

Delegating performance evaluation

IGOR LETINA

Department of Economics, University of Bern

SHUO LIU

Guanghua School of Management, Peking University

NICK NETZER

Department of Economics, University of Zurich

We study optimal incentive contracts with multiple agents when performance is evaluated by a reviewer. The reviewer may be biased in favor of the agents, but the degree of bias is unknown to the principal. We show that a contest, which is a contract in which the principal fixes a set of prizes to be allocated to the agents, is optimal. By using a contest, the principal can commit to sustaining incentives despite the reviewer's potential leniency bias. The optimal effort profile can be uniquely implemented by an all-pay auction with a cap. Our analysis has implications for various applications, such as the design of worker compensation or the allocation of research grants.

KEYWORDS. Subjective performance evaluation, mechanism design without money, optimal delegation, optimality of contests.

JEL CLASSIFICATION. D02, D82, M52.

Igor Letina: igor.letina@vwi.unibe.ch

Shuo Liu: shuo.liu@gsm.pku.edu.cn

Nick Netzer: nick.netzer@econ.uzh.ch

We are grateful for useful comments and suggestions to the three anonymous referees, as well as Jean-Michel Benkert, Elchanan Ben-Porath, Lorenzo Casaburi, Yeon-Koo Che, David Dorn, Christian Ewerhart, Alexander Frankel, William Fuchs, Sergiu Hart, Andreas Hefti, Johannes Hörner, Navin Kartik, Christian Kellner, Botond Kőszegi, Wladislaw Mill, Harry Pei, Ron Siegel, Florian Scheuer, Armin Schmutzler, Philipp Strack, Dezső Szalay, and seminar participants at Bar Ilan University, Bielefeld University, BI Oslo, Carnegie Mellon University, Cheung Kong GSB, Columbia University, Hebrew University of Jerusalem, Hong Kong University of Science and Technology, LUISS University Rome, McMaster University, National University of Singapore, Northwestern University, NYU Abu Dhabi, Peking University, ShanghaiTech University, Sun Yat-sen University, Toulouse School of Economics, the Universities of Bern, Bonn, Cologne, Konstanz, St. Gallen, and Zurich, 18th Annual SAET Conference, 20th Colloquium on Personnel Economics, 28th International Conference on Game Theory, 2018 Asian Meeting of the Econometric Society, 2018 China Meeting of the Econometric Society, Organizational Economics Workshop Konstanz 2018, SIOE Annual Conference 2017, SIRE Workshop on Behavioural Economics and Mechanism Design 2016, Swiss Economists Abroad Meeting 2016, Swiss Theory Day 2016, Verein für Socialpolitik Meeting 2016, Verein für Socialpolitik Theoretischer Ausschuss 2017, Zurich Workshop in Economics 2016, and Zurich Workshop on Applied Behavioral Theory 2017. Shuo Liu would like to acknowledge the hospitality of Columbia University, where some of this work was carried out, and financial support by the Swiss National Science Foundation (Doc.Mobility grant P1ZHP1_168260).

1. INTRODUCTION

Principals often lack the information or expertise needed to evaluate the performance of their agents. A common response to this problem is to rely on a better informed party for performance evaluation. For example, funding agencies ask a committee of referees to assess which research projects should be funded. Within a firm, the CEO usually relies on a mid-level supervisor to judge which of the subordinates should get a bonus.

If the preferences of the principal and the evaluator, henceforth called the reviewer, are not aligned, this can lead to distorted evaluations. In particular, there is much evidence that supervisors tend to be too lenient when judging the performance of their subordinates.¹ Anticipating leniency, the subordinates will exert less effort. We study optimal contracts that take into account the reviewer's incentives when reporting the observed performance of the agents.²

We consider a particular way in which the reviewer's preferences are not aligned with the principal's. While the reviewer takes into account the effect of his actions on the principal's payoff—maybe because he owns shares of the company, has career concerns, or cares intrinsically—he also cares about the agents (with equal weight for each agent). A reviewer who cares about the agents will be reluctant to hurt them even if they do not exert sufficient effort. He may thus exhibit a leniency bias. In our model, exactly how much the reviewer cares about the agents is private information of the reviewer. This is again inspired by the empirical literature, which shows that the degree of leniency varies and appears to depend on the supervisor's personality and social ties between the supervisor and the team.³

One could imagine that the principal tries to correct the leniency bias by paying transfers to the reviewer conditional on the evaluation that he performs. Our paper belongs to a literature on mechanism design without money where conditional monetary transfers (to the reviewer) are ruled out by assumption. The approach restricts the class of mechanisms over which the principal optimizes, but we believe that the restriction is very plausible in our setting of performance evaluation. A conditional transfer would mean paying the reviewer more when he reports a more negative evaluation. This is typically considered a conflict of interest and is either explicitly illegal or judged as repugnant.⁴

¹Bol (2011) cites studies documenting leniency bias going back to the 1920's, while citations to similar findings in the 1940's can be found in Prendergast (1999). Berger et al. (2013) provide experimental evidence of the leniency bias.

²We do not consider the problem of incentivizing the reviewer to exert costly effort in order to learn some state of nature. This is an interesting but distinct incentive problem which is studied in Aghion and Tirole (1997), Strausz (1997), Szalay (2005), Rahman (2012), Pei (2015), and Liu and Migrow (2019). We also do not examine why the principal does not perform the task of the reviewer herself. We treat the existence of the reviewer as given in the same way that we treat the existence of the agents as given. Furthermore, Strausz (1997) shows that it can be optimal to delegate evaluation to a reviewer even when the principal could monitor an agent herself.

³Bernardin et al. (2000) document that the degree of leniency bias in an experiment depends on personality traits of the reviewer. Bol (2011) and Breuer et al. (2013) find evidence of leniency bias which depends on the strength of the employee-manager relationship.

⁴In Prendergast and Topel (1996), both the principal and a reviewer receive a signal about a worker's effort. Their main result is that leniency need not be costly for the principal if she can charge the reviewer

Our main result is that a contest among the agents is an optimal mechanism. That is, the principal defines a set of prizes and the reviewer's report only determines how these prizes are allocated to the agents. The reviewer cannot influence the overall size or the division of the agents' compensation. This strongly limits the degree of leniency he can exercise. In particular, some agents are always punished by being assigned a small prize, which is crucial for the preservation of incentives. The downside of the contest mechanism is that a small prize has to be assigned (at random) even when all agents provide the sufficient level of effort. When the agents are risk-averse, this will be inefficient.

More formally, we solve a mechanism design problem where the principal commits to a mapping from the reviewer's report (about his type and the agents' efforts) to the payments of the agents. To understand why a contest is optimal, consider how these payments can vary with the reviewer's report. First, when there is uncertainty about the reviewer's type, the only incentive-compatible mechanisms are those where the total sum of payments is independent of the reported effort profile, for each type of the reviewer. If this was not the case, a reviewer who cares sufficiently much about the agents would distort his report to increase the payments to the agents, while a reviewer who cares sufficiently much about the principal would distort his report to decrease these payments. Then, since the sum of payments is independent of the reported effort profile, varying this sum across different reviewer types also cannot be strictly profitable for the principal.⁵ Consequently, the optimal mechanism must provide a "commitment to pay" a fixed total amount of money to the agents.⁶ Second, the division of the fixed amount of money between the agents can vary with the reviewer's report only to a limited extent. If given the opportunity, a reviewer who cares about the agents would always prefer to distribute the payments equally among all the agents instead of inducing effort-contingent payments. Thus the optimal mechanism must additionally provide a "commitment to punish" shirking agents.⁷ A contest, which is a mechanism in which

for exercising leniency. In [Giebe and Gürtler \(2012\)](#), the principal offers a menu of contracts to screen lenient and nonlenient reviewers. They show that if the nonlenient type is common enough, the optimal solution can be to pay a flat wage to the reviewer, and rely on the nonlenient type for punishment of an agent who shirks. The main difference between these papers and ours is that we consider multiple agents and do not allow contracts which condition payment to the reviewer on the reported evaluation.

⁵[Krähmer and Kováč \(2016\)](#) and [Tanner \(2013\)](#) obtain related no-screening results for delegation.

⁶The literature has observed previously that a commitment to pay can be important when the agents' efforts are not verifiable. For instance, [Malcomson \(1984, 1986\)](#) argues that piece-rate contracts are not credible because the principal would always claim low performance *ex post* in order to reduce payments. [Levin \(2002\)](#) shows that a constant sum of payments arises in equilibrium of a repeated game, where the principal's wage promises have to be self-enforcing. A contest-type compensation structure arises in an optimal equilibrium only if the agents are risk-neutral. [Svensson \(2003\)](#) notes that conditional foreign aid, where countries are rewarded with aid for successful reforms, can be ineffective if the aid officer is lenient. Instead of evaluating each country separately, he proposes that each aid officer should be given a fixed budget and discretion in how to allocate the aid across a group of countries.

⁷The problem of committing to punishment is related to [Konrad \(2001\)](#) and [Netzer and Scheuer \(2010\)](#). They study the problem of a planner who would like to implement redistribution after agents have chosen their actions, the anticipation of which may destroy incentives to choose costly but socially desirable actions. In the context of optimal income taxation, [Konrad \(2001\)](#) shows that private information about labor

the payments to the agents are always just permutations of a fixed profile of prizes, provides these two types of commitment and enables the principal to incentivize both the reviewer and the agents in an optimal way.

Our second result characterizes the prize structure of an optimal contest. Given n agents, an optimal contest will have $n - 1$ equal positive prizes and one zero prize. Thus, while the contest acts as a commitment to punish, the punishment is kept at the minimum required to incentivize effort. The mechanism allocates a zero prize to only one agent, so that the optimal contest exhibits a “loser-takes-nothing” rather than a “winner-takes-all” structure. In equilibrium, when all agents have provided sufficient effort, the agent who receives the zero prize is determined randomly. Thus all agents are facing the same risk. If agents are risk-averse, they respond to this risk by reducing the amount of effort they are willing to exert. A corollary of this result is that the first-best is implementable if and only if the agents are risk-neutral.

Our third result shows that a familiar all-pay auction with a cap achieves the optimum, and it does so in unique equilibrium. As in a standard all-pay auction with $n - 1$ identical prizes, the agent with the lowest effort receives the zero prize. Ties are broken randomly. However, efforts are capped at the desired equilibrium level. This removes the possibility for the agents to exert slightly more than the equilibrium effort in order to guarantee themselves a positive prize with probability one. In addition to the all-pay auction, it can be shown that the optimum can also be achieved with an imperfectly discriminating contest, such as the well-known Tullock contest. Thus, the essential feature of our main result is the fixed profile of prizes, and not the exact form of the contest success function.

Our results are interesting for at least three reasons. First, contests are a commonly used and often-studied incentive scheme, and much work has been devoted to their optimal design.⁸ There is, however, not much work on the more general question whether and under which conditions contests are optimal in some class of mechanisms. In their seminal paper, Lazear and Rosen (1981) compare piece-rate contracts and contests. When agents are risk-neutral, both mechanisms can induce the first-best outcome. For risk-averse agents, they show that either mechanism sometimes dominates the other, but they do not establish results on unconstrained optimality. A defining feature of contests is that the payoff of the agents depends on how well they perform relative to each other. Green and Stokey (1983) and Nalebuff and Stiglitz (1983) use this fact to show that contests can do better than individual contracts when agents are risk-averse and there is a random common shock to their outputs. If the relationship between effort and output is ambiguous and the agents are ambiguity-averse, Kellner (2015) shows that contests can be optimal because they filter out the ambiguity. In our paper, contests provide a

productivity provides a commitment against excessive redistribution. In the context of insurance and labor markets, Netzer and Scheuer (2010) show that adverse selection provides a commitment by generating separating market equilibria.

⁸For instance, Glazer and Hassin (1988) and Moldovanu and Sela (2001) study the design of prizes for a given contest success function. Jia et al. (2013) survey papers that study the design of a contest success function for given prizes.

commitment for lenient reviewers to punish shirking agents. Contests are optimal because they incentivize the agents at least as efficiently as any other contract that can provide such a commitment.

Second, the assumed absence of conditional monetary transfers to the reviewer relates our approach to the literature on optimal delegation. In our model, the reviewer's report determines the agents' compensation, and thus one can think of the reviewer as choosing the agents' compensation directly from a limited set of options that the principal makes available. The literature on optimal delegation studies how such delegation sets should be designed.⁹ In the existing delegation literature, the principal wants to base her decision on some stochastic state of nature which is assumed to be exogenous and observed by an expert. In our setting, the state of nature on which the expert has private information (the performance of the agents) is endogenous and determined in anticipation of the report that the expert will make. As a matter of fact, the goal of the principal is exactly to incentivize the agents to exert effort. Thus we contribute to the literature on optimal delegation by introducing an endogenous state of nature that is affected by the mechanism. The most closely related paper from the delegation literature is Frankel (2014). Like in our paper, he considers a multidimensional problem with uncertainty about the expert's preferences. In contrast to our paper, he assumes that the state of nature is exogenous. Another difference is that Frankel (2014) derives max–min mechanisms, which are optimal for the worst possible realization of the expert's bias, while we are interested in mechanisms that maximize the principal's expected payoff given her beliefs about the expert's type. Frankel (2014) shows that, when the set of possible preferences of the expert is rich enough, a ranking mechanism is max–min optimal. When the multidimensional state of nature is an effort profile, like in our setting, this ranking mechanism would correspond to a standard all-pay auction.¹⁰ Such an all-pay auction without a cap is not optimal in our setting with endogenous efforts. However, there is a range of different contests that are optimal. Furthermore, since the principal's payoff turns out to be independent of the reviewer's type in these optimal contests, they not only maximize expected payoffs but are also max–min optimal.

