American Economic Review 2021, 111(3): 757–786
https://doi.org/10.1257/aer.20200555

Rank Uncertainty in Organizations†

By Marina Halac, Elliot Lipnowski, and Daniel Rappoport*

A principal incentivizes a team of agents to work by privately offering them bonuses contingent on team success. We study the principal's optimal incentive scheme that implements work as a unique equilibrium. This scheme leverages rank uncertainty to address strategic uncertainty. Each agent is informed only of a ranking distribution and his own bonus, the latter making work dominant provided that higher-rank agents work. If agents are symmetric, their bonuses are identical. Thus, discrimination is strictly suboptimal, in sharp contrast with the case of public contracts (Winter 2004). We characterize how agents' ranking and compensation vary with asymmetric effort costs. (*JEL* D23, D62, D81, D82, D86)

Project success in organizations relies on the contribution of multiple workers whose tasks are complementary. Workers who are rewarded based on overall project outcomes may thus be reluctant to do their share unless they expect that others will as well. We study the optimal provision of incentives that addresses this strategic risk. A firm seeking to ensure effort at the lowest cost must not only offer high rewards to make effort worthwhile, but also fine-tune rewards to keep workers optimistic that others will also want to work toward their common goal. How does the firm design incentives to manage workers' expectations and facilitate coordination? Is transparency about workers' rank and pay always good? Is pay inequality a feature of optimal incentive provision?

It is intuitive that transparency about coworker effort would help coordination, as workers could then condition their behavior on that of others.[1] Effort, however, is often not directly observable to the firm or other workers, and responding to others' behavior is infeasible when tasks must be undertaken in parallel. Yet, the firm can still provide assurance to a worker that his coworker will work, for example by offering the latter a high reward for project success that makes working a dominant

* Halac: Department of Economics, Yale University, and CEPR (email: marina.halac@yale.edu); Lipnowski: Department of Economics, Columbia University (email: e.lipnowski@columbia.edu); Rappoport: University of Chicago, Booth School of Business (email: daniel.rappoport@chicagobooth.edu). Jeffrey Ely was the coeditor for this article. We are grateful to Eyal Winter for his valuable feedback and to Ale Bonatti and Bob Gibbons for organizing the 2018 "Lab on Organizational Economics," where we started this project. We thank Heski Bar-Isaac, Navin Kartik, Ilan Kremer, Raphaël Levy, Teddy Mekonnen, Meg Meyer, Jacopo Perego, Ilya Segal, Raphaël Soubeyran, Asher Wolinsky, Takuro Yamashita, and various seminar and conference audiences for helpful comments.

† Go to https://doi.org/10.1257/aer.20200555 to visit the article page for additional materials and author disclosure statements.

[1] See Winter (2006, 2010) for theoretical analysis and discussions of the empirical evidence.

choice. In fact, based on this logic, Winter's (2004) seminal paper finds that a hierarchy of workers is optimal: high-rank workers are offered steep incentives to work no matter what others do, while low-rank workers are offered just enough to make them work knowing that higher-rank workers will work. Consistent with a broader literature on contracting with externalities (e.g., Segal 1999, 2003), Winter (2004) obtains that any optimal incentive scheme must be discriminatory, offering differential rewards even when workers are symmetric.

In this paper, we show that the conclusion above hinges on the presumption that contracts within an organization are public, so that workers' relative ranks are transparent. Such a presumption is at odds with evidence and current debates on firm practices. Firms not only rarely disclose information about employees' contractual terms, but also tend to discourage or even prohibit employees from discussing this information with each other (e.g., Gely and Bierman 2003; Hegewisch, Williams, and Drago 2011). Thus, there is often little if any transparency, and this secrecy inside firms is further sustained by social norms (e.g., Edwards 2005, Cullen and Perez-Truglia 2018). Taking the canonical team setting of Winter (2004), we show that limiting workers' information about their position relative to others is indeed beneficial to a firm. With private contract offers, the firm leverages rank uncertainty to better address strategic uncertainty and ensure effort at a lower cost. Furthermore, we prove that the firm's optimal incentive scheme is unique and does not discriminate between symmetric workers. That is, under private contracting, discrimination is not only unnecessary but strictly suboptimal for the firm.

Our model is cast in a moral-hazard-in-teams setting (Holmström 1982) with effort complementarities. A principal incentivizes a number of agents to work toward a common project. Each agent privately chooses whether to work or shirk, where working is costly and increases the likelihood of project success at a rate that is increasing with other agents' work. The principal can commit to make bonus payments to the agents conditional on the project succeeding, with the bonus offers being private. The principal's goal is to uniquely implement work at the least total cost, that is, to specify a least-cost incentive scheme such that all agents working is the unique Bayesian Nash equilibrium.

We begin by showing that an optimal incentive scheme specifies ranks as in Winter (2004). However, a key difference is that the ranking can be uncertain. Formally, the principal uses a ranking scheme, specifying a distribution of type profiles and informing each agent only of his own realized type. Given his belief over other agents' types, each agent type is then offered the minimum bonus for success that makes him want to work under the hypothesis that agents of lower type work and the rest do not. Since agents' efforts are strategic complements, an agent who believes himself to have a low type relative to others must be offered a relatively high bonus. Thus, a relatively low type corresponds to being assigned a relatively high rank.

Using the structure of ranking schemes, we characterize the principal's optimal value in Theorem 1. The principal's problem can be viewed as choosing an average ranking distribution plus an information structure, the latter determining what each agent learns about the realized ranking from his own type. Theorem 1 shows that providing agents with any private information about the ranking is suboptimal, and therefore the principal's problem reduces to only optimizing over the average ranking distribution. To prove this result, we show that the minimum bonus payment required for an agent to work when only lower types work is convex in the agent's belief about the ranking. Hence, it is optimal to induce the same belief for all types of a given agent.

In Theorem 2, we show that our characterization of the principal's optimal value permits a characterization of optimal incentive schemes. Specifically, we show that there is a unique profile of optimal agent bonuses that solve the principal's problem in Theorem 1, and any optimal sequence of incentive schemes must induce a bonus distribution that converges to a degenerate distribution on this optimal bonus profile. Therefore, optimal incentive contracts are (approximately) unique, with agents' bonus payments being a continuous function of the parameters of the model.

An important corollary of Theorem 2 concerns the case in which agents are symmetric. The principal in this case finds it optimal to induce uniform beliefs over all possible agent rankings, and the unique optimal limiting bonuses are equal across the agents. The intuition turns on how the principal provides assurance to the agents that other agents will work. Making a worker more optimistic that his coworkers will work allows the principal to incentivize his effort at a lower cost; however, this requires that some coworker become more pessimistic, with his incentives then more sensitive to his beliefs about others. Thus, to minimize bonus payments, the principal builds assurance between the workers in a mutual way. The optimal incentive scheme makes symmetric workers hold the same belief about their rank, and therefore demand the same bonus payment to work.

