, x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]})}{\mathbb{P}(\hat{x}_A^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}) \mathbb{P}(x_B^t | \mathbf{x}_C^{[t']}, \mathbf{x}_B^{[t-1]})} \right] \\ &= \mathbb{E}_{\mathbb{P}} \left[ \log \frac{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, \theta_{A^t}(\mathbf{x}_A^{[t_A]}))}{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]})} \right] \end{aligned} \quad (12)$$

By a similar calculation, we have that reviewer  $A$ 's expected payment under a truth-telling strategy is  $u_t(\tau) := \mathbb{E}_{\mathbb{P}} \left[ \log \frac{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, x_A^t)}{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]})} \right]$ . Then, the difference between truth-telling payment and the best response payment is:

$$\begin{aligned}
u_t(\tau) - u_t(\theta) &= \mathbb{E}_{\mathbb{P}} \left[ \log \frac{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, x_A^t)}{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, \theta_{A^t}(\mathbf{x}_A^{[t_A]}))} \right] \\
&= \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^{[t-1]}, \mathbf{x}_C^{[t_C]}} \left[ \mathbb{E}_{X^t} \left[ \log \frac{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, x_A^t)}{\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, \theta_{A^t}(\mathbf{x}_A^{[t_A]}))} \right] \middle| \mathbf{x}_A^{[t_A]} = \mathbf{x}_A^{[t_A]}, \mathbf{x}_B^{[t-1]} = \mathbf{x}_B^{[t-1]}, \mathbf{x}_C^{[t_C]} = \mathbf{x}_C^{[t_C]} \right] \\
&= \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^{[t-1]}, \mathbf{x}_C^{[t_C]}} \left[ \text{KL} \left( \mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, x_A^t) \middle| \mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, \theta_{A^t}(\mathbf{x}_A^{[t_A]})) \right) \right] \tag{13}
\end{aligned}$$

Since  $\text{KL}$  is a divergence, it is greater than or equal to zero, so we must have  $u_t(\tau) - u_t(\theta) \geq 0$ . Finally, observe that if the strategy  $\theta$  is not truthful, i.e.,  $\theta_{A^t}(\mathbf{x}_A^{[t_A]}) \neq x_A^t$ , then by the second-order stochastic relevance there is some  $x_C^t \in \mathbf{x}_C^{[t_C]}$  for which  $\mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, x_A^t) \neq \mathbb{P}(x_B^t | \mathbf{x}_C^{[t_C]}, \mathbf{x}_B^{[t-1]}, \hat{x}_A^t)$ , so the  $\text{KL}$  divergence is not identically zero over all signals in the expectation. Consequently, the inequality is strict, implying  $u(\tau) > u(\theta)$ . Thus, the best response strategy is truthtelling ( $\theta = \tau$ ), and reviewer  $A$  strictly maximizes their target payment by truthfully reporting their signal  $\hat{x}_A^t = x_A^t$  on all complete criteria.

We have shown that when peers report their predictions and private review scores truthfully, a reviewer should truthfully report their prediction to strictly maximize their expert payment and truthfully report their private review score to strictly maximize their target payment. These incentives are symmetrical for reviewers  $B$  and  $C$ , so truthful reporting is a strict Bayesian Nash Equilibrium.  $\square$

**Lemma 1** Suppose reviewers  $B$ ,  $C$  complete all criteria and truthfully report their private review scores and predictions. Then, reviewer  $A$ 's marginal expected payment for completing criterion  $t_A$  and honestly reporting private review scores and predictions is  $\Delta R_A(t_A) := R_A(t_A) - R_A(t_A - 1) = \Delta R_A^{\text{expert}}(t_A) + \Delta R_A^{\text{target}}(t_A)$ , where:

1.  $\Delta R_A^{\text{expert}}(t_A) := R_A^{\text{expert}}(t_A) - R_A^{\text{expert}}(t_A - 1) = \sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{\text{expert}}(t_A)$
2.  $\Delta R_A^{\text{target}}(t_A) := R_A^{\text{target}}(t_A) - R_A^{\text{target}}(t_A - 1) = \sum_{t=1}^T \beta_A^t \Delta R_{A^t}^{\text{target}}(t_A)$
3.  $\Delta R_{A^t}^{\text{expert}}(t_A) := I(X_B^t; X_A^{t_A} | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t]}) + I(X_B^t; X_A^{t_A} | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t-1]})$
4.  $\Delta R_{A^t}^{\text{target}}(t_A) := \begin{cases} 0, & t < t_A \\ I(X_B^t; X_A^{t_A} | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) - I(X_B^t; g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}), & t = t_A \\ \mathbb{E} \left[ \text{KL} \left( \mathbb{P}(X_B^t | g_A^{t|t_A}(\mathbf{x}_A^{[t_A]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]}) \middle| \mathbb{P}(X_B^t | g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]}) \right) \right], & t > t_A \end{cases}$

Furthermore, the marginal expected expert and target payments on each criterion are always non-negative, i.e.,  $\Delta R_{A^t}^{\text{expert}}(t_A), \Delta R_{A^t}^{\text{target}}(t_A) \geq 0$  for all  $t, t_A \in [T]$  and the inequality is strict when  $t = t_A$ . Similar results hold for reviewers  $B$  and  $C$  when peers complete all criteria and report truthfully.

*Proof.* Assuming peers  $B$ ,  $C$  complete all  $T$  criteria and report truthfully, we can write reviewer  $A$ 's strategy profile of effort  $e_A(t_A)$  and truthful reports as  $s_A = (e_A(t_A), \mathbf{x}_k^{[t_A]} \cup \{g_A^{t'|t_A}(\mathbf{x}_A^{[t_A]})\}_{t' \in [t_A+1, T]}, \{\mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t]})\}, \{\mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t-1]})\}_{t \in [T]})$ , and the expected payment is:

$$\begin{aligned}
R_A &= \underbrace{\sum_{t=1}^T \alpha_A^t \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t]}} \left[ \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t]}) \right) + \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t-1]}) \right) \right]}_{R_A^{\text{expert}}} \\
&\quad + \underbrace{\sum_{t=1}^T \beta_A^t \mathbb{E}_{\mathbf{x}_A^t, \mathbf{x}_B^{[t]}, \mathbf{x}_C^{[T]}} \left[ \log \left( \frac{\mathbb{P}(X_A^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t]})}{\mathbb{P}(X_A^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right) \right]}_{R_A^{\text{target}}} \tag{14}
\end{aligned}$$

The expected payment  $R_A$  is technically a function of all reviewers' completed criteria  $R_A(t_A, t_B, t_C)$ , but we simply write it as  $R_A(t_A)$  as other reviewers' are assumed to have completed  $T$  criteria. The marginal payment for completing criterion  $t_A$  is  $R_A(t_A) - R_A(t_A - 1) = \Delta R_A(t_A) = \Delta R_A^{\text{expert}}(t_A) + \Delta R_A^{\text{target}}(t_A)$ . The marginal expert payment  $\Delta R_A^{\text{expert}}(t_A)$  is:

$$\begin{aligned}
\Delta R_A^{\text{expert}}(t_A) &= R_A^{\text{expert}}(t_A) - R_A^{\text{expert}}(t_A - 1) \\
&= \sum_{t=1}^T \alpha_A^t \left( \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t]}} \left[ \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t]}) \right) + \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t-1]}) \right) \right] \right) \\
&\quad - \sum_{t=1}^T \alpha_A^t \left( \mathbb{E}_{\mathbf{x}_A^{[t_A-1]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t]}} \left[ \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t]}) \right) + \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t-1]}) \right) \right] \right) \\
&= \sum_{t=1}^T \alpha_A^t \left( \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t]}} \left[ \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t]}) \right) - \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t]}) \right) \right] \right. \\
&\quad \left. + \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t-1]}} \left[ \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t-1]}) \right) - \log \left( \mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t-1]}) \right) \right] \right) \\
&= \sum_{t=1}^T \alpha_A^t \left( \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t]}} \left[ \log \left( \frac{\mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t]})}{\mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t]})} \right) \right] + \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^t, \mathbf{x}_C^{[t-1]}} \left[ \log \left( \frac{\mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A]}, \mathbf{x}_C^{[t-1]})}{\mathbb{P}(X_B^t | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t-1]})} \right) \right] \right) \\
&= \sum_{t=1}^T \alpha_A^t \left( I(X_B^t; X_A^{t_A} | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t]}) + I(X_B^t; X_A^{t_A} | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t-1]}) \right)
\end{aligned} \tag{15}$$