Third, our results have direct implications for the controversial debate over the use of so-called “forced rankings,” a review system which was most famously used by General Electric under Jack Welch during their fast growth in the 1980's and 1990's.¹¹ Our contribution to this discussion is (i) to show that forced rankings are optimal for motivating effort under the assumptions of our model, (ii) to show how optimal forced rankings should be constructed, and (iii) to show that some elements of forced rankings which

⁹See, for example, Holmström (1977, 1984), Melumad and Shibano (1991), Alonso and Matouschek (2008), Armstrong and Vickers (2010), Amador and Bagwell (2013), and Frankel (2014).

¹⁰See Charkraborty and Harbaugh (2007) for a model of multidimensional cheap talk in which ranking statements are credible in equilibrium, and Siegel (2009) for a general treatment of contests with all-pay structure.

¹¹For example, see “Rank and Yank’ Retains Vocal Fans” (<https://www.wsj.com/articles/SB10001424052970203363504577186970064375222>), L. Kwok, *The Wall Street Journal*, January 31, 2012, and “For Whom the Bell Curve Tolls” (https://www.washingtonpost.com/news/on-leadership/wp/2013/11/20/for-whom-the-bell-curve-tolls/?utm_term=.c09d9ad39dc8), J. McGregor, *The Washington Post*, November 20, 2013.

are usually criticized are actually necessary for incentivizing effort. In particular, forced rankings are criticized for forcing managers to assign low rankings even when all workers are performing well: “What happens if you’re working with a superstar team? You’ve just forced a distribution that doesn’t exist. You create this stupid world where [great] people are punished.”¹² Our results show that, far from being stupid, not rewarding some workers even when they perform well is necessary, since if the managers were given an opportunity to reward all workers, they may use it irrespective of actual performance, which would destroy the incentive effect of the evaluation system.

We make several simplifying assumptions in our model. The most important ones are that the reviewer perfectly observes the efforts of the agents, that the reviewer has no intrinsic motive to reward higher efforts and also does not exhibit favoritism toward any agent, and that all agents have identical cost functions. In the conclusion, we will discuss how these assumptions can be relaxed and to what extent our results are robust. On a more general note, while our paper identifies a novel benefit of relative performance evaluation, it abstracts away from some of its costs, namely collusion among the agents and sabotage. While these costs of relative performance evaluation are not the focus of the paper, we will return to them in the conclusion where a more substantive discussion is provided.

2. THE MODEL

2.1 Environment

There is a set of agents $I = \{1, \dots, n\}$ with $n \geq 2$. Each agent $i \in I$ chooses an effort level $e_i \geq 0$ and obtains a monetary transfer $t_i \geq 0$. The agents have an outside option of zero. The payoff of agent i is given by

$$\pi_i(e_i, t_i) = u(t_i) - c(e_i).$$

The utility function $u : \mathbb{R}_+ \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing, weakly concave, and satisfies $u(0) = 0$. The cost function $c : \mathbb{R}_+ \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing, strictly convex, and satisfies $c(0) = 0$, $c'(0) = 0$, and $\lim_{e_i \rightarrow \infty} c'(e_i) = \infty$. The assumption of additive separability of transfers and efforts is standard in contract theory, mechanism design, and contest theory. We will discuss the robustness of our results with respect to nonseparable preferences in the conclusion.

Denote effort profiles by $e = (e_1, \dots, e_n) \in E$ and transfer profiles by $t = (t_1, \dots, t_n) \in T$. We assume that $E = \mathbb{R}_+^n$ and $T = \{t \in \mathbb{R}_+^n \mid \sum_{i=1}^n t_i \leq \bar{T}\}$, where $\bar{T} > 0$ can be arbitrarily large. Our results hold no matter whether the budget \bar{T} is eventually binding or not.

There is a *principal*, whose payoff from an allocation (e, t) is

$$\pi_P(e, t) = z(e) - \sum_{i=1}^n t_i,$$

¹²Quote of a management adviser in “For Whom the Bell Curve Tolls” (https://www.washingtonpost.com/news/on-leadership/wp/2013/11/20/for-whom-the-bell-curve-tolls/?utm_term=.c09d9ad39dc8), J. McGregor, *The Washington Post*, November 20, 2013.

where $z : E \rightarrow \mathbb{R}_+$ is interpreted as the production function that converts efforts into output. For clarity of exposition, we will focus only on the case where $z(e) = \sum_{i=1}^n e_i$. Our main results continue to hold if we assume more generally that z is symmetric, weakly concave, and strictly increasing in each of its arguments.

EXAMPLE. We will use a parameterized example to illustrate our results throughout the paper. In this example, each of the n agents has the payoff function

$$\pi_i(e_i, t_i) = t_i^\alpha - \gamma e_i^\beta,$$

where $0 < \alpha \leq 1$ parameterizes risk-aversion, $\beta > 1$ describes the degree of cost convexity, and $\gamma > 0$ determines the relative weight of effort costs.¹³ We will always assume that \bar{T} is large enough to be nonbinding in the example. The first-best effort level e^{FB} is what the principal would demand from each agent if she could perfectly control effort and would only have to compensate the agent for his cost, thus paying $t^{\text{FB}} = u^{-1}(c(e^{\text{FB}}))$. In our example, maximization of $e - u^{-1}(c(e))$ yields

$$e^{\text{FB}} = \left(\frac{\alpha}{\beta \gamma^{1/\alpha}} \right)^{\frac{\alpha}{\beta-\alpha}} \quad \text{and} \quad t^{\text{FB}} = \left(\frac{\alpha}{\beta \gamma^{1/\beta}} \right)^{\frac{\beta}{\beta-\alpha}}.$$

The principal's first-best profit is $n(e^{\text{FB}} - t^{\text{FB}})$.

The effort profile of the agents is neither verifiable to outside parties (e.g., a court) nor observable to the principal. The efforts can, however, be observed by a *reviewer*. The payoff of the reviewer from an allocation (e, t) is given by

$$\pi_R(e, t, \theta) = \pi_P(e, t) + \theta \sum_{i=1}^n \pi_i(e, t),$$

where θ is a parameter that captures how much the reviewer cares about the well-being of the agents, and thus by how much the reviewer's preferences are misaligned with those of the principal. The parameter θ can be thought of as a fundamental preference or as a reduced-form representation of concerns due to other interactions with the agents. We assume that θ is private information of the reviewer, observable neither to the principal nor to the agents. It is drawn according to a commonly known continuous distribution with full support on $\Theta = [\underline{\theta}, \bar{\theta}]$, where $\underline{\theta} < \bar{\theta}$. We describe this distribution by an (absolutely continuous) probability measure τ over Θ . Our results will be independent of the shape of this distribution. In particular, τ could be arbitrarily close to a probability measure with atoms. When discussing the intuition behind our results, such as the issues arising due to a leniency bias, we implicitly assume that θ is positive, but our results also hold when θ can take negative values.

¹³Since we assume $\alpha > 0$, the example includes some but not all CRRA utility functions. This is because we require $u(0) = 0$, which is not satisfied by all CRRA functions.

2.2 Direct mechanisms

We solve a mechanism design problem without monetary transfers from the principal to the reviewer.¹⁴ The timing is as follows. First, the principal designs a direct mechanism (which is w.l.o.g. due to the revelation principle). Second, the agents choose their efforts simultaneously. Third, the reviewer privately observes these efforts and his type and makes a report.¹⁵

Formally, the principal commits to a direct mechanism $\Phi = (\mu^{e,\theta})_{(e,\theta) \in E \times \Theta}$, which maps each possible report $(e, \theta) \in E \times \Theta$ of the reviewer into a probability measure $\mu^{e,\theta} \in \Delta T$ on the set of transfers to the agents. We will simply refer to such a mechanism as a *contract*.¹⁶ Since the reviewer is always free to make any report that he prefers, the following incentive-compatibility constraint makes sure that the reviewer has an incentive to report truthfully:

$$\Pi_R(e, \mu^{e,\theta}, \theta) \geq \Pi_R(e, \mu^{e',\theta'}, \theta) \quad \forall (e, \theta), (e', \theta') \in E \times \Theta, \quad (\text{IC-R})$$

where

$$\Pi_R(e, \mu^{e',\theta'}, \theta) = \mathbb{E}_{\mu^{e',\theta'}} \left[\pi_P(e, t) + \theta \sum_{i=1}^n \pi_i(e_i, t_i) \right].$$

We say that a contract Φ is *credible* if it satisfies (IC-R).

Given a credible contract, the agents choose their efforts simultaneously, anticipating that they will be rewarded according to the contractual specifications. Denote by $\sigma_i \in \Delta \mathbb{R}_+$ agent i 's mixed strategy for his effort provision. We also write $e_i \in \Delta \mathbb{R}_+$ for Dirac measures that represent pure strategies. Strategy profiles are given by $\sigma = (\sigma_1, \dots, \sigma_n) \in (\Delta \mathbb{R}_+)^n$. We also use σ to denote the induced product measure in ΔE . We say that a contract Φ *implements* a strategy profile σ if it is credible and satisfies

$$\Pi_i((\sigma_i, \sigma_{-i}), \Phi) \geq \Pi_i((\sigma'_i, \sigma_{-i}), \Phi) \quad \forall \sigma'_i \in \Delta \mathbb{R}_+, \forall i \in I, \quad (\text{IC-A})$$

¹⁴To be precise, we assume that the principal does not pay the reviewer conditional on the report that the reviewer makes. The principal could pay a fixed fee. In addition, our model can be easily reformulated to include payments to the reviewer based on the profitability of the firm, such as when the reviewer owns shares or has career concerns that align his interests with the performance of the firm. Any such term in the reviewer's payoff function can be captured by π_P and a normalization of θ . Another instrument that our approach rules out is the possibility of direct communication between the principal and the agents.

¹⁵A different timing would be that the reviewer observes his type already before he interacts with the agents and observes their efforts. In that case, the principal could design a menu of mechanisms from which the reviewer selects before observing the agents' efforts. Our results are robust to this different timing provided that the reviewer's selection is not observable to the agents. If it was observable, additional signalling issues would arise, because the agents may coordinate on different effort profiles contingent on the reviewer's observable action. We refrain from studying these issues here, but we remark that, even in that case, the additional gain (if any) from a more complicated mechanism is limited if the agents' risk-aversion is moderate or the number of agents is large (see Section 3.2).

¹⁶We impose the following regularity condition on contracts: For each measurable set $A \subseteq T$, $\mu^{e,\theta}(A)$ is a measurable function of (e, θ) . This ensures that expected payoffs are well-defined.

where

$$\Pi_i(\sigma, \Phi) = \mathbb{E}_\sigma[\mathbb{E}_\tau[\mathbb{E}_{\mu^{e,\theta}}[u(t_i)]]] - \mathbb{E}_{\sigma_i}[c(e_i)].$$

Since a deviation to an effort of zero always guarantees each agent a payoff of at least zero, the agents' participation constraints can henceforth be ignored.

The principal maximizes her expected payoff by choosing a contract Φ to implement some strategy profile σ . Formally, the principal's problem is given by

$$\max_{(\sigma, \Phi)} \Pi_P(\sigma, \Phi) \quad \text{s.t.} \quad (\text{IC-R}), (\text{IC-A}), \tag{P}$$

where

$$\Pi_P(\sigma, \Phi) = \mathbb{E}_\sigma \left[\sum_{i=1}^n e_i \right] - \mathbb{E}_\sigma \left[\mathbb{E}_\tau \left[\mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n t_i \right] \right] \right].$$

A contract Φ^* is *optimal* if there exists σ^* such that (σ^*, Φ^*) solves (P).

Finally, we introduce a specific class of direct mechanisms that will be referred to as *contests*. A contest is described by a fixed profile of prizes $y = (y_1, \dots, y_n) \in T$, some of which could be zero, and a rule how these prizes are allocated to the n agents as a function of their efforts. More formally, let $P(y)$ denote the set of permutations of y .¹⁷ Then a contest C_y with prize profile y is a contract that satisfies (i) $\mu^{e,\theta}(P(y)) = 1$ for all $(e, \theta) \in E \times \Theta$, and (ii) $\mu^{e,\theta} = \mu^{e,\theta'}$ for all $\theta, \theta' \in \Theta$ and $e \in E$. Condition (i) states that, for each possible report of the reviewer, the realized transfer profile will always be a permutation of the same prize profile. Condition (ii) requires that the prizes are allocated according to a (possibly probabilistic) contest success function that depends on the agents' efforts but not on the type of the reviewer. This formulation includes all common contests, such as all-pay auctions and Tullock contests.

Note that every contest is credible. This is because, once the agents' efforts are sunk, any permutation of the prizes generates the same payoff for the principal and the same sum of utilities for the agents. Formally, in any contest C_y we obtain that

$$\Pi_R(e, \mu^{e',\theta'}, \theta) = \sum_{i=1}^n e_i - \sum_{i=1}^n y_i + \theta \left(\sum_{i=1}^n u(y_i) - \sum_{i=1}^n c(e_i) \right),$$

which is independent of (e', θ') . However, the set of credible contracts is substantially larger than the set of contests. For instance, it is possible to select from a much larger set of transfer profiles, not just permutations of a given prize profile, and still keep both the expected sum of transfers and the expected sum of the agents' utilities constant. Furthermore, credible contracts do not have to be independent of θ but can screen different reviewer types.