These results bring a new perspective on the current debate about firm compensation practices. Concerns about pay discrimination have motivated recent regulation to improve transparency inside firms (Trotter, Zacur, and Stickney 2017), and have led a number of companies to announce more open internal disclosure policies.[2] Our analysis shows that, from the view of optimal incentive provision, lack of transparency does not generate discriminatory incentives. In fact, together with the results of Winter (2004), we obtain precisely the opposite: a firm's optimal incentive scheme is discriminatory if and only if contracts are required to be public. As a practical implication, this suggests that either measures aimed at improving transparency may be counterproductive at reducing discrimination, or factors others than optimal incentive provision are behind firms' discriminatory practices.[3]

We characterize the optimal degree of rank uncertainty and derive comparative statics of optimal contracts for agents who may be asymmetric in their costs of effort. We show that for any set of agents who are sufficiently symmetric (in a sense that we make precise), the principal strictly benefits from making them uncertain about their relative ranks.[4] An agent is ranked lower than another agent with certainty if and only if his effort cost is comparatively higher enough, in which case the principal seeks to assure him that the lower-cost agent will work. Regarding the optimal bonus profile, we find that higher-cost agents are offered larger bonuses.

[2] See Sue Shellenbarger, "Open Salaries: The Good, the Bad and the Awkward," *Wall Street Journal*, January 12, 2016.

[3] In particular, transparency might help address *taste-based* discrimination; see, for example, Bennedsen et al. (forthcoming).

[4] Moreover, the principal finds it optimal to make this uncertainty complete, so that no ranking among the agents can be ruled out. Importantly, recall that the agents face this uncertainty both ex ante and interim, as they learn minimal information about others from their own contract offers.

However, since the principal also tailors ranking beliefs in order to relax the more demanding incentive constraints of higher-cost agents, she compensates lower-cost agents with higher markups, namely higher bonuses per effort cost. Overall, our model predicts an organizational hierarchy in which workers of similar skill are assigned to the same level and only those of sufficiently higher skill to higher levels, and in which higher skill is rewarded with higher rents.

We close the paper by discussing a number of model extensions. In one such extension, we address a key element of our model, that the principal can withhold information from one agent about another's contract terms. We show that if agents could voluntarily, verifiably disclose their contract terms to other agents, then the principal would be unable to leverage rank uncertainty and outperform the public-contracts benchmark. As such, our results can help understand why firms often maintain strict pay secrecy policies. As noted above, many private sector employers in the United States prohibit the discussion of salary information, and workers report that discussing this information can lead to punishment: see Gely and Bierman (2003); Hegewisch, Williams, and Drago (2011); and Rosenfeld (2017), among others. One potential explanation why firms maintain these policies is that they put workers at an informational disadvantage when negotiating contract terms. But workers in many cases do not have the bargaining power to exploit information about others' contracts, and when they do, it is unclear that transparency would help them (cf. Cullen and Pakzad-Hurson 2020). We study a model in which the firm has all the bargaining power and show that limiting workers' information is beneficial to the firm even in such a case. Pay secrecy benefits the firm not by improving its bargaining position, but by improving its ability to address strategic risk in incentive provision.

Related Literature.—Our paper belongs to the literature on contracting with externalities in multi-agent settings, pioneered by Segal (1999, 2003). Most of this literature, including Segal (2003); Winter (2004); Bernstein and Winter (2012); and Halac, Kremer, and Winter (2020), studies optimal unique-implementation mechanisms as we do. Except for Winter (2004), however, these papers focus on settings in which agents' choices are observable and bilaterally contractible. Moreover, all of this work requires that contracts be public, thus abstracting from the possibility of information design.[5]

Winter (2004) provides a benchmark for our study by analyzing how to uniquely induce a team of agents to work when only the team's overall success is observable and contractible. Restricting attention to publicly observed offers, he finds that any optimal incentive scheme must be discriminatory if agents' efforts are strategic complements. This result is related to the optimality of "divide and conquer" strategies in Segal (2003). A key lesson from our analysis is that this result is overturned when contract offers can be private.[6]

[5] Segal (1999, Section IV) also considers privately observed contract offers but in a setting in which the principal is unable to commit to a scheme. As Segal (1999) studies equilibria in pure strategies, his agents face no on-path uncertainty about others' offers.

[6] While the substantive message is overturned, the mechanics are not. One can view our principal's optimal incentive scheme as a divide and conquer strategy that proceeds "type by type" as opposed to "agent by agent."

Our point that introducing uncertainty can reduce a principal's cost of ensuring effort is related to results in Eliaz and Spiegler (2015) and Moriya and Yamashita (2020). Both of these articles, however, feature different constraints on the space of contracts. In the context of Winter (2004), the analysis in Eliaz and Spiegler (2015) amounts to letting the principal randomize over each agent's bonus offer, but, crucially, with bonus distributions that are required to be independent and identical across the (symmetric) agents. Moriya and Yamashita (2020) extend Winter's (2004) setting by letting effort incentives depend on the realization of an exogenous binary state. Taking bonus offers to be deterministic (hence public) and symmetric, they show that the principal's optimal scheme may inform agents asymmetrically about the realized state.

More broadly, our paper is related to recent work studying multi-agent persuasion under adversarial equilibrium selection.[7] Inostroza and Pavan (2020) characterize optimal disclosure policies in a global game, where discriminatory disclosures can be beneficial by a similar "divide and conquer" logic as in the literature above. Mathevet, Perego, and Taneva (2020) examine the distributions of agents' beliefs that an information designer can induce, providing a representation for optimal distributions. The analysis of Oyama and Takahashi (2020) yields conditions for unique implementation in binary-action supermodular games of complete information, using contagion arguments (as in Rubinstein 1989, Carlsson and van Damme 1993, Kajii and Morris 1997, and many others) similar to those in our Lemma 1, and Morris, Oyama, and Takahashi (2020) address information design for such games under incomplete information.[8] While our paper differs in many aspects, our main departure from this literature is that we consider the joint design of both transfers and information about said transfers.

There is recent work studying, from different angles, pay transparency and discrimination in firms. Cullen and Pakzad-Hurson (2020) analyze the effects of increasing transparency on wage bargaining. Both theoretically and empirically, they find that higher transparency can allow workers to renegotiate to a common wage, while also reducing average wages as firms bargain more aggressively. Habibi (2020) uses a signaling model to study a firm's decision to be transparent, finding that either decision may be optimal depending on the firm's technology.

Finally, there is a literature on organizational hierarchies. One strand of this literature focuses on the role of hierarchy in the processing and aggregation of information and knowledge (see the survey by Garicano and Van Zandt 2013). Our paper is more closely related to another strand that examines the incentive function of hierarchy (see the survey by Mookherjee 2013), although the focus there is on monitoring and delegation of authority, from which we abstract. A hierarchy in our setting instead arises as a means of determining and (selectively) communicating the strength of workers' incentives.