Defining  $\Delta R_{A^t}^{\text{expert}}(t_A) := I(X_B^t; X_A^{t_A} | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t]}) + I(X_B^t; X_A^{t_A} | \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t-1]})$ , we have  $\Delta R_A^{\text{expert}}(t_A) = \sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{\text{expert}}(t_A)$ . Next, we compute the marginal target payment as follows (splitting the sum over  $T$  criteria into three: marginal target payment on criteria less than  $t_A$ , criterion  $t_A$ , and criteria greater than  $t_A$ ):

$$\begin{aligned}
\Delta R_A^{\text{target}}(t_A) &= R_A^{\text{target}}(t_A) - R_A^{\text{target}}(t_A - 1) \\
&= \sum_{t=1}^{t_A-1} \beta_A^t \left( \mathbb{E} \left[ \log \left( \frac{\mathbb{P}(X_A^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t]})}{\mathbb{P}(X_A^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right) \right] - \mathbb{E} \left[ \log \left( \frac{\mathbb{P}(X_A^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t]})}{\mathbb{P}(X_A^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right) \right] \right) \\
&\quad + \beta_A^{t_A} \left( \mathbb{E} \left[ \log \left( \frac{\mathbb{P}(X_A^{t_A} | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A]})}{\mathbb{P}(X_A^{t_A} | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]})} \right) \right] - \mathbb{E} \left[ \log \left( \frac{\mathbb{P}(g_A^{t_A|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A]})}{\mathbb{P}(g_A^{t_A|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]})} \right) \right] \right) \\
&\quad + \sum_{t=t_A+1}^T \beta_A^t \left( \mathbb{E} \left[ \log \left( \frac{\mathbb{P}(g_A^{t|t_A}(\mathbf{x}_A^{[t]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t]})}{\mathbb{P}(g_A^{t|t_A}(\mathbf{x}_A^{[t]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right) \right] - \mathbb{E} \left[ \log \left( \frac{\mathbb{P}(g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t]})}{\mathbb{P}(g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right) \right] \right) \\
&= \beta_A^{t_A} \left( I(X_B^{t_A}; X_A^{t_A} | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) - I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) \right) \\
&\quad + \sum_{t=t_A+1}^T \beta_A^t \left( \mathbb{E} \left[ \log \frac{\mathbb{P}(X_B^t | g_A^{t|t_A}(\mathbf{x}_A^{[t]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})}{\mathbb{P}(X_B^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right] - \mathbb{E} \left[ \log \frac{\mathbb{P}(X_B^t | g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})}{\mathbb{P}(X_B^t | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right] \right) \\
&= \beta_A^{t_A} \left( I(X_B^{t_A}; X_A^{t_A} | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) - I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) \right) \\
&\quad + \sum_{t=t_A+1}^T \beta_A^t \left( \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^{[t]}, \mathbf{x}_C^{[T]}} \left[ \log \frac{\mathbb{P}(X_B^t | g_A^{t|t_A}(\mathbf{x}_A^{[t]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})}{\mathbb{P}(X_B^t | g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]})} \right] \right) \\
&= \beta_A^{t_A} \left( I(X_B^{t_A}; X_A^{t_A} | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) - I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{x}_A^{[t_A-1]}) | \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) \right) \\
&\quad + \sum_{t=t_A+1}^T \beta_A^t \left( \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \mathbf{x}_B^{[t-1]}, \mathbf{x}_C^{[T]}} \left[ \text{KL} \left( \mathbb{P}(X_B^t | g_A^{t|t_A}(\mathbf{x}_A^{[t]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]}) \parallel \mathbb{P}(X_B^t | g_A^{t|t_A-1}(\mathbf{x}_A^{[t_A-1]}), \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t-1]}) \right) \right] \right)
\end{aligned} \tag{16}$$