¹⁷Profile t is a permutation of y if there exists a bijective mapping $s : I \rightarrow I$ such that $t_i = y_{s(i)} \forall i \in I$.

3. OPTIMAL CONTRACTS

3.1 *The optimality of contests*

To illustrate the key incentive problem in our model, suppose first that the preference parameter θ was known to the principal and the agents. The following example shows that, in this case, there may exist a credible contract which is not a contest but which implements the first-best effort levels and extracts the entire surplus.

EXAMPLE. Consider our previous example for the special case of $n = 2$. Suppose the reviewer's type was common knowledge. First, assume $\theta = 0$ so that there is also no misalignment of preferences between the principal and the reviewer. Consider a contract Φ^{FB} where, if both agents exert e^{FB} , each of them is paid t^{FB} . If one agent deviates, that agent is paid 0 while the nondeviating agent is paid $2t^{\text{FB}}$. In case both agents deviate, they are again each paid t^{FB} . It is easy to verify that this contract is credible, because the sum of transfers is constant across $(t^{\text{FB}}, t^{\text{FB}})$, $(2t^{\text{FB}}, 0)$, and $(0, 2t^{\text{FB}})$, which makes the reviewer indifferent between these transfer profiles. It is also easy to verify that this contract implements $(e^{\text{FB}}, e^{\text{FB}})$, because both agents receive a payoff of zero in equilibrium and a payoff of at most zero after any unilateral deviation. Thus the first-best is achievable. Observe that Φ^{FB} is not a contest, because the three transfer profiles are not all permutations of each other. We will show in Section 3.2 that the first-best is not achievable by a contest if the agents are risk-averse. Hence, Φ^{FB} performs strictly better than any contest in this example. This shows that nonverifiability of effort alone does not make contests optimal.¹⁸

Now assume that $\theta > 0$ and adjust the contract Φ^{FB} as follows. The payment $2t^{\text{FB}}$ to a nondeviating agent is replaced by some t^{nd} , while everything else is kept unchanged. If t^{nd} is chosen such that the reviewer is indifferent between the transfer profiles $(t^{\text{FB}}, t^{\text{FB}})$, $(t^{\text{nd}}, 0)$ and $(0, t^{\text{nd}})$, credibility is restored and the first-best can be implemented. For instance, with $\alpha = 1/2$, $\beta = 2$, and $\gamma = 1$ we have $e^{\text{FB}} \approx 0.63$ and $t^{\text{FB}} \approx 0.16$. For a reviewer of known type $\theta = 3$, we would then obtain $t^{\text{nd}} \approx 1.15$.¹⁹ This shows that the misalignment of preferences per se does not make contests optimal either.

The contracts described in the example no longer work if θ is the reviewer's private information. Just consider the optimal contract for type $\theta = 0$. Any reviewer with type $\theta' > 0$ will strictly prefer to report in a way that induces the transfers $(t^{\text{FB}}, t^{\text{FB}})$, no matter what efforts the agents have exerted. This is the leniency bias. The crucial point is that *only* the reviewer of type $\theta = 0$ will follow the contract. All other reviewer types, no matter how close they are to type $\theta = 0$, will not take the efforts of the agents into

¹⁸This argument is similar to the optimality of bonus pools in Rajan and Reichelstein (2006). It is also related to MacLeod (2003), who considers transfers to a third party in order to keep the principal's expenditure constant. Fuchs (2015) shows that there is an additional commitment advantage of bonus payment schemes—they allow the manager to credibly reveal the fit between the worker and the firm.

¹⁹The indifference condition is $-2t^{\text{FB}} + \theta 2u(t^{\text{FB}}) = -t^{\text{nd}} + \theta u(t^{\text{nd}})$. Given our parameters, it has a second solution $t^{\text{nd}} \approx 3.72$, which would work as well. Note that the indifference condition is not guaranteed to have a solution for all parameter values, so our simple construction of a first-best contract does not work for all values of $\theta > 0$.

account when deciding on their optimal report. Thus, the contracts outlined in the example above are knife-edge cases. They fail as soon as the reviewer’s type is drawn from an arbitrary continuous distribution. This creates a stark contrast between the case of perfect information and the case of arbitrarily small uncertainty about the reviewer’s type.

Our first main result shows that, despite the fact that the set of possible contracts is very large, optimal contracts with uncertainty about θ take a very simple form.

THEOREM 1. *The set of optimal contracts contains a contest.*

We prove [Theorem 1](#) by a series of six lemmas. Since we have shown that the principal may be able to implement the first-best if she knew the reviewer’s private type θ , it would seem reasonable to expect that the principal could benefit from screening these types. In fact, it is not difficult to construct contracts that screen the reviewer’s type, by varying the sum of transfers and the sum of utilities given to the agents. However, [Lemmas 1–3](#) below demonstrate that it is not possible for the principal to benefit from screening. [Lemmas 4 and 5](#) then show that the principal cannot benefit from implementing mixed or asymmetric effort profiles. Finally, [Lemma 6](#) shows that using contests is without loss of optimality.

To begin with, we fix an arbitrary contract $\Phi = (\mu^{e,\theta})_{(e,\theta) \in E \times \Theta}$ and denote

$$S_t(e, \theta) = \mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n t_i \right], \quad S_u(e, \theta) = \mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n u(t_i) \right]$$

and $S(e, \theta) = -S_t(e, \theta) + \theta S_u(e, \theta)$. We can then write the credibility constraint (IC-R) as

$$-S_t(e, \theta) + \theta S_u(e, \theta) \geq -S_t(e', \theta') + \theta S_u(e', \theta') \quad \forall (e, \theta), (e', \theta') \in E \times \Theta.$$

Our first lemma provides a characterization of this multidimensional constraint.

LEMMA 1. *A contract Φ is credible if and only if the conditions (i)–(iii) hold:*

- (i) $\forall \theta \in \Theta, S(e, \theta) = S(e', \theta) \forall e, e' \in E.$
- (ii) $\forall e \in E, S_u(e, \theta)$ is nondecreasing in $\theta.$
- (iii) $\forall e \in E, S(e, \theta) = S(e, \underline{\theta}) + \int_{\underline{\theta}}^{\theta} S_u(e, s) ds \forall \theta \in \Theta.$

Conditions (ii) and (iii) are familiar from the mechanism design literature. They have to hold separately for each fixed effort profile e . Condition (i) concerns the effort dimension and shows that the payoff of any reviewer has to be constant for any reported e .

Observe that combining (i) and (iii) implies $\int_{\underline{\theta}}^{\theta} S_u(e, s) ds = \int_{\underline{\theta}}^{\theta} S_u(e', s) ds$ for all e, e' and θ . This means that not only $S(e, \theta)$ but also its constituent $S_u(e, \theta)$ has to be constant across different effort profiles e (for almost every θ). But then the other constituent $S_t(e, \theta)$ cannot vary with e either. The next lemma summarizes this important implication of the credibility constraint.

LEMMA 2. *A contract Φ is credible only if there exists a pair of functions $x : \Theta \rightarrow \mathbb{R}_+$ and $\hat{x} : \Theta \rightarrow \mathbb{R}_+$ such that, $\forall e \in E$,*

$$S_l(e, \theta) = x(\theta), \quad S_u(e, \theta) = \hat{x}(\theta)$$

for almost all $\theta \in \Theta$.

We now show that there is no gain for the principal from screening the reviewer's private type by using a complex contract where $\mu^{e,\theta}$ varies with θ .

LEMMA 3. *For every contract Φ that implements a strategy profile σ , there exists a contract $\hat{\Phi}$ that also implements σ , yields the same expected payoff to the principal, and, $\forall \theta, \theta' \in \Theta$, satisfies $\hat{\mu}^{e,\theta} = \hat{\mu}^{e,\theta'} \forall e \in E$.*

The proof of the lemma is constructive and shows how the contract $\hat{\Phi}$ without screening can be obtained from an arbitrary credible contract Φ . Additional intuition can be obtained by looking again at the example from the beginning of this section. There the lenient reviewer of type $\theta = 3$ is compensated for revealing a shirking agent by a larger sum of transfers that he can allocate to the nonshirking agents off the equilibrium path. However, by Lemma 2, the sum of transfers $S_l(e, \theta)$ cannot vary with the reported effort profile e , so the principal cannot just compensate the reviewer off the equilibrium path. Screening would require that the expected sum of transfers also varies with type θ on the equilibrium path, which is not beneficial to the principal.

Given this result, we from now on focus without loss of generality on contracts where the agents' transfers depend on their efforts only, which we write as $\Phi = (\mu^e)_{e \in E}$. The next lemma shows that the principal does not benefit from implementing mixed strategies.

LEMMA 4. *For every contract Φ that implements a strategy profile σ , there exists a contract $\hat{\Phi}$ that implements the pure-strategy profile $\bar{e} = (\bar{e}_1, \dots, \bar{e}_n)$, where $\bar{e}_i = \mathbb{E}_{\sigma_i}[e_i] \forall i \in I$, and yields the same expected payoff to the principal.*

The intuition behind this result is simple: any randomness in transfers that is achieved by mixed strategies can equivalently be generated by the contract. On the other hand, since c is convex, the agents benefit from exerting the average effort \bar{e}_i instead of σ_i , while the principal is indifferent as to whether she obtains the efforts in expectation or deterministically.

The next lemma states that it is without loss to restrict attention to the implementation of symmetric effort profiles.

LEMMA 5. *For every contract Φ that implements a pure-strategy profile $\bar{e} = (\bar{e}_1, \dots, \bar{e}_n)$, there exists a contract $\hat{\Phi}$ that implements the symmetric pure-strategy profile $\hat{e} = (\hat{e}_1, \dots, \hat{e}_n)$, where $\hat{e}_1 = \dots = \hat{e}_n = \frac{1}{n} \sum_{i=1}^n \bar{e}_i$, and yields the same expected payoff to the principal.*

The next lemma completes the proof of [Theorem 1](#) by demonstrating that the principal can achieve the same payoff with a contest as with any nonscreening contract that implements a symmetric pure-strategy effort profile. Thus, the principal can obtain her maximal payoff with a contest.²⁰

LEMMA 6. *For every contract Φ that implements a symmetric pure-strategy profile \hat{e} , there exists a contest C_y that also implements \hat{e} and yields the same expected payoff to the principal.*

To prove this lemma, we construct a contest which implements the effort profile \hat{e} . This contest features $n - 1$ identical prizes and one prize that is smaller. The small prize is used to punish agents who deviate in either direction from \hat{e} . In equilibrium, when the effort profile \hat{e} is realized, the small prize is allocated randomly among the agents. As we will show in the next section, this particular prize structure is a general feature of optimal contests, while the specific (nonmonotonic) allocation rule is not required to achieve the optimum.

3.2 Optimal contests

From the previous section, we know that the principal can restrict attention to contests when designing an optimal contract. In this section, we characterize general features of all optimal contests. When describing a contest C_y , in the following we always assume w.l.o.g. that the prize profile y is ordered such that $y_1 \geq y_2 \geq \dots \geq y_n$.

THEOREM 2. *A contest is optimal if and only if conditions (i) and (ii) hold:*

(i) *The prizes satisfy $y_n = 0$ and $\sum_{k=1}^n y_k = x^*$, with $x^* = \min\{\bar{x}, \bar{T}\}$ and \bar{x} given by*

$$u' \left(\frac{\bar{x}}{n-1} \right) = c' \left(c^{-1} \left(\frac{n-1}{n} u \left(\frac{\bar{x}}{n-1} \right) \right) \right).$$

If the agents are risk-averse, then the prize profile is unique and given by

$$y = (x^*/(n-1), \dots, x^*/(n-1), 0).$$

(ii) *The contest implements (e^*, \dots, e^*) , where e^* is given by*

$$e^* = c^{-1} \left(\frac{n-1}{n} u \left(\frac{x^*}{n-1} \right) \right).$$

Condition (i) in the theorem shows that the lowest prize will be zero in any optimal contest. This is not obvious, since the agents can be risk-averse and in equilibrium all agents face the risk of receiving the zero prize. The intuition is that in equilibrium an agent receives the zero prize with probability $1/n$, while a shirking agent would receive

²⁰To be exact, [Theorem 1](#) follows only after it has been shown that problem (P) has a solution, so that an optimal contract exists. This will be shown in the proof of [Theorem 2](#) in the next section.

the zero prize with strictly larger probability (possibly one). Thus, any increase in y_n decreases the difference between the equilibrium and the deviation payoffs and, therefore, decreases the amount of effort that can be demanded in equilibrium. When the agents are risk-averse, the optimal prize profile will feature $n - 1$ identical positive prizes in addition to the zero prize, which keeps the risk imposed on the agents in equilibrium at a minimum. This prize structure is not novel, and its good incentive properties have been observed before in different settings.²¹ We also characterize the optimal total prize sum x^* , which is given by the point \bar{x} where marginal cost and benefit of inducing effort are equalized, or by the exogenous budget \bar{T} whenever it is sufficiently tight.

Condition (ii) in the theorem shows that every optimal contest extracts the entire surplus from the agents, because it implements a symmetric pure-strategy effort profile such that each agent's equilibrium payoff is zero. This condition puts limits on the set of possible contest success functions which can be used to achieve the optimum. We will explore these limits in Section 3.3.

Having characterized the optimal contests, we turn to the question of efficiency loss.

COROLLARY 1. *If the agents are risk-neutral, the principal can achieve the first-best. If the agents are risk-averse, the principal cannot achieve the first-best.*

The efficiency loss is driven entirely by risk-aversion of the agents. The loss is a direct consequence of the necessity to assign a prize of zero. Since the principal has to commit to the low prize, it must be delivered even in equilibrium. Risk-averse agents have to be compensated for this, which increases the cost of inducing effort. Hence, the commitment problem prevents the principal from achieving the first-best. However, the loss will be small if the agents are only mildly risk-averse, as the following example illustrates.