[7] Our principal's goal of inducing all agents to work as the unique equilibrium is equivalent to inducing all agents to work under adversarial selection, i.e., in the lowest-effort equilibrium of the induced game.

[8] The latter generalizes the results of Moriya and Yamashita (2020) in an application. See also Hoshino (2020), which relates persuasion under adversarial selection to risk dominance, and Doval and Ely (2020), which characterizes, for given actions and payoffs, all the equilibrium outcomes that are consistent with some choice of information structure and extensive form.

I. Simple Example

This section explains the intuition for our main results in a simplified example. Consider a firm with two workers and a production technology as in Winter (2004). Each worker performs a different task. If a worker works, he completes his task with probability 1; if instead he shirks, the probability that he completes his task is only $p \in (0,1)$. The firm's project succeeds if and only if both workers' tasks are completed. Each worker's cost of working is $c > 0$, and workers make their work decisions simultaneously.

The firm observes only whether the project as a whole succeeds and can commit to make bonus payments to the workers given success. The firm's bonus offers induce a game between the workers, where a worker's payoff is equal to his expected bonus payment, minus the cost c if he works. The firm's problem is to find a least-cost incentive scheme such that both workers working is the unique (Bayesian) Nash equilibrium of the induced game.

Suppose first that the firm commits to public bonus contracts. Note that conditional on the other worker working, a worker is willing to work so long as he is offered a bonus for success no smaller than b_L satisfying $b_L - c = pb_L$, i.e., $b_L = c/(1-p)$. However, if the other worker shirks, then the required bonus is no smaller than b_H satisfying $pb_H - c = p^2 b_H$, i.e., $b_H = c/[p(1-p)] > b_L$. This implies that, while offering bonuses $b_1 = b_2 = b_L$ would suffice to induce an equilibrium in which both workers work, these (or slightly higher) bonuses would not exclude an equilibrium in which both workers shirk. To uniquely implement work, the firm must make it dominant for one of the workers to work, namely offer a bonus above b_H to one of the workers. A bonus above b_L for the other worker then ensures that he also works, knowing that the first worker always works. It follows that a least-cost public scheme that guarantees work entails a cost (slightly above) $b_H + b_L$ to the firm. As stressed by Winter (2004), such a scheme is discriminatory, as workers are treated asymmetrically despite them being identical.

It turns out that the firm can do better by introducing uncertainty. For example, consider a scheme in which worker 1 may be offered a high bonus (slightly above) b_H or a low bonus (slightly above) b_L, each with equal probability. Worker 1 is informed of his bonus before choosing whether to work, but worker 2 only knows worker 1's bonus distribution. Note that if worker 1 is offered b_H, then it is dominant for him to work. Since worker 2 knows that this happens with probability $1/2$, the firm can ensure that worker 2 works by offering him a bonus (slightly above) b_M satisfying

$$\left(\frac{1}{2} + \frac{1}{2}p\right) b_M - c = \left(\frac{1}{2}p + \frac{1}{2}p^2\right) b_M,$$

that is,

$$b_M = \frac{c}{\frac{1}{2}(1-p) + \frac{1}{2}p(1-p)}.$$

Moreover, since worker 2 works under bonus b_M, it is then optimal for worker 1 to also work when he is offered bonus b_L. It follows that the firm guarantees work with this scheme at a cost (slightly above) $(b_H + b_L)/2 + b_M$. Since $b_M < (b_H + b_L)/2$, the firm's cost is lower than the cost computed above under public offers. Furthermore, since $b_L < b_M < b_H$, the firm's scheme is less discriminatory compared to that above, not only ex ante but also ex post.

Intuitively, given effort complementarities, the firm wishes to provide assurance to the workers that their coworker will work. A public scheme favors one of the workers in order to guarantee his effort and provide assurance to the other worker. By creating rank uncertainty, however, the firm can build this assurance in a mutual way, so that both workers provide assurance to each other.[9] This not only makes the workers' contracts less asymmetric, but also allows the firm to ensure effort at a lower cost. The reason is that a worker's incentive to work increases at a decreasing rate with his belief that the other worker works; hence, assurance is most valuable when a worker is pessimistic about his coworker's effort.

While the scheme described above improves upon the public-contracts scheme and illustrates our main points, it is not an optimal scheme in this setting. A simple way to further decrease the firm's cost is to add another bonus realization for worker 2 so that worker 1 is also uncertain about his relative rank. More specifically, consider a uniform distribution over bonus profiles (slightly above) (b_H, b_M), (b_M, b_M), and (b_M, b_L). By analogous reasoning to that above, proceeding one bonus realization at a time, one can verify that both workers working is the unique equilibrium. Moreover, the cost to the firm is $2[(2/3)b_M + (1/3)(b_H + b_L)/2]$which is lower than the previous cost (again because $b_M < (b_H + b_L)/2$). We return to this example and expand on this idea when describing an optimal scheme in Subsection IIIB.

This discussion serves as a prelude for our analysis in the next sections. Our characterization of the firm's optimal incentive scheme will show more generally how creating rank uncertainty can be beneficial to the firm, and how offering discriminatory incentives becomes detrimental.

II. Model

We study a principal's design of incentives for N agents tasked with working on a common project. With some abuse of notation, let the set of agents be $N = \{1, \ldots, N\}$. Each agent $i \in N$ can either work at cost $c_i > 0$ or shirk. Depending on the agents' work decisions, the project can either succeed or fail: the production technology is given by $P : 2^N \rightarrow [0,1]$, where $P(J)$ is the probability that the project succeeds conditional on all agents $i \in J$ working and all others shirking. We make the following assumption on P.

ASSUMPTION 1: *P is strictly increasing and strictly supermodular*, i.e., *for any J, $J' \subseteq N$:*

(i) $P(J) > P(J')$ *if* $J \supsetneq J'$;

(ii) $P(J \cup J') - P(J) > P(J') - P(J \cap J')$ *if* J, J' *are not nested.*

[9] An analogous form of mutual assurance arises in the literature on virtual implementation. For instance, the analysis of Abreu and Matsushima (1992) ensures truthful implementation by breaking a collective decision into a product of many distinct decisions which are determined at various stages of a deletion sequence. Similarly, our principal ensures work by breaking the agents' effort decisions into those of distinct types whose choices are determined at various stages of a deletion sequence.

Part (i) says that effort is productive, that is, the project succeeds with higher probability if more agents choose to work rather than shirk. Part (ii) says that efforts are complementary, that is, an agent's marginal product (the marginal effect of his effort on the probability of project success) is larger when more other agents are working.