Defining  $\Delta R_{A^t}^{\text{target}}(t_A) = 0$  for  $t < t_A$ ,  $\Delta R_{A^t A}^{\text{target}}(t_A) = I(X_B^{t_A}; X_A^{t_A} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) - I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{X}_A^{[t_A-1]}) | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]})$ , and for  $t > t_A$  as  $\Delta R_{A^t}^{\text{target}} = \mathbb{E} [\text{KL}(\mathbb{P}(X_B^t | g_A^{t|t_A}(\mathbf{X}_A^{[t_A]}), \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t-1]})) || \mathbb{P}(X_B^t | g_A^{t|t_A-1}(\mathbf{X}_A^{[t_A-1]}), \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t-1]})]$ , we get  $\Delta R_A^{\text{target}}(t_A) = \sum_{t=1}^T \beta_A^t \Delta R_{A^t}^{\text{target}}(t_A)$ . Thus, we see that the marginal expected payment is:

$$\begin{aligned}
R_A(t_A) - R_A(t_A-1) &= \Delta R_A(t_A) \\
&= \Delta R_A^{\text{expert}}(t_A) + \Delta R_A^{\text{target}}(t_A) \\
&= \sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{\text{expert}}(t_A) + \beta_A^t \Delta R_{A^t}^{\text{target}}(t_A)
\end{aligned} \tag{17}$$

Finally we show non-negativity of marginal expected expert and target payments  $\Delta R_k^{\text{expert}}, \Delta R_k^{\text{target}} \geq 0$ , and positivity on criterion  $t_A$  when the marginal criterion completed is  $t_A$ , i.e.,  $\Delta R_{A^t A}^{\text{expert}}(t_A) > 0$  and  $\Delta R_{A^t A}^{\text{target}}(t_A) > 0$ .

We easily get the non-negativity of the marginal expected expert payment  $\Delta R_{A^t}^{\text{expert}}(t_A) \geq 0$  from the non-negativity of mutual information. Now consider the marginal expected expert payment on criterion  $t_A$  in particular ( $\Delta R_{A^t A}^{\text{expert}}(t_A)$ ). By second order stochastic relevance (Assumption 3), for any  $x_A^{t_A} \neq \tilde{x}_A^{t_A}$  there is some  $x_C^{[t_A]} \subset \mathbf{x}_C^{[t_A]}$  such that  $\mathbb{P}(X_B^{t_A} | x_A^{t_A}, \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t_A]}) \neq \mathbb{P}(X_B^{t_A} | \tilde{x}_A^{t_A}, \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[t_A]}) \Rightarrow X_B^{t_A} \not\perp\!\!\!\perp X_A^{t_A} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[t_A]}$ . Thus,  $I(X_B^{t_A}; X_A^{t_A} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[t_A]}) > 0$ . This gives us that  $\Delta R_{A^t A}^{\text{expert}}(t_A) > 0$  for all  $t_A \in [T]$ .

Next, we consider the marginal expected target payment. By the non-negativity of mutual information and KL divergence, we have that  $\Delta R_{A^t}^{\text{target}}(t_A) \geq 0$ . We show that the inequality is strict when  $t = t_A$ , i.e.,  $\Delta R_{A^t A}^{\text{target}}(t_A) = I(X_B^{t_A}; X_A^{t_A} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) - I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{X}_A^{[t_A-1]}) | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) > 0$  as follows:

$$\begin{aligned}
I(X_B^{t_A}; g_A^{t_A|t_A-1}(\mathbf{X}_A^{[t_A-1]}) | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) &\leq I(X_B^{t_A}; \mathbf{X}_A^{[t_A-1]} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) \\
&= I(X_B^{t_A}; \mathbf{X}_A^{[T]} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) - I(X_B^{t_A}; \mathbf{X}_A^{[t_A, T]} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) \\
&= I(X_B^{t_A}; X_A^{t_A} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) - I(X_B^{t_A}; \mathbf{X}_A^{[t_A, T]} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) \\
&< I(X_B^{t_A}; X_A^{t_A} | \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]})
\end{aligned} \tag{18}$$

For the first inequality, we used the information monotonicity of mutual information. For the equality in the second step, we used the chain rule of mutual information. For the equality in the third step, we use the conditional independence assumption (Assumption 2). For the final inequality, we used the fact that second-order stochastic relevance (Assumption 3) implies that for any  $x_A^{t_A} \neq \tilde{x}_A^{t_A}$  there is some  $x_C^{[T]} \subset \mathbf{x}_C^{[T]}$  such that  $\mathbb{P}(X_B^{t_A} | x_A^{t_A}, \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) \neq \mathbb{P}(X_B^{t_A} | \tilde{x}_A^{t_A}, \mathbf{x}_A^{[t_A-1]}, \mathbf{x}_C^{[T]}, \mathbf{x}_B^{[t_A-1]}) \Rightarrow X_B^{t_A} \not\perp\!\!\!\perp X_A^{t_A} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]} \Rightarrow I(X_B^{t_A}; \mathbf{X}_A^{[t_A, T]} | \mathbf{X}_A^{[t_A-1]}, \mathbf{X}_C^{[T]}, \mathbf{X}_B^{[t_A-1]}) > 0$ .