EXAMPLE. Consider again our example for $n = 2$. Applying the results from Theorem 2, it can be shown that $e^* = 2^{(\alpha-1)/(\beta-\alpha)} e^{\text{FB}}$ and $x^* = 2^{(\beta-1)/(\beta-\alpha)} t^{\text{FB}}$ holds in any optimal contest. The ratio of second-best to first-best profits of the principal is therefore

$$R = \frac{2e^* - x^*}{2e^{\text{FB}} - 2t^{\text{FB}}} = \frac{2^{\frac{\beta-1}{\beta-\alpha}} (e^{\text{FB}} - t^{\text{FB}})}{2(e^{\text{FB}} - t^{\text{FB}})} = 2^{\frac{\alpha-1}{\beta-\alpha}}.$$

This ratio is increasing in α , with $R \rightarrow 1$ in the limit as $\alpha \rightarrow 1$, so second-best profits approach first-best profits when the agents' risk-aversion vanishes.

²¹For example, Nalebuff and Stiglitz (1983) show that a single negative prize (a penalty) provides better incentives than a single positive prize. Glazer and Hassin (1988) show that the optimal way to split a budget into prizes is to have $n - 1$ positive prizes and one prize equal to zero. The reason for optimality of this prize profile is the same as in our model, namely the assumption that agents are risk-averse. Fang et al. (forthcoming) obtain a similar result for risk-neutral agents and convex effort costs. It is difficult to directly compare these papers to ours, since they consider only one specific contest success function that is taken as given. The contest success function that is optimal in our setting is different from those ones studied in these papers (see Subsection 3.3).

The loss will also be small if there are many agents, because the probability of not receiving any of the $n - 1$ identical prizes is $1/n$ in equilibrium and vanishes as the number of agents grows. This has been observed before by Glazer and Hassin (1988), and similar arguments can be found in Green and Stokey (1983) and Nalebuff and Stiglitz (1983). In our setting, these arguments imply that using mechanisms without conditional monetary transfers to the reviewer comes with little loss compared to an unconstrained mechanism design approach if the agents' risk-aversion is small or if the number of agents is large.

3.3 Unique implementation

We say that a contract Φ *uniquely implements* some pure-strategy effort profile e if it (i) implements e and (ii) does not implement any other (possibly mixed) strategy profile $\sigma \neq e$. The next theorem states that the second-best effort profile from Theorem 2 can be uniquely implemented by an all-pay auction with a cap.²²

An all-pay auction is one of the canonical contest types (see Konrad 2009, Chapter 2.1). It is perfectly discriminating, in the sense that the agent with the highest effort wins the highest prize with probability one, the agent with the second highest effort wins the second prize, and so on. Ties are broken randomly. A cap is a maximum admissible effort level (see Che and Gale 1998). In our setting, a cap could be implemented in two different ways. First, the contract could simply not differentiate effort levels at or above the cap, i.e., an agent who exerts effort exactly at the cap and an agent who exerts effort above the cap are treated the same and have the same chance of winning each of the prizes. With our optimal prize profile, this amounts to the rule that an agent with lowest effort receives the low prize, except if all agents have reached a prespecified performance threshold (the cap), in which case the low prize is allocated randomly. Second, an even simpler implementation is to put an actual upper bound on the effort that each agent can provide, e.g. by enforcing maximal work hours, limiting the accumulation of overtime, or imposing page limits and deadlines on grant proposals. This can often be done even though the principal cannot observe or control the actual effort. For example, by limiting the work day to 8 hours, the principal imposes an upper bound on the maximal effort by limiting the opportunities to exert effort. Making sure that an agent actually exerts effort during those 8 hours is a much more difficult problem. Hence, while imposing upper bounds on effort is often feasible, enforcing lower bounds is typically not possible. See Gaviious et al. (2002) for many other examples of actual caps that are imposed in different contests.

The following result shows that a cap at the optimal effort level e^* generates a unique equilibrium in the all-pay auction, which is in pure strategies.

THEOREM 3. *The all-pay auction with prize profile $y = (x^*/(n-1), \dots, x^*/(n-1), 0)$ and a cap at e^* uniquely implements the effort profile (e^*, \dots, e^*) .*

²²Uniqueness, as defined here, refers to the agents' choice of efforts in the given contest. The reviewer will always be indifferent among several actions, by Lemma 1. In particular, a "babbling" equilibrium exists in any credible contract. In such an equilibrium, the reviewer's report is uninformative about the agents' efforts and they exert zero effort. Hence, our notion of uniqueness is the strongest possible in this setting.

To see why all agents exerting e^* is an equilibrium, observe that upward deviations are ineffective (or impossible) while downward deviations guarantee the zero prize. The intuition for the result that no other pure-strategy equilibria exist is similar to that for all-pay auctions without caps: For every effort profile $e \neq (e^*, \dots, e^*)$, either an upward deviation discretely increases the probability of winning, or a downward deviation decreases costs without changing the probability of winning. The last step in the proof is to show that the cap destroys any potential mixed-strategy equilibrium. Note that the all-pay auction without a cap does not have an equilibrium in pure strategies. Therefore, the max–min optimal ranking mechanism in Frankel (2014) would not be optimal in our setting with endogenous efforts, because it is unable to implement the optimal effort profile.

An existing literature has investigated the effect of caps on equilibria and revenues in all-pay auctions (see Che and Gale 1998 for an early contribution and Olszewski and Siegel 2019 for a very general recent treatment). In particular, it is known that caps can generate pure-strategy equilibria. In our setting, e^* is the largest possible cap for which a pure-strategy equilibrium still exists. This is intuitive. The principal wants to make the contest as competitive as possible without destroying the pure-strategy equilibrium.

The all-pay auction with a cap is not the only contest success function that implements the optimum. It can be shown that a suitably defined Tullock contest is also able to implement the optimal effort profile.²³ This provides a novel foundation for the imperfectly discriminating Tullock contest success function. The noise in the prize allocation is explicitly generated within the optimal contract, rather than being the result of uncontrolled factors in the allocation process. Combined with the above Theorem 3, this also shows that the essential feature of our main result is the fixed profile of prizes, and not the exact way in which these prizes are awarded to the agents.

4. DELEGATION MECHANISMS

In our motivating example of performance pay within firms, the mid-level supervisor is often not asked to report to the principal but can himself determine which subordinate should get a bonus. That is, the principal *delegates* the choice of the agents' compensation to the reviewer. The classical delegation literature studies the design of a delegation set of actions that the principal makes available to the person who takes the final decision.

In our context, a delegation set is a set $D \subseteq \Delta T$ of (potentially stochastic) payments to the agents, from which the reviewer is free to choose. From our mechanism design perspective, a delegation set is just an indirect mechanism that must yield the same equilibrium outcome as some direct mechanism. Conversely, given an incentive-compatible direct mechanism, i.e., a credible contract Φ , a natural delegation set to consider is

$$D^\Phi = \{\mu \in \Delta T \mid \exists (e, \theta) \text{ s.t. } \mu = \mu^{e, \theta}\}.$$

If the principal makes the actions in D^Φ available to the reviewer, then the (IC-R) constraint implies that a reviewer of any type θ finds it optimal to choose $\mu^{e, \theta}$ from D^Φ when

²³Details are available upon request.

he observes any effort profile e .²⁴ Hence, the delegation set D^Φ can be used to replicate the outcome of the direct mechanism Φ .

The delegation set D^Φ is the smallest but generally not the only set that can be used to replicate the outcome of Φ . It is possible to add additional actions which the reviewer never finds strictly more attractive than the action $\mu^{e,\theta}$ that he is supposed to take. This is of particular interest when we consider the replication of a contest C_y by means of a delegation set. The smallest delegation set D^{C_y} , as defined above, does generally not include all possible random allocations of the prize profile y , only those that the specific contest success function prescribes for some effort profile e . For instance, an all-pay auction prescribes only deterministic allocations or uniform randomization in case of ties. A larger and more natural delegation set is the set

$$D^y = \{\mu \in \Delta T \mid \mu(P(y)) = 1\}$$

of all possible random allocations of the prize profile y . This literally reflects the idea that the reviewer is constrained only by the prizes but not in how he distributes them. It follows immediately from Lemma 2 that D^y can be used as a delegation set for *any* contest C_y , because *all* random allocations of y generate the same payoffs for the principal and the same sum of utilities for the agents.

Given that our optimal prize structure has $n - 1$ identical prizes, one can contemplate an even larger delegation set which forces the reviewer to pay zero to one agent but gives him freedom to allocate x^* arbitrarily among the remaining agents. Any reviewer with $\theta \geq 0$ would find it optimal to hand out $n - 1$ equal shares. Hence, if $\theta \geq 0$ is guaranteed, this larger delegation set is another indirect replication of the optimal contest.²⁵

5. CONCLUSION

We have analyzed a three-tiered structure consisting of a principal, a reviewer, and n agents. The principal designs a reward scheme to induce the agents to exert effort. However, the principal does not observe the efforts, so she relies on the reviewer to evaluate the agents' performance. The reviewer has private information about the utility weights he puts on the payoffs of the principal and the agents. Our main result is that a simple mechanism, a contest, is optimal. By providing commitment for a lenient reviewer to punish shirking agents, a contest efficiently incentivizes the agents to exert costly effort.

²⁴Formally, constraints (IC-R) and (IC-A) characterize perfect Bayesian equilibria of the following game. Given the delegation set D^Φ , the agents simultaneously choose their efforts e . Nature then determines the reviewer's type θ . The reviewer finally observes e and θ and selects an action from D^Φ . Constraint (IC-R) prescribes sequential rationality for the reviewer's (singleton) information sets, while (IC-A) prescribes sequential rationality (with weakly consistent beliefs) for the information sets in which the agents choose.

²⁵The larger set no longer works if $\theta < 0$ is possible, in which case the reviewer would find it optimal to pay x^* to a single agent regardless of actual performance. There is not much scope for giving more discretion to the reviewer than D^y if the delegation set is required to work for all possible reviewer type distributions. It is straightforward to show that we can only add stochastic transfers to D^y which generate the same expected sum of utilities of the agents as y and have an expected sum weakly larger than x^* .

As we have shown in an earlier working paper version (Letina et al. 2018), our results are robust in a range of extensions. First, the assumption that the reviewer perfectly observes the effort of each agent can be relaxed. As long as effort observation is not too incomplete or too noisy, the principal will still be able and find it optimal to implement the second-best effort from the noiseless environment, using the same optimal prize profile but adjusting the contest success function. Second, our optimal contest can remain effective with more general preferences of the reviewer. For instance, the reviewer may have concerns for fairness, resulting in a desire to give larger prizes to agents who have exerted higher efforts. This creates no fundamental problems for our all-pay auction with a cap, which is (almost) perfectly discriminating. A more challenging problem is favoritism. The reviewer will not be indifferent between all prize permutations if he prefers some agents over others. To address this problem, the principal could structure the observation process such that the chosen efforts are observable but not the identity of the agent who chose each effort. This is commonly referred to as a blind reviewing process and is indeed used for performance evaluation (e.g., see Goldin and Rouse 2000). Even if blind reviews are impossible, a contest may still outperform alternative contracts, because it restricts the scope of favoritism more than, e.g., a bonus pool contract. Finally, our model can be generalized to incorporate heterogeneous abilities of the agents. A loser-gets-nothing contest remains optimal if there are only two agents, or if the heterogeneity in the agents' abilities is not too large.²⁶

Of course, contests are not flawless mechanisms. Our model ignores three potential drawbacks: sabotage, collusion, and inequality concerns. First, since winning a contest requires outperforming the competitors, an agent may try to sabotage the effort of his competitors instead of exerting productive effort.²⁷ Whether this is an issue will depend on how difficult it is to sabotage the competitors, and how well the reviewer can observe attempts at sabotage. In our setting, it is credible to assign the zero prize whenever any sabotage attempt is observed. Hence, if sabotage is observable with a sufficiently high probability, it can be dealt with just like deviations from the optimal effort profile. In other cases, sabotage will limit the extent to which a contest can be used in practice.

Second, collusion is an attractive option for contestants.²⁸ In our all-pay auction with a cap, collusion on an effort level lower than the cap may be fragile, because a marginal increase in effort would guarantee the deviator a positive prize. Furthermore, as Bandiera et al. (2005) point out, the ability of coworkers to monitor each other's output is necessary for collusion. Whether collusion can arise in a contest will therefore again depend on the specific application. It is important to note that contest mechanisms are not unique in their susceptibility to collusion. Laffont and Martimort (1997,

²⁶With an arbitrary number of agents and arbitrary heterogeneity in effort costs, a *generalized contest* is always optimal. Such a contest is described by two prize profiles y and y^d , where the prizes y are allocated in equilibrium and the prizes y^d are allocated after deviations. In general, the prize profile used in equilibrium may feature more than two prize levels.

²⁷Berger et al. (2013) find that, when given the opportunity, subjects in their laboratory experiment do indeed engage in some sabotage.

²⁸Using personnel data from a large UK farm, Bandiera et al. (2005) show that workers paid under a relative incentive scheme do collude. Fairburn and Malcomson (2001) consider collusion between workers and the supervisor and show how it can reduce incentives for agents to exert effort.