Agents' work choices are private and only the project outcome (whether the project succeeds or fails) is contractible. Additionally, agents are protected by limited liability, requiring any payments from the principal to be non-negative. Without loss, the principal therefore only offers the agents success-contingent bonuses, with project failure resulting in no payment to any agent. We allow the principal's bonus offers to be private. Formally, an incentive scheme is $\sigma = \langle q, B \rangle$, where $q \in \Delta(\mathbb{N}^N)$ is a finite-support prior and $B = (B_1, \ldots, B_N)$ is a bonus rule, with $B_i : \mathrm{supp}(q_i) \to \mathbb{R}_+$. For convenience, let $T_i^q := \mathrm{supp}(q_i)$ denote the support of the marginal of q along dimension i and $T^q := \prod_{i \in N} T_i^q$. The interpretation is that the principal privately informs each agent $i \in N$ of his *type* $t_i \in T_i^q$ and, through the bonus rule, of his success-contingent bonus $B_i(t_i)$.[10] Hence, an agent may face uncertainty about the contracts of other agents but is completely informed about his own.[11] The extent of the uncertainty depends on the principal's choice of q, specifically the correlation between agents' types. For example, agents face no uncertainty about other agents' terms under public contracting, i.e., if the prior q satisfies $\left|\{t_{-i} : q(t_i, t_{-i}) > 0\}\right| = 1$ for every agent $i \in N$ and type $t_i \in T_i^q$.

An incentive scheme $\sigma = \langle q, B \rangle$ defines a Bayesian game between the agents. In this game, $\langle (T_i^q)_{i \in N}, q \rangle$ is a common-prior type space; each agent simultaneously makes a type-contingent decision of whether to work or shirk; and an agent's payoff is equal to his expected bonus payment net of any cost of effort he bears. An equilibrium induced by an incentive scheme is a Bayesian Nash equilibrium of the Bayesian game defined by the scheme. The principal's goal is to offer a least-cost incentive scheme that ensures work, namely that induces all types of all agents to work as the unique equilibrium. For expositional convenience when dealing with indifferent agents, we require schemes to induce such a unique equilibrium only once the bonus offers are increased by any positive amount. More precisely, say that $\sigma = \langle q, B \rangle$ *uniquely implements work* (*UIW*) if, for every $\varepsilon > 0$, all agent types working is the unique equilibrium of the game induced by scheme $\langle q, B + \varepsilon \rangle$ (in which each type t_i of agent i is offered bonus $B_i(t_i) + \varepsilon$).[12] The principal then solves the following problem:

$$\inf_{\sigma \text{ } UIW} W(\sigma), \tag{1}$$

[10] An essentially equivalent formulation would have the principal choose a distribution over bonus profiles directly. Our present formulation highlights that the principal separately chooses an agent's bonus and his information about others' realized terms via his type. We restrict attention to finite type spaces for notational simplicity, but the analysis is substantively the same, and the principal's optimal value identical, for more general common-prior type spaces. Given the finite type restriction, since the type itself is simply a label, it is immaterial that agents' types are labeled with natural numbers.

[11] This realistic assumption is not without loss. In Section V, we show how results change without it, i.e., when agent i's bonus can be conditioned on the private contract of agent j.

[12] As will be clear, our characterization results in Theorem 1 and Theorem 2 would remain true as stated if ε were replaced with 0 in this definition. Our solution concept corresponds to that in Winter (2004) except that we have a Bayesian game as the principal's offers are private.

where $W(\sigma)$ is the principal's total expected cost given incentive scheme σ and all agents working:

$$W(\sigma) = P(N) \sum_{i \in N} \sum_{t_i \in T_i^q} q_i(t_i) B_i(t_i).$$

It will follow from our results that the principal's problem does not generally admit a minimum. Hence, we will focus on approximately optimal incentive schemes. Let W^* denote the principal's optimal value, i.e., the value of program (1). We define an *optimal sequence of incentive schemes* as a sequence $(\sigma^m)_m$ such that each σ^m uniquely implements work and $W(\sigma^m)$ converges to W^* as $m \to \infty$.

A feature of approximately optimal incentive schemes in which we are particularly interested is the distribution of bonuses offered to each agent. Given an incentive scheme $\sigma = \langle q, B \rangle$, let $\beta^\sigma \in \Delta(\mathbb{R}_+^N)$ be the distribution of bonus profiles, i.e., the finite-support distribution with[13]

$$\beta^\sigma(b) := q\{t \in T^q : B_1(t_1) = b_1, \ldots, B_N(t_N) = b_N\} \quad \text{for all } b \in \mathbb{R}_+^N.$$

Given a sequence $(\sigma^m)_m$ of incentive schemes, $\beta^* \in \Delta(\mathbb{R}_+^N)$ is the limit bonus distribution of this sequence if it is the limit of the sequence $(\beta^{\sigma^m})_m$.[14] Finally, say β^* is an optimal bonus distribution if it is the limit bonus distribution of an optimal sequence of incentive schemes.

Remark 1: We focus our exposition around the principal inducing every agent to work as the unique equilibrium. In our setting, this requirement is equivalent to inducing every agent to work as the unique rationalizable outcome. This follows from the fact that, given the supermodular technology P and the agents' limited liability, the game played by the agents under any incentive scheme is a supermodular game, and thus the results of Milgrom and Roberts (1990) apply.[15] In particular, the strategy profile in which every agent type chooses his lowest (i.e., least-work) interim correlated rationalizable action is an equilibrium, and if this lowest rationalizable strategy profile does not consist of everybody working, then it cannot be the unique equilibrium for everybody to work.

III. Optimal Incentives

We characterize optimal incentives by solving the principal's problem in (1). We proceed as follows. First, we show that it is without loss for the principal to focus on a simple class of incentive schemes which we call ranking schemes. Second, using

[13]Note that using this definition, the principal's total expected cost in (1) can be written as $W(\sigma) = P(N) \sum_{i \in N} \int b_i \, d\beta_i^\sigma(b_i)$, where $\beta_i^\sigma \in \Delta(\mathbb{R}_+)$ is the marginal distribution over the bonus offered to agent i.

[14]Here, we take convergence in the weak* sense.

[15]Milgrom and Roberts (1990) study games of complete information. To apply their results to our setting, however, note that our game can be viewed in "agent form" by taking each of the finitely many types of each agent as a distinct player. Interim correlated rationalizable play in the induced Bayesian game is equivalent to correlated rationalizable play in the agent-form game, and the latter is a game of strategic complements.

the structure of ranking schemes, we derive an auxiliary optimization program that fully characterizes the principal's value of (1). Finally, we show that this auxiliary program also characterizes the solution to (1), by pinning down the limiting bonus distribution of an optimal sequence of incentive schemes. We discuss the implications of optimal bonuses for pay discrimination.

A. *Ranking Agents*

It will be useful for our analysis to define a class of incentive schemes that we call ranking schemes.