The above derivation applies similarly for reviewer  $B$  when peers  $A, C$  complete all criteria report truthfully, and for reviewer  $C$  when peers  $A, B$  complete all criteria and report truthfully.  $\square$

**Theorem 2** Suppose mechanism hyperparameters  $\{\alpha_A^{[T]}, \beta_A^{[T]}, \alpha_B^{[T]}, \beta_B^{[T]}, \alpha_C^{[T]}, \beta_C^{[T]}\}$  are solutions to the following linear program:

$$\begin{aligned}
\min \quad & \sum_{k \in \{A, B, C\}} \sum_{t=1}^T \alpha_k^t R_{k^t}^{\text{expert}}(T) + \beta_k^t R_{k^t}^{\text{target}}(T) \\
\text{s.t.} \quad & \sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{\text{expert}}(t_A) + \beta_A^t \Delta R_{A^t}^{\text{target}}(t_A) > \Delta e_A(t_A) \text{ for } 1 \leq t_A \leq T \\
& \sum_{t=1}^T \alpha_B^t \Delta R_{B^t}^{\text{expert}}(t_B) + \beta_B^t \Delta R_{B^t}^{\text{target}}(t_B) > \Delta e_B(t_B) \text{ for } 1 \leq t_B \leq T \\
& \sum_{t=1}^T \alpha_C^t \Delta R_{C^t}^{\text{expert}}(t_C) + \beta_C^t \Delta R_{C^t}^{\text{target}}(t_C) > \Delta e_C(t_C) \text{ for } 1 \leq t_C \leq T \\
& \alpha_k^t, \beta_k^t \geq 0
\end{aligned} \tag{19}$$

Then, for any reviewer  $k$ , if their peers complete all criteria and report predictions and review scores truthfully, reviewer  $k$  strictly maximizes their expected utility by also completing all criteria and reporting predictions and review scores truthfully. In other words, completing all criteria and reporting truthfully is a strict Bayesian-Nash equilibrium. We refer to this equilibrium as the fully-informative equilibrium.

*Proof.* We require the hyperparameters to be non-negative so reviewers are paid and not penalized for maximizing their expert and target payments. When reviewer  $k$ 's peers are completing all criteria and reporting truthfully, we want the mechanism to pay enough to cover reviewer  $k$ 's cost of effort and incentivize them to complete all criteria and report truthfully as well. By Theorem 1, we know that truthful reporting on shared completed criteria is a strict BNE, so when a reviewer  $k$ 's peers complete all criteria and report truthfully, every additional criterion completed by reviewer  $k$  becomes a shared complete criterion and their payoff is maximized by truthfully reporting the review score on the marginal criterion. Thus, the question is simply how to set the hyperparameters to incentivize completion of all criteria (when peers complete all criteria and report truthfully). We analyse reviewer  $A$ 's incentives when reviewers  $B, C$  complete all criteria and report truthfully, and this applies symmetrically to reviewers  $B, C$ .