1998) argue convincingly that collusion is an issue in a standard mechanism design problem with hidden information. In addition to agents colluding among themselves, agents could also collude with the supervisor, as in Faure-Grimaud et al. (2003). Interestingly, both Laffont and Martimort (1998) and Faure-Grimaud et al. (2003) find that delegation to the supervisor can be a response to the possibility of collusion. An additional way to mitigate the danger of collusion, suggested by Laffont and Martimort (1999), is to introduce multiple supervisors.

Finally, a contest generates ex post inequality even when all agents have exerted the expected level of effort, which can be perceived as unfair and can generate additional welfare losses (on top of the welfare loss from ex ante uncertainty). While some inequality is necessary to sustain incentives in our setting, the loser-gets-nothing contest generates a particular type of inequality. Inequality concerns may put additional constraints on the prize profile that can be used in practice.

Despite these potential drawbacks, contests are widely used in reality. Understanding precisely when the benefits of contests outweigh their shortcomings is a promising area for future research.

APPENDIX: PROOFS

A.1 Proof of Theorem 1

A.1.1 *Proof of Lemma 1 If-statement.* We first show that (IC-R) is implied by (i)–(iii). Note that (IC-R) can be rewritten as

$$S(e, \theta) - S(e', \theta') \geq (\theta - \theta') S_u(e', \theta') \quad \forall (e, \theta), (e', \theta') \in E \times \Theta.$$

Using (i) and (iii), this is equivalent to the requirement that, $\forall e' \in E$ and $\forall \theta, \theta' \in \Theta$,

$$\int_{\theta'}^{\theta} (S_u(e', s) - S_u(e', \theta')) ds \geq 0,$$

and this inequality indeed holds since $S_u(e', \theta)$ is nondecreasing in θ by (ii).

Only-if-statement. We now proceed to prove that (IC-R) implies (i)–(iii). Note that for the special case $\theta' = \theta$, (IC-R) is reduced to the requirement that $S(e, \theta) \geq S(e', \theta)$ $\forall e, e' \in E$. Interchanging e and e' , we immediately obtain (i). Next, consider the special case where $e' = e$. For this case, (IC-R) requires that, $\forall \theta, \theta' \in \Theta$,

$$S(e, \theta) \geq -S_t(e, \theta') + \theta S_u(e, \theta') \tag{IC}_{\theta, \theta'}$$

and

$$S(e, \theta') \geq -S_t(e, \theta) + \theta' S_u(e, \theta). \tag{IC}_{\theta', \theta}$$

Summing up $(IC)_{\theta, \theta'}$ and $(IC)_{\theta', \theta}$ we obtain

$$(\theta - \theta')(S_u(e, \theta) - S_u(e, \theta')) \geq 0 \quad \forall \theta, \theta' \in \Theta.$$

Thus, $S_u(e, \theta)$ must be nondecreasing in θ , which is condition (ii). The envelope formula in (iii) follows directly from Theorem 2 of Milgrom and Segal (2002), where absolute continuity of $S(e, \theta)$ holds because the set of transfer profiles is bounded.

A.1.2 *Proof of Lemma 2* By Lemma 1, credibility implies that, $\forall e, e' \in E$ and $\forall \theta \in \Theta$,

$$\delta(e, e', \theta) = S(e, \theta) - S(e', \theta) = \int_{\underline{\theta}}^{\theta} (S_u(e, s) - S_u(e', s)) ds = 0.$$

This implies that, for any fixed $e, e' \in E$, $S_u(e, \theta) = S_u(e', \theta)$ for almost every $\theta \in \Theta$. It then also immediately follows that $S_t(e, \theta) = S_t(e', \theta)$ for almost every $\theta \in \Theta$. Now choose an arbitrary $e' \in E$ and define the functions x and \hat{x} by

$$x(\theta) = S_t(e', \theta), \quad \hat{x}(\theta) = S_u(e', \theta) \quad \forall \theta \in \Theta.$$

It follows that, for any $e \in E$, $S_t(e, \theta) = x(\theta)$ and $S_u(e, \theta) = \hat{x}(\theta)$ for almost all $\theta \in \Theta$.

A.1.3 *Proof of Lemma 3* Suppose $\Phi = (\mu^{e,\theta})_{(e,\theta) \in E \times \Theta}$ implements σ . In particular, Φ is credible, so by Lemma 2 there exists a pair of function x and \hat{x} such that, $\forall e \in E$,

$$\mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n t_i \right] = x(\theta), \quad \mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n u(t_i) \right] = \hat{x}(\theta),$$

for almost every $\theta \in \Theta$. Since Φ implements σ , we also have $\forall i \in I$ and $\forall \sigma'_i \in \Delta \mathbb{R}_+$,

$$\mathbb{E}_{\sigma} [\mathbb{E}_{\tau} [\mathbb{E}_{\mu^{e,\theta}} [u(t_i)]] - c(e_i)] \geq \mathbb{E}_{(\sigma'_i, \sigma_{-i})} [\mathbb{E}_{\tau} [\mathbb{E}_{\mu^{e,\theta}} [u(t_i)]] - c(e_i)].$$

The expected payoff of the principal with (σ, Φ) is given by

$$\Pi_P(\sigma, \Phi) = \mathbb{E}_{\sigma} \left[\sum_{i=1}^n e_i - \mathbb{E}_{\tau} \left[\mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n t_i \right] \right] \right] = \mathbb{E}_{\sigma} \left[\sum_{i=1}^n e_i \right] - \mathbb{E}_{\tau} [x(\theta)].$$

For every $e \in E$, define a probability measure $\mu^e \in \Delta T$ such that

$$\mu^e(A) = \mathbb{E}_{\tau} [\mu^{e,\theta}(A)]$$

for all measurable subsets $A \subseteq T$.²⁹ Now construct an alternative contract $\hat{\Phi}$ by setting $\hat{\mu}^{e,\theta} = \mu^e$ for all $(e, \theta) \in E \times \Theta$. This contract satisfies the property of θ -independence stated in the lemma. Since, $\forall (e, \theta) \in E \times \Theta$,

$$\begin{aligned} \hat{S}_t(e, \theta) &= \mathbb{E}_{\hat{\mu}^{e,\theta}} \left[\sum_{i=1}^n t_i \right] = \mathbb{E}_{\mu^e} \left[\sum_{i=1}^n t_i \right] = \mathbb{E}_{\tau} \left[\mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n t_i \right] \right] = \mathbb{E}_{\tau} [x(\theta)], \\ \hat{S}_u(e, \theta) &= \mathbb{E}_{\hat{\mu}^{e,\theta}} \left[\sum_{i=1}^n u(t_i) \right] = \mathbb{E}_{\mu^e} \left[\sum_{i=1}^n u(t_i) \right] = \mathbb{E}_{\tau} \left[\mathbb{E}_{\mu^{e,\theta}} \left[\sum_{i=1}^n u(t_i) \right] \right] = \mathbb{E}_{\tau} [\hat{x}(\theta)], \end{aligned}$$

by Lemma 1 it is straightforward to check that $\hat{\Phi}$ is credible. Furthermore, note that

$$\Pi_i(\sigma', \hat{\Phi}) = \mathbb{E}_{\sigma'} [\mathbb{E}_{\tau} [\mathbb{E}_{\hat{\mu}^{e,\theta}} [u(t_i)]] - c(e_i)]$$

²⁹The assumption discussed in footnote 16 ensures that the expectation (as well as the ones in the proof of the next lemma) is well-defined. It is also easy to show that μ^e is indeed a probability measure.

$$\begin{aligned}
&= \mathbb{E}_{\sigma'} [\mathbb{E}_{\tau} [\mathbb{E}_{\mu^e} [u(t_i)]]] - c(e_i) \\
&= \mathbb{E}_{\sigma'} [\mathbb{E}_{\tau} [\mathbb{E}_{\mu^{e,\theta}} [u(t_i)]]] - c(e_i) \\
&= \Pi_i(\sigma', \Phi)
\end{aligned}$$

for all σ' and $i \in I$, which implies that $\hat{\Phi}$ implements σ because Φ implements σ .

Finally, from the above arguments we also obtain that the principal's expected payoff is $\mathbb{E}_{\sigma}[\sum_{i=1}^n e_i] - \mathbb{E}_{\tau}[x(\theta)]$ with both (σ, Φ) and $(\sigma, \hat{\Phi})$.

A.1.4 Proof of Lemma 4 Suppose $\Phi = (\mu^e)_{e \in E}$ implements σ . We first construct a probability measure $\eta \in \Delta T$ by

$$\eta(A) = \mathbb{E}_{\sigma}[\mu^e(A)]$$

for all measurable subsets $A \subseteq T$. Furthermore, for each $i \in I$ we construct a probability measure $\eta^{(i)} \in \Delta T$ by setting

$$\eta^{(i)}(A) = \mathbb{E}_{\sigma}[\mu^{(0, e_{-i})}(A)]$$

for all measurable subsets $A \subseteq T$. We now construct an alternative contract $\hat{\Phi} = (\hat{\mu}^e)_{e \in E}$ as follows. For $e = \bar{e}$, we let $\hat{\mu}^e = \eta$. For any $e = (e_i, \bar{e}_{-i})$ with $e_i \neq \bar{e}_i$, we let $\hat{\mu}^e = \eta^{(i)}$. For all remaining e , we let $\hat{\mu}^e = \mu^e$.

We first show that $\hat{\Phi}$ is credible. Since Φ is credible and its transfers are independent of θ , by Lemma 2 there exist $x, \hat{x} \in \mathbb{R}_+$ such that $\mathbb{E}_{\mu^e}[\sum_{i=1}^n t_i] = x$ and $\mathbb{E}_{\mu^e}[\sum_{i=1}^n u(t_i)] = \hat{x}$ for all $e \in E$. First consider $\hat{\mu}^e$ for $e = \bar{e}$. We obtain

$$\mathbb{E}_{\hat{\mu}^{\bar{e}}} \left[\sum_{i=1}^n t_i \right] = \mathbb{E}_{\eta} \left[\sum_{i=1}^n t_i \right] = \mathbb{E}_{\sigma} \left[\mathbb{E}_{\mu^e} \left[\sum_{i=1}^n t_i \right] \right] = \mathbb{E}_{\sigma}[x] = x$$

and, by the analogous argument, $\mathbb{E}_{\hat{\mu}^{\bar{e}}}[\sum_{i=1}^n u(t_i)] = \hat{x}$. Now consider $\hat{\mu}^e$ for $e = (e_i, \bar{e}_{-i})$ with $e_i \neq \bar{e}_i$. We obtain

$$\mathbb{E}_{\hat{\mu}^{(e_i, \bar{e}_{-i})}} \left[\sum_{i=1}^n t_i \right] = \mathbb{E}_{\eta^{(i)}} \left[\sum_{i=1}^n t_i \right] = \mathbb{E}_{\sigma} \left[\mathbb{E}_{\mu^{(0, e_{-i})}} \left[\sum_{i=1}^n t_i \right] \right] = \mathbb{E}_{\sigma}[x] = x$$

and, by the analogous argument, $\mathbb{E}_{\hat{\mu}^{(e_i, \bar{e}_{-i})}}[\sum_{i=1}^n u(t_i)] = \hat{x}$. Since $\hat{\Phi}$ and Φ are identical for all other e , we can conclude that $\mathbb{E}_{\hat{\mu}^e}[\sum_{i=1}^n t_i] = x$ and $\mathbb{E}_{\hat{\mu}^e}[\sum_{i=1}^n u(t_i)] = \hat{x}$ for all $e \in E$. It is then straightforward to check that $\hat{\Phi}$ is credible by using Lemma 1.

We next show that, in $\hat{\Phi}$, for each agent $i \in I$ it is a best response to play \bar{e}_i when the remaining agents are playing \bar{e}_{-i} , which implies that $\hat{\Phi}$ implements \bar{e} . This claim holds because, $\forall i \in I$ and $\forall e'_i \neq \bar{e}_i$,

$$\begin{aligned}
\Pi_i(\bar{e}, \hat{\Phi}) &= \mathbb{E}_{\eta}[u(t_i)] - c(\bar{e}_i) \\
&= \mathbb{E}_{\sigma}[\mathbb{E}_{\mu^e}[u(t_i)]] - c(\mathbb{E}_{\sigma}[e_i]) \\
&\geq \mathbb{E}_{\sigma}[\mathbb{E}_{\mu^e}[u(t_i)]] - \mathbb{E}_{\sigma}[c(e_i)]
\end{aligned}$$

$$\begin{aligned} &\geq \mathbb{E}_\sigma[\mathbb{E}_{\mu^{(0, e_{-i})}}[u(t_i)]] \\ &\geq \mathbb{E}_\sigma[\mathbb{E}_{\mu^{(0, e_{-i})}}[u(t_i)]] - c(e'_i) \\ &= \mathbb{E}_{\eta^{(i)}}[u(t_i)] - c(e'_i) = \Pi_i((e'_i, \bar{e}_{-i}), \hat{\Phi}), \end{aligned} \tag{1}$$

where the first inequality follows the convexity of c and the second inequality follows from the fact that Φ implements σ .

Finally, from the above arguments we also obtain that the principal's expected payoff is $\sum_{i=1}^n \bar{e}_i - x$ with both (σ, Φ) and $(\bar{e}, \hat{\Phi})$.