DEFINITION 1: *An incentive scheme* $\sigma = \langle q, B\rangle$ *is a* ranking scheme *if:*

(*i*) *No two agents are assigned the same type. That is, for all* $i, j \in N$,

$$q\{t \in T^q : t_i = t_j\} = 0 \quad \text{if } i \neq j.$$

(*ii*) *The bonus for each agent type makes him indifferent between working and shirking given the belief that all other agents of lower type work and the rest shirk. That is, for all* $i \in N$ *and all* $t_i \in T_i^q$,

$$B_i(t_i) \sum_{t_{-i}} q_i(t_{-i}|t_i)\left[P\{j \in N : t_j \leq t_i\} - P\{j \in N : t_j < t_i\}\right] = c_i.$$

Note that the bonuses in a ranking scheme $\sigma = \langle q, B\rangle$ are pinned down, given the prior q, by the indifference conditions in part (ii) of Definition 1. Thus, we will sometimes refer to a ranking scheme simply by a distribution q over type profiles (satisfying part (i) of Definition 1), with the understanding that the associated bonus rule is as specified in this definition.

Ranking schemes have a simple structure. Under any such scheme $\sigma = \langle q, B\rangle$, each type of every agent is indifferent over working if all (strictly) lower types of other agents work and the rest do not. This means that under $\langle q, B + \varepsilon\rangle$, for any $\varepsilon > 0$, each type strictly prefers to work if all lower types work and the rest do not. Moreover, by supermodularity of P, working is then dominant for each type given that *at least* all lower types work. It therefore follows, by induction on the type, that any ranking scheme uniquely implements work: under $\langle q, B + \varepsilon\rangle$, the lowest type finds work dominant; the second lowest type finds work dominant given that the lowest type works; the third lowest type finds work dominant given that the lowest and second lowest types work; and so on.

The next lemma shows that ranking schemes are useful not only because they ensure that all agents work, but because they can also be constructed to ensure work at the least total cost to the principal. Hence, to solve the principal's problem in (1), it is without loss to focus on ranking schemes.

LEMMA 1: *Every ranking scheme uniquely implements work. Moreover, if an incentive scheme* σ *uniquely implements work, then there is a ranking scheme* σ^* *such* *that, for each agent $i \in N$, the marginal bonus distribution β_i^{σ} first-order-stochastically dominates $\beta_i^{\sigma^*}$. Hence, the principal's optimal value satisfies*

$$\inf_{\sigma\ UIW} W(\sigma) = \inf_{\sigma \text{ is a ranking scheme}} W(\sigma).$$

For intuition, fix an incentive scheme $\sigma = \langle q, B \rangle$ that uniquely implements work. We can construct a ranking scheme that (weakly) improves the principal's objective in three steps. First, we iteratively assign each type of each agent an "index": the index-1 types are those who find work weakly dominant; the index-2 types are those who find work weakly dominant given that index-1 types work; the index-3 types are those who find work weakly dominant given that index-1 and index-2 types work; and so on. We observe that *every* type will be assigned such an index, for otherwise one can construct an equilibrium for $\langle q, B + \varepsilon \rangle$, with $\varepsilon > 0$ small enough, in which the unindexed types shirk. Second, we relabel types, so that no two agents ever have the same type and, moreover, $t_i < t_j$ whenever t_i has a strictly lower index than t_j. We can then assign bonuses to relabeled types as required by the definition of a ranking scheme, namely as specified in part (ii) of Definition 1. Third, we show that, by supermodularity of P, this adjustment to the bonuses systematically (weakly) lowers each type's bonus payment, thereby first-order-stochastically reducing the bonus distribution.[16] The resulting ranking scheme reduces the principal's total expected cost.

Agents' incentives to work in a ranking scheme depend on the probability that they assign to other agents having realized a lower type than their own. To make these incentives precise, we introduce the following notation. Let Π denote the set of all permutations on N, i.e., all $\pi \in N^N$ such that $\pi_i \neq \pi_j$ for distinct $i, j \in N$. For a type profile t with $t_i \neq t_j$ for all $i, j \in N$, let the *agent ranking* given t be $\pi(t) \in \Pi$ with $\pi_i(t) = |\{j \in N : t_j \leq t_i\}|$. Then given a ranking scheme q, agent $i \in N$, and type $t_i \in T_i^q$, let t_i's *ranking belief* $\mu_i^q(\cdot|t_i) \in \Delta\Pi$ be given by

$$\mu_i^q(\hat{\pi}|t_i) := q_i\left(\left\{t_{-i} : \pi(t_i, t_{-i}) = \hat{\pi}\right\}|t_i\right) \quad \text{for all } \hat{\pi} \in \Pi. \tag{2}$$

Finally, let the *ranking distribution* $\mu^q \in \Delta\Pi$ be given by

$$\mu^q(\hat{\pi}) := q\left(\left\{t : \pi(t) = \hat{\pi}\right\}\right) \quad \text{for all } \hat{\pi} \in \Pi.$$

For any agent $i \in N$ and ranking belief $\mu_i \in \Delta\Pi$ that he might hold, we can compute explicitly the bonus payment required by this agent to work. Specifically, if agent i believes that the set of other agents who work is $J \subseteq N \backslash \{i\}$ with probability $\mu_i\{\pi \in \Pi : \text{for all } j \in N, \pi_j < \pi_i \text{ if and only if } j \in J\}$, then the success-contingent bonus that keeps him exactly indifferent between working and shirking is equal to

$$\frac{c_i}{\sum_{\pi \in \Pi} \mu_i(\pi)\left[P\{j \in N : \pi_j \leq \pi_i\} - P\{j \in N : \pi_j < \pi_i\}\right]}.$$

[16] This reduction in bonuses for a given type could be due to slack in the original incentive scheme given that lower index types work, or because relabeling changes the relevant incentive constraint to one that assumes more other types are working.

Moreover, since the principal makes the bonus payments only when the project succeeds, and the probability of project success (given that work is uniquely implemented) is equal to $P(N)$, it follows that the principal's expected bonus payment to agent i, as a function of this agent's interim ranking belief μ_i, is

$$f_i(\mu_i) := \frac{c_i P(N)}{\sum_{\pi \in \Pi} \mu_i(\pi)\big[P\{j \in N : \pi_j \leq \pi_i\} - P\{j \in N : \pi_j < \pi_i\}\big]}. \tag{3}$$

The function $f_i : \Delta\Pi \to \mathbb{R}_{++}$ describes the expected fees associated with providing incentives to agent i for any given ranking belief that he might hold. Note that since P is supermodular, an agent who believes himself to be placed earlier (respectively, later) in the agent ranking on average will require higher (respectively, lower) bonuses to be willing to work. We thus say that agents who are placed earlier in the ranking are "higher-rank" agents.