Reviewer  $A$ 's expected utility for completing  $t_A$  criteria and reporting signals truthfully is  $U_A = R_A(t_A) - e_A(t_A)$ , so to ensure that the mechanism always covers the cost of effort, we want the marginal utility of completing an additional criterion  $t_A$  to be positive, i.e.,  $R_A(t_A) - e_A(t_A) > R_A(t_A-1) - e_A(t_A-1)$  for all  $1 \leq t_A \leq T$ . Rearranging this, we can write  $R_A(t_A) - R_A(t_A-1) > e_A(t_A) - e_A(t_A-1) \Rightarrow \Delta R_A(t_A) > \Delta e_A(t_A)$ .  $\Delta e_A(t_A)$  is the marginal cost of completing criterion  $t_A$ , and by Equation 17, the marginal expected payment for completing criterion  $t_A$  when peers complete all criteria is  $\Delta R_A(t_A) = \sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{\text{expert}}(t_A) + \beta_A^t \Delta R_{A^t}^{\text{target}}(t_A)$ . Thus, if the hyperparameters  $\{\alpha_A^{[T]}, \beta_A^{[T]}\}$  satisfy  $\sum_{t=1}^T \alpha_A^t \Delta R_{A^t}^{\text{expert}}(t_A) + \beta_A^t \Delta R_{A^t}^{\text{target}}(t_A) > \Delta e_A(t_A)$  for all  $1 \leq t_A \leq T$ , then reviewer  $A$  maximizes their utility by completing all criteria and reporting truthfully. Satisfying these constraints is feasible since  $\Delta R_{A^t}^{\text{expert}}(t_A), \Delta R_{A^t}^{\text{target}}(t_A) > 0$  for  $t = t_A$  as shown in Lemma 1.

By the same argument, we desire  $\{\alpha_B^{[T]}, \beta_B^{[T]}\}$  satisfy  $\Delta R_B(t_B) > \Delta e_B(t_B)$ , and  $\{\alpha_C^{[T]}, \beta_C^{[T]}\}$  satisfy  $\Delta R_C(t_C) > \Delta e_C(t_C)$ . Then, for any reviewer  $k$ , if their peers complete all criteria and report truthfully, their marginal utility of completing an additional criterion is always positive, so reviewer  $k$  strictly maximizes their expected utility by also completing all criteria and reporting truthfully. Hence, completing all criteria and reporting truthfully is a strict Bayesian-Nash equilibrium.

In setting the hyperparameters, we also wish to minimize the mechanism's cost for having a paper reviewed, so we wish to minimize the total payout  $R_A(T) + R_B(T) + R_C(T) = \min \sum_{k \in \{A, B, C\}} \sum_{t=1}^T \alpha_k^t R_{k^t}^{\text{expert}}(T) + \beta_k^t R_{k^t}^{\text{target}}(T)$ . The optimal choice of hyperparameters is found by minimizing the mechanism's cost under the aforementioned constraints by solving the given linear program where the constraints are modified to a weak inequality, and adding an  $\epsilon$  to the solution of the LP to strictly satisfy the constraints.

□

**Lemma 2** Let  $s_k = (e_k(t_k), \hat{\mathbf{x}}_k^{[T]}, \hat{\mathbf{p}}_{k' \leftarrow k}^+, \hat{\mathbf{p}}_{k' \leftarrow k})$  be reviewer  $k$ 's strategy, where  $k'$  is reviewer  $k$ 's target. If  $\mathbf{s} = (s_A, s_B, s_C)$  is an equilibrium strategy profile, reviewer  $A$ 's expected payment is:

$$\mathbb{E}[r_A] = \sum_{t=1}^T \alpha_A^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right) \right] \right) + \beta_A^t I \left( \hat{X}_B^t; \hat{X}_A | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]} \right) \quad (20)$$

(expected payment for peers  $B, C$  computed analogously). Reviewers' total expected payment is:

$$\begin{aligned} \mathbb{E}[r_A + r_B + r_C] = & \sum_{t=1}^T \alpha_A^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right) \right] \right) + \beta_A^t I \left( \hat{X}_B^t; \hat{X}_A | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]} \right) \\ & + \alpha_B^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t-1]} \right) \right) \right] \right) + \beta_B^t I \left( \hat{X}_C^t; \hat{X}_B | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{X}}_C^{[t-1]} \right) \\ & + \alpha_C^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t-1]} \right) \right) \right] \right) + \beta_C^t I \left( \hat{X}_A^t; \hat{X}_C | \mathbf{X}_B^{[t_B]}, \hat{\mathbf{X}}_A^{[t-1]} \right) \end{aligned} \quad (21)$$

*Proof of Lemma 2.* First, we observe that since  $\mathbf{s}$  is an equilibrium strategy profile, all reviewers' predictions are truthful, i.e.,  $\hat{\mathbf{p}}_{B \leftarrow A}^+ = \mathbb{Q} \left( \hat{X}_B^t | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{X}}_C^{[t]} \right)$  and  $\hat{\mathbf{p}}_{B \leftarrow A}^- = \mathbb{Q} \left( \hat{X}_B^t | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{X}}_C^{[t-1]} \right)$ , and so on for the other reviewers' expert predictions. Predictions are truthful in any equilibrium as otherwise reviewers could unilaterally improve their expert payment under the strictly proper log scoring rule by deviating to a truthful prediction. The expected payment for reviewers (expectation of Equation 5) for criterion  $t$  is:

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[r_{A^t}] &= \mathbb{E}[\alpha_A^t (S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B \leftarrow A}^+) + S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B \leftarrow A}^-)) + \beta_A^t (S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A \leftarrow C}^+) - S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A \leftarrow C}^-))] \\ \mathbb{E}_{\mathbb{Q}}[r_{B^t}] &= \mathbb{E}[\alpha_A^t (S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C \leftarrow B}^+) + S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C \leftarrow B}^-)) + \beta_B^t (S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B \leftarrow A}^+) - S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B \leftarrow A}^-))] \\ \mathbb{E}_{\mathbb{Q}}[r_{C^t}] &= \mathbb{E}[\alpha_C^t (S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A \leftarrow C}^+) + S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A \leftarrow C}^-)) + \beta_C^t (S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C \leftarrow B}^+) - S_{\log}(\hat{x}_C^t, \hat{\mathbf{p}}_{C \leftarrow B}^-))] \end{aligned} \quad (22)$$

Now, observe that reviewer  $A$ 's expected expert payment (first two terms in expectation) on some criterion  $t$  can be written as follows:

$$\begin{aligned} \alpha_A^t \mathbb{E}[S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B \leftarrow A}^+) + S_{\log}(\hat{x}_B^t, \hat{\mathbf{p}}_{B \leftarrow A}^-)] &= \alpha_A^t \mathbb{E} \left[ \mathbb{Q} \left( \hat{X}_B^t | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) + \mathbb{Q} \left( \hat{X}_B^t | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right] \\ &= -\mathbb{E}_{\mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]}} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) \right) \right] - \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]}} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right) \right] \end{aligned} \quad (23)$$

where  $H(\cdot)$  is the Shannon entropy. For reviewer  $A$ 's target payment (last two terms in expectation), we can use Equation 10 to write:

$$\mathbb{E}_{\mathbb{Q}}[S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A \leftarrow C}^+) - S_{\log}(\hat{x}_A^t, \hat{\mathbf{p}}_{A \leftarrow C}^-)] = I \left( \hat{X}_B^t; \hat{X}_A | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]} \right) \quad (24)$$

**Remark 2.** Each reviewer's target payment on criterion  $t$  is exactly the HMIP payment for that criterion under the conditional independence assumption (Assumption 4.3 in Kong and Schoenebeck (2018)).

We can then write reviewer  $A$ 's total expected payment on criterion  $t$  as follows (and the same applies to  $B$  and  $C$ ):

$$\mathbb{E}_{\mathbb{Q}}[r_{A^t}] = \alpha_A^t \left( -\mathbb{E}_{\mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]}} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) \right) \right] - \mathbb{E}_{\mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]}} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right) \right] \right) + \beta_A^t I \left( \hat{X}_B^t; \hat{X}_A | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]} \right) \quad (25)$$

Making the analogous computation for reviewers  $B$  and  $C$ , we can compute the total expected payment for the reviewers on a single criterion  $t$ :

$$\begin{aligned} \mathbb{E}_{\mathbb{Q}}[r_{A^t} + r_{B^t} + r_{C^t}] = & \alpha_A^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right) \right] \right) + \beta_A^t I \left( \hat{X}_B^t; \hat{X}_A | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]} \right) \\ & + \alpha_B^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t-1]} \right) \right) \right] \right) + \beta_B^t I \left( \hat{X}_C^t; \hat{X}_B | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{X}}_C^{[t-1]} \right) \\ & + \alpha_C^t \left( -\mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( \mathbb{Q} \left( \hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t-1]} \right) \right) \right] \right) + \beta_C^t I \left( \hat{X}_A^t; \hat{X}_C | \mathbf{X}_B^{[t_B]}, \hat{\mathbf{X}}_A^{[t-1]} \right) \end{aligned} \quad (26)$$