A.1.5 Proof of Lemma 5 Suppose $\Phi = (\mu^e)_{e \in E}$ implements \bar{e} . We now construct an alternative contract $\hat{\Phi} = (\hat{\mu}^e)_{e \in E}$ as follows. For $e = \hat{e}$, we define $\hat{\mu}^e$ by generating a profile of prizes $t = (t_1, \dots, t_n)$ according to $\mu^{\bar{e}}$ and then allocating these prizes randomly and uniformly among the agents. For any $e = (e_i, \hat{e}_{-i})$ with $e_i \neq \hat{e}_i$, we let $\hat{\mu}^e$ be given as follows. A number j is drawn uniformly from I and then a profile of prizes $t = (t_1, \dots, t_n)$ is generated according to $\mu^{(0, \bar{e}_{-j})}$. The deviating agent i gets the prize t_j and the remaining $n - 1$ prizes are allocated randomly and uniformly among the nondeviating agents. Note that, by construction, this punishment rule for unilateral deviations does not depend on the identity of the agent being punished. For all remaining e , we let $\hat{\mu}^e = \mu^e$.

We first show that $\hat{\Phi}$ is credible. By Lemma 2, credibility and θ -independence of Φ imply that there exists $x, \hat{x} \in \mathbb{R}_+$ such that $\mathbb{E}_{\mu^e}[\sum_{i=1}^n t_i] = x$ and $\mathbb{E}_{\mu^e}[\sum_{i=1}^n u(t_i)] = \hat{x}$ for all $e \in E$. Now first consider $\hat{\mu}^e$ for $e = \hat{e}$. We obtain

$$\begin{aligned} \mathbb{E}_{\hat{\mu}^{\hat{e}}} \left[\sum_{i=1}^n t_i \right] &= \mathbb{E}_{\mu^{\bar{e}}} \left[\sum_{i=1}^n t_i \right] = x, \\ \mathbb{E}_{\hat{\mu}^{\hat{e}}} \left[\sum_{i=1}^n u(t_i) \right] &= \mathbb{E}_{\mu^{\bar{e}}} \left[\sum_{i=1}^n u(t_i) \right] = \hat{x}. \end{aligned}$$

Now consider $\hat{\mu}^e$ for any $e = (e_i, \hat{e}_{-i})$ with $e_i \neq \hat{e}_i$. We obtain

$$\begin{aligned} \mathbb{E}_{\hat{\mu}^{(e_i, \hat{e}_{-i})}} \left[\sum_{i=1}^n t_i \right] &= \sum_{j=1}^n \frac{1}{n} \mathbb{E}_{\mu^{(0, \bar{e}_{-j})}} \left[\sum_{i=1}^n t_i \right] = \frac{1}{n} \sum_{j=1}^n x = x, \\ \mathbb{E}_{\hat{\mu}^{(e_i, \hat{e}_{-i})}} \left[\sum_{i=1}^n u(t_i) \right] &= \sum_{j=1}^n \frac{1}{n} \mathbb{E}_{\mu^{(0, \bar{e}_{-j})}} \left[\sum_{i=1}^n u(t_i) \right] = \frac{1}{n} \sum_{j=1}^n \hat{x} = \hat{x}. \end{aligned}$$

Since $\hat{\Phi}$ and Φ are identical for all other e , we can conclude that $\mathbb{E}_{\hat{\mu}^e}[\sum_{i=1}^n t_i] = x$ and $\mathbb{E}_{\hat{\mu}^e}[\sum_{i=1}^n u(t_i)] = \hat{x}$ for all $e \in E$. It is then straightforward to check that $\hat{\Phi}$ is credible by using Lemma 1.

We next show that, in $\hat{\Phi}$, for each agent $i \in I$ it is a best response to play \hat{e}_i when the remaining agents are playing \hat{e}_{-i} , which implies that $\hat{\Phi}$ implements \hat{e} . To prove this claim, note that

$$\mathbb{E}_{\mu^{\bar{e}}} [u(t_i)] - c(\bar{e}_i) \geq \mathbb{E}_{\mu^{(0, \bar{e}_{-i})}} [u(t_i)]$$

holds for all $i \in I$ because Φ implements \bar{e} . Summing over all $i \in I$ and dividing by n yields

$$\mathbb{E}_{\mu^{\bar{e}}} \left[\frac{1}{n} \sum_{k=1}^n u(t_k) \right] - \frac{1}{n} \sum_{k=1}^n c(\bar{e}_k) \geq \frac{1}{n} \sum_{k=1}^n \mathbb{E}_{\mu^{(0, \bar{e}_k)}} [u(t_k)].$$

We now obtain, $\forall i \in I$ and $\forall e_i \neq \hat{e}_i$,

$$\begin{aligned} \Pi_i(\hat{e}, \hat{\Phi}) &= \mathbb{E}_{\mu^{\hat{e}}} [u(t_i)] - c(\hat{e}_i) \\ &= \mathbb{E}_{\mu^{\bar{e}}} \left[\frac{1}{n} \sum_{k=1}^n u(t_k) \right] - c(\hat{e}_i) \\ &\geq \mathbb{E}_{\mu^{\bar{e}}} \left[\frac{1}{n} \sum_{k=1}^n u(t_k) \right] - \frac{1}{n} \sum_{k=1}^n c(\bar{e}_k) \\ &\geq \frac{1}{n} \sum_{k=1}^n \mathbb{E}_{\mu^{(0, \bar{e}_k)}} [u(t_k)] - c(e_i) \\ &= \mathbb{E}_{\mu^{(e_i, \hat{e}_{-i})}} [u(t_i)] - c(e_i) = \Pi_i((e_i, \hat{e}_{-i}), \hat{\Phi}), \end{aligned}$$

where the first inequality follows from convexity of c . Hence the claim follows.

Finally, from the above arguments we also obtain that the principal's expected payoff is $\sum_{i=1}^n \bar{e}_i - x$ with both (\bar{e}, Φ) and $(\hat{e}, \hat{\Phi})$.

A.1.6 Proof of Lemma 6 Suppose $\Phi = (\mu^e)_{e \in E}$ implements the symmetric profile \hat{e} . From the proof of Lemma 5 we know that it is without loss of generality to assume that Φ has the following form. If $e = \hat{e}$, a profile of prizes $t = (t_1, \dots, t_n)$ is generated according to some probability measure π and these prizes are randomly and uniformly allocated to the agents. If $e = (e_i, \hat{e}_{-i})$ with $e_i \neq \hat{e}_i$ for some $i \in I$, a profile of prizes $t^d = (t_1^d, \dots, t_n^d)$ is generated according to some (i -independent) probability measure ρ and agent i gets t_n^d , while the remaining $n - 1$ prizes are randomly and uniformly allocated among the other agents. For all other effort profiles e , the transfer rule can be chosen as for \hat{e} . Thus, we have

$$\mathbb{E}_{\mu^{\hat{e}}} [u(t_i)] = \mathbb{E}_{\pi} \left[\frac{1}{n} \sum_{k=1}^n u(t_k) \right], \quad \mathbb{E}_{\mu^{(e_i, \hat{e}_{-i})}} [u(t_i)] = \mathbb{E}_{\rho} [u(t_n^d)].$$

Furthermore, by Lemma 2, credibility and θ -independence of Φ imply that there exist $x, \hat{x} \in \mathbb{R}_+$ such that $\mathbb{E}_{\mu^e} [\sum_{i=1}^n t_i] = x$ and $\mathbb{E}_{\mu^e} [\sum_{i=1}^n u(t_i)] = \hat{x}$ for all $e \in E$.

Now construct a contest C_y with prize profile y as follows. Define \underline{t}^d as the certainty equivalent of a deviating agent's random transfers in contract Φ , i.e., $u(\underline{t}^d) = \mathbb{E}_{\rho} [u(t_n^d)]$. Note that $\underline{t}^d \leq \mathbb{E}_{\rho} [t_n^d]$ by concavity of u . Then define the prize profile

$$y = \left(\frac{x - \underline{t}^d}{n - 1}, \dots, \frac{x - \underline{t}^d}{n - 1}, \underline{t}^d \right).$$

The allocation rule of C_y is as follows. If $e = \hat{e}$, the prizes are randomly and uniformly allocated among all agents. If $e = (e_i, \hat{e}_{-i})$ with $e_i \neq \hat{e}_i$ for some $i \in I$, the deviating agent i obtains \underline{t}^d and all other agents obtain $(x - \underline{t}^d)/(n - 1)$. For all other effort profiles e , the prizes are again randomly and uniformly allocated among all agents.

Since C_y is a contest, it is credible. Furthermore, $\forall i \in I$ and $\forall e_i \neq \hat{e}_i$,

$$\begin{aligned}
 \Pi_i(\hat{e}, C_y) &= \left(\frac{n-1}{n}\right) u\left(\frac{x - \underline{t}^d}{n-1}\right) + \frac{1}{n} u(\underline{t}^d) - c(\hat{e}_i) \\
 &\geq \left(\frac{n-1}{n}\right) u\left(\frac{\mathbb{E}_\rho\left[\sum_{k=1}^n t_k^d\right] - \mathbb{E}_\rho[t_n^d]}{n-1}\right) + \frac{1}{n} \mathbb{E}_\rho[u(t_n^d)] - c(\hat{e}_i) \\
 &= \left(\frac{n-1}{n}\right) u\left(\mathbb{E}_\rho\left[\sum_{k=1}^{n-1} \frac{1}{n-1} t_k^d\right]\right) + \frac{1}{n} \mathbb{E}_\rho[u(t_n^d)] - c(\hat{e}_i) \\
 &\geq \left(\frac{n-1}{n}\right) \mathbb{E}_\rho\left[u\left(\sum_{k=1}^{n-1} \frac{1}{n-1} t_k^d\right)\right] + \frac{1}{n} \mathbb{E}_\rho[u(t_n^d)] - c(\hat{e}_i) \\
 &\geq \left(\frac{n-1}{n}\right) \mathbb{E}_\rho\left[\sum_{k=1}^{n-1} \frac{1}{n-1} u(t_k^d)\right] + \frac{1}{n} \mathbb{E}_\rho[u(t_n^d)] - c(\hat{e}_i) \\
 &= \mathbb{E}_\rho\left[\frac{1}{n} \sum_{k=1}^{n-1} u(t_k^d)\right] + \mathbb{E}_\rho\left[\frac{1}{n} u(t_n^d)\right] - c(\hat{e}_i) \\
 &= \mathbb{E}_\rho\left[\frac{1}{n} \sum_{k=1}^n u(t_k^d)\right] - c(\hat{e}_i) \\
 &= \mathbb{E}_\pi\left[\frac{1}{n} \sum_{k=1}^n u(t_k)\right] - c(\hat{e}_i) \\
 &\geq \mathbb{E}_\rho[u(t_n^d)] - c(e_i) \\
 &= u(\underline{t}^d) - c(e_i) = \Pi_i((e_i, \hat{e}_{-i}), C_y),
 \end{aligned}$$

where the first inequality follows from $\underline{t}^d \leq \mathbb{E}_\rho[t_n^d]$, the second and third inequalities follow from concavity of u , and the last inequality follows from the fact that Φ implements \hat{e} . We can thus conclude that C_y also implements \hat{e} .

Finally, from the above arguments we also obtain that the principal's expected payoff is $\sum_{i=1}^n \hat{e}_i - x$ with both (\hat{e}, Φ) and (\hat{e}, C_y) .

A.2 Proof of Theorem 2

Only-if-statement. Suppose (σ^*, C_y^*) solves (P). We first claim that σ^* must be a pure-strategy effort profile. By contradiction, suppose there exists $j \in I$ such that σ_j^* is not a

Dirac measure. We can now proceed exactly as in the proof of Lemma 4 to construct a contract $\hat{\Phi}$ (in fact, a contest) that implements a pure-strategy profile \bar{e} . The only difference to the proof of Lemma 4 is that we let $\bar{e}_j = \mathbb{E}_{\sigma_j^*}[e_j] + \epsilon$ for some $\epsilon > 0$ (but still $\bar{e}_i = \mathbb{E}_{\sigma_i^*}[e_i]$ for all $i \neq j$). Credibility of $\hat{\Phi}$ and the fact that \bar{e}_i is a best response to \bar{e}_{-i} for all $i \neq j$ follow exactly as in the proof of Lemma 4. Since c is strictly convex and σ_j^* is not a Dirac measure, the first inequality in (1) is strict for j when $\epsilon = 0$. It then follows that \bar{e}_j is a best response to \bar{e}_{-j} for sufficiently small $\epsilon > 0$. Since the principal's payoff with $(\bar{e}, \hat{\Phi})$ is increased by ϵ , (σ^*, C_y^*) cannot have been a solution to (P).