B. *Principal's Value*

By Lemma 1, it is without loss for the principal to restrict attention to ranking schemes when solving for a least-cost incentive scheme that uniquely implements work. Moreover, we have shown that given a ranking scheme q, the bonus offer for type $t_i \in T_i^q$ of agent $i \in N$ can be computed as $(1/P(N)) f_i\big(\mu_i^q(\cdot|t_i)\big)$, where the function f_i is defined in (3) and $\mu_i^q(\cdot|t_i)$ is t_i's ranking belief. Hence, the principal's total expected cost under ranking scheme q can be written as

$$W(q) = \sum_{t \in T^q, i \in N} q(t) f_i\big(\mu_i^q(\cdot|t_i)\big). \tag{4}$$

It is clear from this expression that the relevant choice for the principal is a profile of distributions over ranking beliefs. However, the principal is constrained: she cannot make an agent believe that he is placed later in the agent ranking (thereby reducing his associated bonus) without making some other agent believe that he is placed earlier. To incorporate this constraint, consider the following interpretation of the principal's problem. The principal first chooses an average ranking distribution and then chooses an information structure for each agent about the realized ranking, with the agent's type corresponding to his signal. Every ranking scheme naturally implies some selection of ranking distribution and information structure, and it turns out that a certain converse also holds. Specifically, the following theorem shows that providing no private information about the ranking to any agent is optimal given a ranking distribution, and moreover that this can be done for any desired ranking distribution. Therefore, the principal's problem can be recast as an unconstrained optimization over ranking distributions themselves.

THEOREM 1: *The principal's optimal value satisfies*

$$\inf_{\sigma \text{ UIW}} W(\sigma) = \min_{\mu \in \Delta\Pi} \sum_{i \in N} f_i(\mu).$$

Theorem 1 reduces the principal's problem to choosing a ranking distribution, as opposed to choosing a profile of distributions over ranking beliefs. This means that it is optimal for the principal not to provide any information to the agents about the realized ranking. The reason lies in the fact that the minimum bonus payment required for an agent to work is convex in the agent's belief about the ranking. Intuitively, since P is supermodular, an agent type who believes himself to have a high rank (i.e., to be placed early in the agent ranking) will have weak incentives to work. This implies that the principal will have to offer him a large bonus, but also that she can obtain large savings from making this type believe that his rank is relatively lower. On the contrary, an agent type who believes himself to have a low rank will have strong incentives to work, so the principal can make him believe that his rank is relatively higher and maintain his incentives at a low cost. As a consequence, the principal will benefit from equalizing the beliefs of different types of the same agent.

While convexity implies that giving agents no private information about the ranking is optimal, it is a priori unclear whether some ranking scheme achieves this, and moreover whether in this way the principal can achieve any ranking distribution without restriction. The proof of Theorem 1 shows that this is indeed true. For any given desired distribution, we construct a sequence of ranking schemes that approximates this distribution and gives agents no information about their rank with probability approaching 1.[17]

To understand our construction, it is instructive to specialize it to the two-worker symmetric example of Section I. Recall that in that example, we considered an incentive scheme that would uniquely implement work by making worker 2 uncertain about his rank. Specifically, worker 2 did not know whether worker 1 was offered a high bonus b_H (making work dominant) or a low bonus b_L (making work dominant only conditional on worker 2 also working), and worker 2's bonus b_M would make him just willing to work for uniform beliefs about worker 1. Given the assumption that the project would succeed for sure under both workers working, the principal's value under this scheme was equal to $(b_H + b_L)/2 + b_M$. Theorem 1 however tells us that the principal can do better; in fact, her optimal value in this simple setting is equal to $2b_M$.

We can approximate the principal's optimal value by constructing a sequence of ranking schemes that makes both workers uncertain about their rank. For each $m \in \mathbb{N}$, let the principal (privately) draw a random variable ℓ which is uniform over $\{0, \ldots, m-1\}$. Independent of ℓ and uniformly across the workers, one worker is chosen to be the "leader" and the other to be the "follower." The leader is told that his type is $\ell + 1$ and the follower is told that his type is $\ell + 2$; crucially, neither is told whether he is the leader or the follower. Consider a worker's ranking belief given his type in $\{1, \ldots, m+1\}$: when his type is 1, he knows he is the leader; when his type is $m+1$, he knows he is the follower; and for every other type, he assigns probability $1/2$ to being the leader. Therefore, the principal can uniquely implement work by specifying bonus b_H for type 1, bonus b_L for type $m+1$, and bonus b_M for every other type. As m becomes large, this ranking scheme pays bonus b_M to each

[17] The type space induced by each such ranking scheme is essentially the same as one constructed in Oyama and Takahashi (2020). We thank an anonymous referee for bringing this to our attention.

worker with arbitrarily high probability, thus approximating the principal's optimal value in Theorem 1.[18]

C. *Bonus Distribution and (No-)Discrimination*

We have shown that the principal's optimal value can be characterized via the auxiliary optimization over ranking distributions given in Theorem 1. We now show that this auxiliary program also permits a characterization of optimal incentives. The following theorem fully characterizes optimal sequences of incentive schemes, by characterizing the limiting bonus distributions they generate.

THEOREM 2: *There is a unique bonus profile* $b^* \in \mathbb{R}^N$ *which minimizes* $\sum_{i \in N} b_i$ *among all*

$$b \in \left\{ \frac{1}{P(N)} \big(f_1(\mu), \ldots, f_N(\mu)\big) : \mu \in \Delta\Pi \right\}.$$

Furthermore, a sequence $(\sigma^m)_m$ *of incentive schemes that uniquely implement work is optimal if and only if, for every* $i \in N$, *the induced bonus distribution of agent i under* σ^m *converges to a degenerate distribution on* b_i^*.

The proof of Theorem 2 makes use of the fact that the function $f_i(\mu_i)$, which describes the cost of incentivizing agent $i \in N$ as a function of his ranking belief, is a strictly convex transformation of an affine, scalar-valued function. Thus, for every $i \in N$ and finite-support distribution $\tau_i \in \Delta\Delta\Pi$, we obtain $\int f_i \, d\tau_i \geq f_i\big(\int \mu_i \, d\tau_i(\mu_i)\big)$, with a strict inequality if τ_i induces a nondegenerate distribution of f_i. This allows us to show that $\sum_{i \in N} f_i(\mu)$ admits a unique minimizing payment profile $\big(f_i(\mu)\big)_{i \in N}$ over all $\mu \in \Delta\Pi$, and that, given an optimal sequence $(q^m)_m$ of ranking schemes, any limit point of the implied sequence $(\tau_i^m)_m$ induces a degenerate distribution of f_i for every $i \in N$. Combining these together with Lemma 1 and Theorem 1 yields Theorem 2.

The result in Theorem 2 shows that optimal incentive contracts are (approximately) unique, with agents' bonus payments being a continuous function of the parameters of the model.[19] It is worth noting that both of these features stand in contrast with the results of Winter (2004). In a setting in which the principal's contract offers are constrained to be public, Winter (2004) finds that the optimal incentive scheme may be arbitrary (for example, when all agents are symmetric), and so agents' payments may change discontinuously with the model parameters (for example, when agents become slightly asymmetric). Instead, allowing contract offers to be private delivers a unique and well-behaved solution.