Thus, the total expected payment for the reviewers for all criteria is:

$$\begin{aligned}
\mathbb{E}_Q[r_A + r_B + r_C] &= \mathbb{E} \left[ \sum_{k \in \{A, B, C\}} \sum_{t=1}^T r_{k^t} \right] \\
&= \sum_{t=1}^T \alpha_A^t \left( -\mathbb{E} \left[ H \left( Q \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( Q \left( \hat{X}_B^t | \mathbf{x}_A^{[t_A]}, \hat{\mathbf{x}}_C^{[t-1]} \right) \right) \right] \right) + \beta_A^t I \left( \hat{X}_B^t; \hat{X}_A^t | \mathbf{X}_C^{[t_C]}, \hat{\mathbf{X}}_B^{[t-1]} \right) \\
&\quad + \alpha_B^t \left( -\mathbb{E} \left[ H \left( Q \left( \hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( Q \left( \hat{X}_C^t | \mathbf{x}_B^{[t_B]}, \hat{\mathbf{x}}_A^{[t-1]} \right) \right) \right] \right) + \beta_B^t I \left( \hat{X}_C^t; \hat{X}_B^t | \mathbf{X}_A^{[t_A]}, \hat{\mathbf{X}}_C^{[t-1]} \right) \\
&\quad + \alpha_C^t \left( -\mathbb{E} \left[ H \left( Q \left( \hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t]} \right) \right) \right] - \mathbb{E} \left[ H \left( Q \left( \hat{X}_A^t | \mathbf{x}_C^{[t_C]}, \hat{\mathbf{x}}_B^{[t-1]} \right) \right) \right] \right) + \beta_C^t I \left( \hat{X}_A^t; \hat{X}_C^t | \mathbf{X}_B^{[t_B]}, \hat{\mathbf{X}}_A^{[t-1]} \right)
\end{aligned} \tag{27}$$

□

### Theorem 3

- (a) *In the highest paying uninformative equilibrium (where all reviewers exert zero effort), every reviewer's expected payment (and hence aggregate reviewer welfare) is zero.*
- (b) *Suppose that in addition to the constraints specified in Theorem 2, the mechanism hyperparameters also satisfy  $R_k(T) - e_k(T) > 0 \Leftrightarrow \sum_{t=1}^T \alpha_k^t R_{k^t}^{\text{expert}}(T) + \beta_k^t R_{k^t}^{\text{target}}(T) > e_k(T)$  for  $k \in \{A, B, C\}$ , i.e., reviewers' expected utility in the fully-informative equilibrium is positive. Then, it is individually rational to participate in the mechanism and complete all criteria and report truthfully, given peers are doing the same. Additionally, the fully-informative equilibrium has higher individual and aggregate utility than the uninformed equilibrium.*

*Proof of Theorem 3.*

- (a) When all reviewers exert zero effort, their reports are not sampled from the common prior. Instead, they may report some fixed a priori most likely sequence of signals. In such an equilibrium, expert predictions without the source will be perfect and the updated predictions will be the same, as the expert can simply predict the constant signals with complete certainty. Thus, the entropy of expert predictions is minimized and expert payments will be 0. Since expert predictions will not change, all target payments will also be 0 in expectation. Intuitively, this is because the signals do not come from the common prior, and do not contain information about other reviewers' signals.

In any other uninformative equilibrium, expected target payments will still be zero since mutual information of uninformed signals not drawn from the common prior will be zero. However, the expert payment can only decrease (the maximum possible expert payment is zero). Thus, in the highest paying uninformative equilibrium, reviewers individually and in aggregate are paid nothing in expectation.

- (b) It is individually rational for a reviewer to participate in the mechanism when their expected utility is positive. We want individual rationality under our desired fully informative equilibria, so we achieve positive expected utility by requiring  $R_k(T) - e_k(T) > 0$  or equivalently  $\sum_{t=1}^T \alpha_k^t R_{k^t}^{\text{expert}}(T) + \beta_k^t R_{k^t}^{\text{target}}(T) > e_k(T)$ .

Denoting the expected utility of some arbitrary uninformative equilibrium as  $R_k(0, 0, 0)$ , we observe that  $R_k(T) > e_k(T) > 0 \geq R_k(0, 0, 0)$ , thus the individual utility of the fully-informative equilibrium is higher than any uninformative equilibrium. This holds for all reviewers, so the aggregate utility of the fully-informative equilibrium is also higher than the uninformative equilibrium.

□