Now suppose (\bar{e}, C_y^*) solves (P), where \bar{e} may still be asymmetric. Denote $x = \sum_{k=1}^n y_k$. We next show that whenever $y_n > 0$, there exists another contest $C_{\tilde{y}}$ with $\sum_{k=1}^n \tilde{y}_k = x$ that implements an effort profile \tilde{e} with $\sum_{i=1}^n \tilde{e}_i > \sum_{i=1}^n \bar{e}_i$, and hence (\bar{e}, C_y^*) cannot have been a solution to (P). Denote by $p_i^k(e)$ the probability that agent i receives prize y_k in C_y^* when the effort profile is e . Note that

$$\Pi_i(\bar{e}, C_y^*) = \sum_{k=1}^n p_i^k(\bar{e})u(y_k) - c(\bar{e}_i) \geq \sum_{k=1}^n p_i^k(0, \bar{e}_{-i})u(y_k) \geq u(y_n),$$

because C_y^* implements \bar{e} . Now consider an agent $j \in I$ for which $p_j^n(\bar{e}) < 1$. Construct a contest $C_{\tilde{y}}$ with a profile of prizes \tilde{y} given by $\tilde{y}_1 = y_1 + \delta$, $\tilde{y}_n = y_n - \delta$, and $\tilde{y}_k = y_k$ for all $k \neq 1, n$, where $\delta \in (0, y_n]$. Note that $\sum_{k=1}^n \tilde{y}_k = x$. Let effort profile \tilde{e} be such that $\tilde{e}_j = \bar{e}_j + \epsilon$ and $\tilde{e}_i = \bar{e}_i$ for all $i \neq j$, where $\epsilon > 0$. Note that $\sum_{i=1}^n \tilde{e}_i > \sum_{i=1}^n \bar{e}_i$. The rule of contest $C_{\tilde{y}}$ is the following. If the effort profile is \tilde{e} , then the prizes \tilde{y} are allocated such that each agent i receives prize \tilde{y}_k with probability $\tilde{p}_i^k(\tilde{e}) = p_i^k(\bar{e})$. If some agent i unilaterally deviates from \tilde{e} , then agent i receives the prize \tilde{y}_n , while the prizes $\tilde{y}_1, \dots, \tilde{y}_{n-1}$ are allocated randomly and uniformly among the remaining agents. Otherwise, the allocation of prizes can be chosen arbitrarily. For sufficiently small $\epsilon > 0$ we then have, $\forall i \in I$ and $\forall e_i \in \mathbb{R}_+$,

$$\begin{aligned} \Pi_i(\tilde{e}, C_{\tilde{y}}) &= \sum_{k=1}^n \tilde{p}_i^k(\tilde{e})u(\tilde{y}_k) - c(\tilde{e}_i) \\ &= \Pi_i(\bar{e}, C_y^*) + p_i^1(\bar{e})(u(y_1 + \delta) - u(y_1)) \\ &\quad + p_i^n(\bar{e})(u(y_n - \delta) - u(y_n)) + c(\bar{e}_i) - c(\tilde{e}_i) \\ &\geq u(y_n) + p_i^1(\bar{e})(u(y_1 + \delta) - u(y_1)) \\ &\quad + p_i^n(\bar{e})(u(y_n - \delta) - u(y_n)) + c(\bar{e}_i) - c(\tilde{e}_i) \\ &\geq u(y_n - \delta) = u(\tilde{y}_n) \geq \Pi_i((e_i, \tilde{e}_{-i}), C_{\tilde{y}}), \end{aligned}$$

where the second inequality holds because

$$u(y_n) + p_i^n(\bar{e})(u(y_n - \delta) - u(y_n)) \geq u(y_n - \delta)$$

for all $i \in I$, with strict inequality for j . Hence $C_{\tilde{y}}$ implements \tilde{e} .

When studying the set of all contest solutions to (P), we thus need to consider only pure-strategy effort profiles \bar{e} and contests C_y with $y_n = 0$. Fix a sum of prizes $x \in [0, \bar{T}]$. Let e^x be the (unique) effort level that solves

$$\frac{n-1}{n}u\left(\frac{x}{n-1}\right) - c(e^x) = 0.$$

Note that, by the assumptions on u and c , the solution e^x is differentiable, strictly increasing and strictly concave in x . We now claim that ne^x is an upper bound on the sum of efforts implementable with a contest C_y that has $\sum_{k=1}^n y_k = x$ and $y_n = 0$, and it can be reached only by implementing the symmetric effort profile (e^x, \dots, e^x) . Suppose first that C_y implements an effort profile \bar{e} with $\sum_{i=1}^n \bar{e}_i \geq ne^x$ but $\bar{e} \neq (e^x, \dots, e^x)$. Note that

$$\Pi_i(\bar{e}, C_y) = \sum_{k=1}^n p_i^k(\bar{e})u(y_k) - c(\bar{e}_i) \geq u(y_n) - c(0) = 0,$$

because C_y implements \bar{e} . Summing these inequalities over all agents we obtain

$$\sum_{i=1}^n \sum_{k=1}^n p_i^k(\bar{e})u(y_k) - \sum_{i=1}^n c(\bar{e}_i) = \sum_{k=1}^{n-1} u(y_k) - \sum_{i=1}^n c(\bar{e}_i) \geq 0.$$

However, due to weak concavity of u and strict convexity of c we also have

$$\sum_{k=1}^{n-1} u(y_k) - \sum_{i=1}^n c(\bar{e}_i) < (n-1)u\left(\frac{x}{n-1}\right) - nc(e^x) = 0,$$

a contradiction. Observe next that (e^x, \dots, e^x) can indeed be implemented. For instance, let $y = (x/(n-1), \dots, x/(n-1), 0)$ and choose the rules of C_y as follows. If the effort profile is (e^x, \dots, e^x) , then the prizes are allocated randomly and uniformly across the agents. If some agent i unilaterally deviates from (e^x, \dots, e^x) , then agent i receives the prize 0, while each other agent receives $x/(n-1)$. Otherwise, the allocation of prizes can be chosen arbitrarily. It follows immediately from the definition of e^x that this contest indeed implements (e^x, \dots, e^x) .

Given any sum of prizes x , the highest payoff that the principal can achieve is thus given by $\Pi_P(x) = ne^x - x$, and the problem is reduced to a choice of $x \in [0, \bar{T}]$. Since Π_P is continuous in x , it follows that a solution exists. Furthermore, since Π_P is differentiable and strictly concave, the first-order condition $\partial \Pi_P / \partial x = 0$ that is stated in part (i) of the theorem uniquely characterizes a value $\bar{x} > 0$ (given the assumptions on u and c), and the optimal value of x is given by $x^* = \min\{\bar{x}, \bar{T}\}$. The resulting implemented optimal effort level is then given by $e^* = e^{x^*}$.

We complete the proof of the only-if-statement by showing that any optimal contest has the profile of prizes $y = (x^*/(n-1), \dots, x^*/(n-1), 0)$ whenever u is strictly concave. By contradiction, let C_y be a contest that implements (e^*, \dots, e^*) with $\sum_{k=1}^n y_k = x^*$ and $y_n = 0$ but $y_1 \neq y_{n-1}$. Proceeding as before, summing the inequalities $\Pi_i((e^*, \dots, e^*), C_y) \geq 0$ over all agents yields $\sum_{k=1}^{n-1} u(y_k) - nc(e^*) \geq 0$. Strict concavity

of u , however, implies that $\sum_{k=1}^{n-1} u(y_k) - nc(e^*) < (n-1)u(x^*/(n-1)) - nc(e^*) = 0$, a contradiction.

If-statement. We showed above that the upper bound on the principal's payoff is given by $ne^* - x^*$. Thus, any contest which implements (e^*, \dots, e^*) with the prize sum x^* attains the upper bound.

A.3 Proof of Corollary 1

Each optimal contest induces individual efforts of e^* and pays a sum of x^* , as characterized in Theorem 2. Now consider the principal's first-best problem. If the agents' efforts were directly observable and verifiable, then the principal could ask for individual efforts of e and would have to compensate the agents with a transfer sum x such that $u(x/n) - c(e) = 0$. Put differently, for a given transfer sum x the maximal achievable individual effort is

$$e^x = c^{-1}\left(u\left(\frac{x}{n}\right)\right),$$

and the first-best problem is to maximize $ne^x - x$ by choice of $x \in [0, \bar{T}]$. With the same arguments as in the proof of Theorem 2, this yields $x^{\text{FB}} = \min\{\tilde{x}, \bar{T}\}$, where \tilde{x} is given by

$$u'\left(\frac{\tilde{x}}{n}\right) = c'\left(c^{-1}\left(u\left(\frac{\tilde{x}}{n}\right)\right)\right).$$

The resulting optimal effort level is

$$e^{\text{FB}} = c^{-1}\left(u\left(\frac{x^{\text{FB}}}{n}\right)\right).$$

Now suppose that the agents are risk-neutral, i.e., the function u is linear. The conditions characterizing (e^*, x^*) in Theorem 2 then coincide with those characterizing $(e^{\text{FB}}, x^{\text{FB}})$ above, which implies $(e^*, x^*) = (e^{\text{FB}}, x^{\text{FB}})$. Then suppose that the agents are risk-averse, i.e., the function u is strictly concave. If $x^* \neq x^{\text{FB}}$ there is nothing to prove. Hence assume $x^* = x^{\text{FB}}$. Inspection of the conditions that define e^* and e^{FB} then immediately reveals that $e^* < e^{\text{FB}}$.

A.4 Proof of Theorem 3

The proof proceeds in two steps. Step 1 shows that (e^*, \dots, e^*) is an equilibrium of the contest C_y^* described in the theorem. Step 2 shows that no other equilibria exist. The structure of the arguments in Step 2 is reminiscent of equilibrium characterization proofs in all-pay auctions with or without caps (e.g. Baye et al. 1996 or Che and Gale 1998). Throughout the proof, we adopt the interpretation of a nonphysical cap, so efforts $e_i > e^*$ are possible but are not differentiated by the reviewer. The result then also follows for the case where e^* is a physical bound on efforts.

Step 1. Consider deviations e'_i of agent i from (e^*, \dots, e^*) . If $e'_i > e^*$, we obtain

$$\begin{aligned} \Pi_i((e^*, \dots, e^*), C_y^*) &= \frac{n-1}{n} u\left(\frac{x^*}{n-1}\right) - c(e^*) \\ &> \frac{n-1}{n} u\left(\frac{x^*}{n-1}\right) - c(e'_i) \\ &= \Pi_i((e^*, \dots, e'_i, \dots, e^*), C_y^*). \end{aligned}$$

If $e'_i < e^*$, we obtain

$$\begin{aligned} \Pi_i((e^*, \dots, e^*), C_y^*) &= \frac{n-1}{n} u\left(\frac{x^*}{n-1}\right) - c(e^*) \\ &= 0 \geq -c(e'_i) = \Pi_i((e^*, \dots, e'_i, \dots, e^*), C_y^*). \end{aligned}$$

Thus, the contest C_y^* implements the effort profile (e^*, \dots, e^*) .

Step 2. By contradiction, suppose C_y^* also implements some other profile $\sigma \neq (e^*, \dots, e^*)$. Denote the support of σ_i by L_i , so $e_i \in L_i$ if and only if every open neighbourhood N of e_i satisfies $\sigma_i(N) > 0$. We first show that it must be that $L_i \subseteq [0, e^*]$ for all $i \in I$. Suppose not, so there exists an agent i and an effort level $e_i > e^*$ such that $\sigma_i((e_i - \epsilon, e_i + \epsilon)) > 0 \forall \epsilon > 0$. Fix $\bar{\epsilon} > 0$ such that $e_i - \bar{\epsilon} > e^*$. Note that the expected payoff of agent i playing $e'_i \geq e^*$ with probability one, while the other agents play σ_{-i} , is

$$\Pi_i(e'_i, \sigma_{-i}) = \left[1 - \prod_{j \neq i} \sigma_j([e^*, \infty)) + \prod_{j \neq i} \sigma_j([e^*, \infty)) \frac{n-1}{n} \right] u\left(\frac{x^*}{n-1}\right) - c(e'_i),$$

where we omit the dependence on C_y^* to simplify notation. Since c is strictly increasing, we have $\Pi_i(e^*, \sigma_{-i}) > \Pi_i(e'_i, \sigma_{-i})$ for all $e'_i > e^*$. Hence $\Pi_i(e^*, \sigma_{-i}) > \Pi_i(\bar{e}_i, \sigma_{-i})$ for all $\bar{e}_i \in (e_i - \bar{\epsilon}, e_i + \bar{\epsilon})$. Since $\sigma_i((e_i - \bar{\epsilon}, e_i + \bar{\epsilon})) > 0$, agent i could strictly increase his expected payoff by shifting the mass from this interval to e^* . Thus, σ is not an equilibrium. From now on, we only consider the cases where $L_i \subseteq [0, e^*] \forall i \in I$. Let $\underline{e}_i = \min L_i$. Since the proposed profile σ is different from (e^*, \dots, e^*) , it must be that $\underline{e} = \min_{i \in I} \underline{e}_i < e^*$.

First, suppose that $\underline{e} > 0$. Furthermore suppose that $\sigma_j(\{\underline{e}\}) > 0$ for exactly one agent $j \in I$, or that $\sigma_i(\{\underline{e}\}) = 0$ for all $i \in I$. In the latter case let j be such that $\underline{e}_j = \underline{e}$. Then there exists some $\bar{\epsilon} > 0$ such that

$$\Pi_j(\underline{e} + \epsilon, \sigma_{-j}) \leq \left[1 - \prod_{i \neq j} \sigma_i((\underline{e} + \epsilon, \infty)) \right] u\left(\frac{x^*}{n-1}\right) - c(\underline{e} + \epsilon) < 0$$

for all $\epsilon < \bar{\epsilon}$. Intuitively, the probability that agent j wins a positive prize approaches zero as ϵ approaches zero (by right continuity of $\sigma_i((\underline{e} + \epsilon, \infty))$ in ϵ and $\sigma_i(\underline{e}, \infty) = 1$), while the cost of effort at \underline{e} is strictly positive. Hence agent j could strictly increase his expected payoff by shifting the mass $\sigma_j([\underline{e}, \underline{e} + \bar{\epsilon})) > 0$ from $[\underline{e}, \underline{e} + \bar{\epsilon})$ to 0. Next suppose that $\sigma_i(\{\underline{e}\}) > 0$ for at least two agents $i = j, k$. Then there exists a small $\epsilon > 0$ such that

$$\Pi_j(\underline{e}, \sigma_{-j}) \leq \left[1 - \left(1 - \frac{1}{2} \sigma_k(\{\underline{e}\}) \right) \prod_{i \neq j, k} \sigma_i((\underline{e}, \infty)) \right] u\left(\frac{x^*}{n-1}\right) - c(\underline{e})$$

$$\begin{aligned} &< \left[1 - \prod_{i \neq j} \sigma_i((\underline{e}, \infty)) \right] u\left(\frac{x^*}{n-1}\right) - c(\underline{e} + \epsilon) \\ &\leq \Pi_j(\underline{e} + \epsilon, \sigma_{-j}). \end{aligned}$$

The intuition is that a small upward deviation from \underline{e} increases the probability of winning discretely, while marginally increasing the effort costs. Hence agent j could strictly increase his expected payoff by shifting the mass $\sigma_j(\{\underline{e}\}) > 0$ from \underline{e} to $\underline{e} + \epsilon$. We conclude that there does not exist an equilibrium $\sigma \neq (e^*, \dots, e^*)$ with $\underline{e} > 0$.