Theorem 2 tells us that, for some parameter values, the principal benefits *strictly* from being able to contract privately with the agents. In particular, we will show in

[18] In this construction, only one low-probability type of each agent is paid enough to make work dominant. Thus, the principal's ability to offer different bonuses to different types, in contrast to the principal of Moriya and Yamashita (2020), is strictly valuable.

[19] Berge's theorem implies that optimal ranking distributions (as in Theorem 1) move upper hemicontinuously with the model's parameters. The uniqueness result of Theorem 2 then implies optimal bonuses are a continuous function of said parameters.

the next section that the unique optimal incentive scheme features rank uncertainty whenever the agents are not too asymmetric. Before turning to those results, we highlight here the implications of Theorem 2 for the case of symmetric agents. We say that agents $i,j \in N$ are *symmetric* if they have the same effort cost and marginal product, i.e., $c_i = c_j$ and $P(J \cup \{i\}) = P(J \cup \{j\})$ for each $J \subseteq N \backslash \{i,j\}$. The next corollary follows directly from Theorem 2.

COROLLARY 1: *Suppose agents $i,j \in N$ are symmetric. If a sequence $(\sigma^m)_m$ of incentives schemes that uniquely implement work is optimal, then the induced bonus distributions of agents i and j under σ^m converge to degenerate distributions on bonuses b_i^* and b_j^* with $b_i^* = b_j^*$.*

The principal's optimal incentive scheme treats symmetric agents symmetrically. Importantly, Corollary 1 says that pay discrimination is not just unnecessary but strictly suboptimal for the principal. Any two symmetric agents receive the same bonus offer, and they are thus also faced with the same uncertainty about their ranks. For example, if all the agents are symmetric, then an optimal incentive scheme induces uniform beliefs over all possible agent rankings, so that the limiting bonus is $b_i^* := (1/P(N))f_1(\mu_!)$ for all $i \in N$, where $\mu_! \in \Delta\Pi$ is uniform.

As discussed in the introduction, these findings provide a key counterpoint to Winter's (2004) seminal discrimination result. When the principal's contract offers are public, Winter (2004) finds that any optimal incentive scheme must treat all agents differently. Such a scheme discriminates between symmetric agents in an arbitrary way (at least ex post) and it treats nearly symmetric agents disproportionately differently (both ex ante and ex post). Intuitively, as in our setting, the principal uses a ranking scheme to uniquely induce the agents to work. However, under public contracts, the agent ranking must be transparent, and it must thus be accompanied by a hierarchy of bonuses: the first agent in the ranking is offered the highest bonus so that working is dominant; the second agent is offered the second highest bonus so that working is dominant knowing that the first agent works; and so on.

Theorem 2 and Corollary 1 show that this logic breaks down when the principal can contract privately with the agents. By announcing only a ranking distribution, the principal can make each agent assign positive probability to other agents working, thereby providing each of them with some assurance that others will work instead of building this assurance hierarchically as under public contracts above. We find that when agents are symmetric, the principal indeed minimizes bonus payments by building assurance symmetrically. The reason is that agents' incentives to work increase at a decreasing rate with their beliefs that other agents will work, so the principal wishes to increase the beliefs of precisely those agents who are most pessimistic. As a consequence, the optimal incentive scheme makes symmetric agents hold the same belief about their rank and, in turn, require the same bonus payment to work.

By reducing the average total payment from the principal to implement the same effort profile, allowing for rank uncertainty also reduces the sum of utilities faced by the agents relative to the case of public contracting. We should note, however, that this conclusion is sensitive to the details of our model as stated, in particular that agents are risk neutral. Indeed, if agents were risk averse over money, then (suitably modifying each f_i) a nearly identical analysis would reproduce the results of this section. In this case, if all agents are symmetric, then optimal incentive schemes can in fact be strictly better for the agents (in the sense of utilitarian welfare) than an optimal public-contracts scheme. Intuitively, agents face no risk or dispersion in pay (given Corollary 1) for a successful project, whereas the public-contracts case causes harmful dispersion in pay between agents.

IV. Rank Uncertainty and Comparative Statics

In this section, we study the optimal degree of rank uncertainty and the comparative statics of optimal contracts. We have seen in Subsection IIIC that if agents are symmetric, then the principal offers them the same bonus payments while making them completely uncertain about their ranks. How does the principal's optimal incentive scheme change if agents become asymmetric? Is it always beneficial to induce rank uncertainty? How do optimal bonuses depend on agents' characteristics?

To address these questions, we focus on asymmetries in agents' costs of effort and take a production technology that is symmetric across the agents. Specifically, let the production function satisfy $P(J) = P\{1,\ldots,|J|\}$ for each $J \subseteq N$, so the probability of project success depends only on how many agents work and not on their identities. Mildly abusing notation, let $P(k) := P\{1,\ldots,k\}$ for each $k \in \{1,\ldots,N\}$ and $P(0) := P(\emptyset)$. With the only remaining heterogeneity being in the agents' costs of effort, we index the agents according to these costs: $c_1 \leq \cdots \leq c_N$. In what follows, we refer to this case as *symmetric production*.

A symmetric production function simplifies the description of the principal's optimal incentive scheme. Recall that, by our previous results, this scheme is associated with an optimal ranking distribution μ that minimizes the principal's cost of providing incentives, $\sum_{i\in N} f_i(\mu)$. With symmetric production, any ranking belief $\mu_i \in \Delta\Pi$ of agent $i \in N$ can be identified with a vector $(\mu_{ij})_{j\in N}$, where μ_{ij} is the probability that agent i is rank j. In turn, any ranking distribution can be described by a doubly stochastic matrix $(\mu_{ij})_{i,j\in N}$, namely a matrix of probabilities each of whose rows and columns sums to 1.[20] (The latter property reflects the principal's constraint that each agent must be assigned to some rank and each rank to some agent.) With some abuse of notation, we will denote the matrix $(\mu_{ij})_{i,j\in N}$ simply by μ and refer to it as a ranking matrix. Using these objects, consider the following definitions.

DEFINITION 2: *Given symmetric production and ranking matrix μ, say:*

(*i*) *Agent i is* ranked higher *than i′ if* $\mu_{i'j} > 0 \Rightarrow \mu_{ij'} = 0$ *for all* $j' \geq j$.

(*ii*) *There is* rank uncertainty *if there exist distinct agents* $i, i' \in N$, *such that neither is ranked higher than the other.*

(*iii*) *There is* complete rank uncertainty *over* $S \subseteq N$ *if each bijection* $\tilde{\pi}: S \to \{1,\ldots,|S|\}$ *has* $j_1 < \cdots < j_{|S|}$ *such that* $\tilde{\mu}_{\pi^{-1}(1)j_1},\ldots,\tilde{\mu}_{\pi^{-1}(|S|)j_{|S|}} > 0$.