Second, suppose that $\underline{e} = 0$. Consider first the case where $\sigma_i(\{0\}) = 0$ for all $i \in I$, that is, no agent places an atom on 0. If there is an agent j such that $e_k > 0$ for all $k \neq j$, then there exists some $\bar{\epsilon} > 0$ such that $\Pi_j(\epsilon, \sigma_{-j}) = -c(\epsilon)$ for all $\epsilon < \bar{\epsilon}$. Agent j could then strictly increase his expected payoff by shifting the mass $\sigma_j((0, \bar{\epsilon})) > 0$ from $(0, \bar{\epsilon})$ to 0. Thus, there have to be at least two agents j and k with $\underline{e}_j = \underline{e}_k = 0$. But in this case, observe that

$$\Pi_j(\epsilon, \sigma_{-j}) \leq \left[\left(1 - \prod_{i \neq j} \sigma_i((\epsilon, \infty)) \right) u\left(\frac{x^*}{n-1}\right) - c(\epsilon) \right]$$

and

$$\lim_{\epsilon \rightarrow 0} \left[\left(1 - \prod_{i \neq j} \sigma_i((\epsilon, \infty)) \right) u\left(\frac{x^*}{n-1}\right) - c(\epsilon) \right] = 0.$$

Thus for every $\bar{\Pi} > 0$ there exists $\bar{\epsilon} > 0$ such that $\Pi_j(\epsilon, \sigma_{-j}) < \bar{\Pi}$ for all $\epsilon < \bar{\epsilon}$. Intuitively, both the probability of winning and the costs approach zero as $\epsilon \rightarrow 0$. However, it must be that $\Pi_j(e^*, \sigma_{-j}) > 0$ since $\Pi_j(e^*, \dots, e^*) = 0$ and the probability that j wins a positive prize is strictly greater if the other agents play σ_{-j} , because at least agent k exerts efforts lower than e^* with strictly positive probability. Hence agent j could strictly increase his expected payoff by shifting the mass $\sigma_j((0, \bar{\epsilon})) > 0$ from $(0, \bar{\epsilon})$ to e^* , for some sufficiently small $\bar{\epsilon} > 0$. The only remaining case is $\sigma_j(\{0\}) > 0$ for at least one agent $j \in I$. Observe that there can only be one such agent, since otherwise a small upward deviation from 0 would lead to a discrete increase in the probability of winning a positive prize, analogous to the argument above. Then it must be that $\Pi_j(\sigma_j, \sigma_{-j}) = 0$ since $\Pi_j(0, \sigma_{-j}) = 0$. This can only be the maximum payoff of agent j if all other agents exert deterministic efforts equal to e^* , since otherwise $\Pi_j(e^*, \sigma_{-j}) > 0$. In this case, agent j is indifferent between playing 0 or e^* , and all other effort levels yield strictly lower payoffs. This implies $\sigma_j(\{0\}) + \sigma_j(\{e^*\}) = 1$. Now consider an agent $k \neq j$. Observe that a deviation by agent k to some ϵ with $0 < \epsilon < e^*$ leads to payoffs

$$\Pi_k(\epsilon, \sigma_{-k}) = \sigma_j(\{0\}) u\left(\frac{x^*}{n-1}\right) - c(\epsilon).$$

Thus a sufficiently small $\epsilon > 0$ will be a profitable deviation whenever

$$\sigma_j(\{0\}) u\left(\frac{x^*}{n-1}\right) > \sigma_j(\{0\}) u\left(\frac{x^*}{n-1}\right) + (1 - \sigma_j(\{0\})) \frac{n-1}{n} u\left(\frac{x^*}{n-1}\right) - c(e^*).$$

This can be reformulated to

$$0 > -\sigma_j(\{0\}) \frac{n-1}{n} u\left(\frac{x^*}{n-1}\right),$$

which always holds because $\sigma_j(\{0\}) > 0$.

REFERENCES

- Aghion, Philippe and Jean Tirole (1997), “Formal and real authority in organizations.” *Journal of Political Economy*, 105, 1–29. [478]
- Alonso, Ricardo and Niko Matouschek (2008), “Optimal delegation.” *Review of Economic Studies*, 75, 259–293. [481]
- Amador, Manuel and Kyle Bagwell (2013), “The theory of optimal delegation with an application to tariff caps.” *Econometrica*, 81, 1541–1599. [481]
- Armstrong, Mark and John Vickers (2010), “A model of delegated project choice.” *Econometrica*, 78, 213–244. [481]
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul (2005), “Social preferences and the response to incentives: Evidence from personnel data.” *Quarterly Journal of Economics*, 120, 917–962. [494]
- Baye, Michael R., Dan Kovenock, and Casper G. De Vries (1996), “The all-pay auction with complete information.” *Economic Theory*, 8, 291–305. [503]
- Berger, Johannes, Christine Harbring, and Dirk Sliwka (2013), “Performance appraisals and the impact of forced distribution—an experimental investigation.” *Management Science*, 59, 54–68. [478, 494]
- Bernardin, H. John, Donna K. Cooke, and Peter Villanova (2000), “Conscientiousness and agreeableness as predictors of rating leniency.” *Journal of Applied Psychology*, 85, 232–236. [478]
- Bol, Jasmijn C. (2011), “The determinants and performance effects of managers’ performance evaluation biases.” *Accounting Review*, 86, 1549–1575. [478]
- Breuer, Kathrin, Petra Nieken, and Dirk Sliwka (2013), “Social ties and subjective performance evaluations: An empirical investigation.” *Review of Managerial Science*, 7, 141–157. [478]
- Charkraborty, Archishman and Richard Harbaugh (2007), “Comparative cheap talk.” *Journal of Economic Theory*, 132, 70–94. [481]
- Che, Yeon-Koo and Ian L. Gale (1998), “Caps on political lobbying.” *American Economic Review*, 88, 643–651. [491, 492, 503]
- Fairburn, James A. and James M. Malcomson (2001), “Performance, promotion, and the Peter principle.” *Review of Economic Studies*, 68, 45–66. [494]

- Fang, Dawei, Thomas Noe, and Philipp Strack (forthcoming), "Turning up the heat: The discouraging effect of competition in contests." *Journal of Political Economy*. [490]
- Faure-Grimaud, Antoine, Jean-Jacques Laffont, and David Martimort (2003), "Collusion, delegation and supervision with soft information." *Review of Economic Studies*, 70, 253–279. [495]
- Frankel, Alexander (2014), "Aligned delegation." *American Economic Review*, 104, 66–83. [481, 492]
- Fuchs, William (2015), "Subjective evaluations: Discretionary bonuses and feedback credibility." *American Economic Journal: Microeconomics*, 7, 99–108. [486]
- Gavious, Arieh, Benny Moldovanu, and Aner Sela (2002), "Bid costs and endogenous bid caps." *RAND Journal of Economics*, 33, 709–722. [491]
- Giebe, Thomas and Oliver Gürtler (2012), "Optimal contracts for lenient supervisors." *Journal of Economic Behavior & Organization*, 81, 403–420. [479]
- Glazer, Amihai and Refael Hassin (1988), "Optimal contests." *Economic Inquiry*, 26, 133–143. [480, 490, 491]
- Goldin, Claudia and Cecilia Rouse (2000), "Orchestrating impartiality: The impact of "blind" auditions on female musicians." *American Economic Review*, 90, 715–741. [494]
- Green, Jerry R. and Nancy L. Stokey (1983), "A comparison of tournaments and contracts." *Journal of Political Economy*, 91, 349–364. [480, 491]
- Holmström, Bengt (1977), *On Incentives and Control in Organizations*. Ph.D. thesis, Stanford University. [481]
- Holmström, Bengt (1984), "On the theory of delegation." In *Bayesian Models in Economic Theory* (Marcel Boyer and Richard Kihlstrom, eds.), 115–141, North-Holland, Amsterdam, Netherlands. [481]
- Jia, Hao, Stergios Skaperdas, and Samarth Vaidya (2013), "Contest functions: Theoretical foundations and issues in estimation." *International Journal of Industrial Organization*, 31, 211–222. [480]
- Kellner, Christian (2015), "Tournaments as a response to ambiguity aversion in incentive contracts." *Journal of Economic Theory*, 159, 627–655. [480]
- Konrad, Kai A. (2001), "Privacy and time-consistent optimal labor income taxation." *Journal of Public Economics*, 79, 503–519. [479]
- Konrad, Kai A. (2009), *Strategy and Dynamics in Contests*. Oxford University Press, New York, New York. [491]
- Krähmer, Daniel and Eugen Kováč (2016), "Optimal sequential delegation." *Journal of Economic Theory*, 163, 849–888. [479]
- Kwoh, Leslie (2012, January 31), "'Rank and yank' retains vocal fans." *The Wall Street Journal*. <http://goo.gl/2UErjD>. [481]

- Laffont, Jean-Jacques and David Martimort (1997), "Collusion under asymmetric information." *Econometrica*, 65, 875–911. [494]
- Laffont, Jean-Jacques and David Martimort (1998), "Collusion and delegation." *RAND Journal of Economics*, 29, 280–305. [495]
- Laffont, Jean-Jacques and David Martimort (1999), "Separation of regulators against collusive behavior." *RAND Journal of Economics*, 30, 232–262. [495]
- Lazear, Edward P. and Sherwin Rosen (1981), "Rank-order tournaments as optimum labor contracts." *Journal of Political Economy*, 89, 841–864. [480]
- Letina, Igor, Shuo Liu, and Nick Netzer (2018), "Delegating performance evaluation." Working Paper No. 266, University of Zurich, Department of Economics. [494]
- Levin, Jonathan (2002), "Multilateral contracting and the employment relationship." *Quarterly Journal of Economics*, 117, 1075–1103. [479]
- Liu, Shuo and Dimitri Migrow (2019), *Designing Organizations in Volatile Markets*. Working Paper No. 319, University of Zurich, Department of Economics. [478]
- MacLeod, Bentley W. (2003), "Optimal contracting with subjective evaluation." *American Economic Review*, 93, 216–240. [486]
- Malcomson, James M. (1984), "Work incentives, hierarchy, and internal labor markets." *Journal of Political Economy*, 92, 486–507. [479]
- Malcomson, James M. (1986), "Rank-order contracts for a principal with many agents." *Review of Economic Studies*, 53, 807–817. [479]
- McGregor, Jena (2013, November 20), "For whom the bell curve tolls." *The Washington Post*. <https://goo.gl/CwdpTO>. [481, 482]
- Melumad, Nahum and Toshiyuki Shibano (1991), "Communication in settings with no transfers." *RAND Journal of Economics*, 22, 173–198. [481]
- Milgrom, Paul R. and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601. [495]
- Moldovanu, Benny and Aner Sela (2001), "The optimal allocation of prizes in contests." *American Economic Review*, 91, 542–558. [480]
- Nalebuff, Barry J. and Joseph E. Stiglitz (1983), "Prizes and incentives: Towards a general theory of compensation and competition." *Bell Journal of Economics*, 14, 21–43. [480, 490, 491]
- Netzer, Nick and Florian Scheuer (2010), "Competitive markets without commitment." *Journal of Political Economy*, 118, 1079–1109. [479, 480]
- Olszewski, Wojciech and Ron Siegel (2019), "Bid caps in large contests." *Games and Economic Behavior*, 115, 101–112. [492]
- Pei, Harry Di (2015), "Communication with endogenous information acquisition." *Journal of Economic Theory*, 160, 132–149. [478]

- Prendergast, Canice (1999), “The provision of incentives in firms.” *Journal of Economic Literature*, 37, 7–63. [478]
- Prendergast, Canice and Robert H. Topel (1996), “Favoritism in organizations.” *Journal of Political Economy*, 104, 958–978. [478]
- Rahman, David (2012), “But who will monitor the monitor?” *American Economic Review*, 102, 2767–2797. [478]
- Rajan, Madhav V. and Stefan Reichelstein (2006), “Subjective performance indicators and discretionary bonus pools.” *Journal of Accounting Research*, 44, 585–618. [486]
- Siegel, Ron (2009), “All-pay contests.” *Econometrica*, 77, 71–92. [481]
- Strausz, Roland (1997), “Delegation of monitoring in a principal-agent relationship.” *Review of Economic Studies*, 64, 337–357. [478]
- Svensson, Jakob (2003), “Why conditional aid does not work and what can be done about it?” *Journal of Development Economics*, 70, 381–402. [479]
- Szalay, Dezső (2005), “The economics of clear advice and extreme options.” *Review of Economic Studies*, 72, 1173–1198. [478]
- Tanner, Noam (2013), “Optimal delegation under uncertain bias.” Unpublished paper, Yale University, Department of Economics. [479]

Co-editor Thomas Mariotti handled this manuscript.

Manuscript received 18 December, 2018; final version accepted 8 October, 2019; available online 14 October, 2019.