[20] Conversely, every doubly stochastic matrix is induced by some ranking distribution by the Birkhoff-von Neumann theorem.

An agent is ranked higher than another under a ranking distribution μ if it is certain from their respective rank distributions that he is assigned to a strictly higher rank (recall that higher ranks are represented by lower rank indices). Accordingly, a ranking distribution μ exhibits rank uncertainty if not all pairs of agents are ranked under μ. Finally, μ exhibits complete rank uncertainty over a set $S \subseteq N$ if no ranking among the agents in this set can be ruled out. The next result describes the kind of rank uncertainty associated with the principal's optimal incentive scheme.

PROPOSITION 1: *Given symmetric production, there is a weak order $\succsim$ on N such that (i) if $i \succ i'$, then every optimal ranking distribution ranks agent i higher than agent i'; and (ii) an optimal ranking distribution exists which induces complete rank uncertainty over each $\sim$ equivalence class. Moreover, $c_i < c_{i'}$ if $i \succ i'$.*

Proposition 1 says that for any two agents, either the lower-cost agent is ranked higher than the higher-cost agent in any optimum, or there exists an optimal ranking distribution in which the two agents belong to a set exhibiting complete rank uncertainty. Importantly, recall that agents face this uncertainty not only ex ante but also interim, as they learn minimal payoff-relevant information about others from their own contract offers.

To prove Proposition 1, we consider an optimal ranking matrix, call it μ^*, with the smallest number of zero entries. Observe that every optimal ranking matrix must have a set of zero entries that contains that of μ^* because the set of optimal ranking matrices is convex. We then establish the result by proving that μ^* is a block diagonal matrix with strictly positive blocks (i.e., inducing complete rank uncertainty over each block). The blocks of μ^* correspond to the equivalence classes under $\succsim$ in the statement of Proposition 1, and thus the order $\succsim$ satisfies $c_i < c_{i'}$ for any agents $i \succ i'$. To see the logic for the latter, consider the principal's cost of providing incentives under ranking matrix μ:

$$\sum_{i \in N} f_i(\mu) = P(N) \sum_{i \in N} \frac{c_i}{\sum_{j=1}^{N} \mu_{ij}\big[P(j) - P(j-1)\big]}. \tag{5}$$

By supermodularity of P, being assigned to lower ranks increases an agent i's expected contribution to project success, $\sum_{j=1}^{N} \mu_{ij}\big[P(j) - P(j-1)\big]$. Since the fee $f_i(\mu)$ required to incentivize an agent i is strictly submodular in this expected contribution and c_i, it follows from (5) that ranking i' lower than i can only be optimal if $c_i < c_{i'}$. That is, the principal optimally assigns lower ranks to higher-cost agents in order to relax their more demanding incentive constraints.

In the online Appendix, we solve explicitly for the order $\succsim$ described in Proposition 1, thus fully characterizing the degree of rank uncertainty that is optimal for the principal.[21] The order $\succsim$ can be derived from the principal's first-order conditions and reflects an intuitive trade-off. On the one hand, because agents'

[21] See Claim 10 in the proof of Proposition 2. While Proposition 1 says that $\succsim$ captures the degree of rank uncertainty in some optimum (and provides limits on it for all optima), one may wonder whether $\succsim$ indeed fully characterizes the degree of rank uncertainty in all optimal ranking distributions. We address this issue in Claim 11 in the proof of Proposition 2: we show that, if μ is any optimal ranking matrix, then $\succsim$ is the finest weak order of the agents such that μ ranks each agent i higher than i' whenever the order does.

incentives to work increase at a decreasing rate with their beliefs about other agents' work, the principal benefits from reducing differences in agents' beliefs. This is the force that leads to complete rank uncertainty when agents are symmetric, and which calls here for agents to be grouped into the same equivalence class. On the other hand, because higher-cost agents have lower incentives to work for any given belief about other agents' work, the principal benefits from assigning them more optimistic beliefs. As discussed above, this force calls for higher-cost agents to be assigned to a lower-ranked class.

The implication of this trade-off is that the optimality of rank uncertainty depends on how symmetric agents' costs are, as summarized by the conditions in the next proposition.

PROPOSITION 2: *Given symmetric production:*

(*i*) *Every optimal ranking distribution exhibits rank uncertainty if and only if*

$$\frac{P(i+1)-P(i)}{\sqrt{c_{i+1}}} > \frac{P(i)-P(i-1)}{\sqrt{c_i}} \quad \text{for some } i \in \{1,\ldots,N-1\}. \tag{6}$$

(*ii*) *There exists an optimal ranking distribution that exhibits complete rank uncertainty over the whole set N of agents if and only if*

$$\frac{P(N)-P(0)}{\sum_{i=1}^{N}\sqrt{c_i}} > \frac{P(n)-P(0)}{\sum_{i=1}^{n}\sqrt{c_i}} \quad \text{for all } n \in \{1,\ldots,N-1\}. \tag{7}$$

Condition (6) in the first part of the proposition says that rank uncertainty is strictly optimal if and only if there is a pair of agents who are sufficiently symmetric. Note that, by our previous results, a ranking distribution with no rank uncertainty would assign any two agents i and $i+1$ to ranks i and $i+1$ respectively. But if (6) holds, then the principal could strictly reduce her total cost in (5) by introducing uncertainty over the ranking of agents i and $i+1$, namely assigning them to ranks $i+1$ and i respectively with positive probability. Condition (7) in the second part of Proposition 2 follows in a similar manner, showing that complete rank uncertainty over the whole set of agents is optimal if and only if all agents are symmetric enough.

Naturally, asymmetries in agents' effort costs affect not only the optimal ranking distribution but also the optimal profile of bonuses. From Corollary 1, we know that the principal pays identical bonuses to any two agents who are symmetric. Suppose instead that the agents have different effort costs. Will the higher-cost agent receive a larger or a smaller bonus than the lower-cost agent? On the one hand, fixing an agent's ranking belief, a higher cost requires a larger bonus to incentivize the agent's effort. On the other hand, as discussed above, the principal optimally tailors ranking beliefs in order to relax the incentive constraints and thus reduce the bonus payments of higher-cost agents.

A priori, it is thus unclear how an agent's optimal bonus changes with his cost of effort. In fact, in the setting of Winter (2004) where contracts are public and thus all agents are ranked, the relationship is non-monotonic: increasing an agent's effort cost can first lead to a reduction of the agent's bonus as the principal lowers his rank, while subsequently leading to an increase of the agent's bonus as his incentive constraint continues to tighten. In our setting with private contracts, in contrast, changes in agents' ranks are "smoothed out" by rank uncertainty, and we can show as a result that the relationship is monotonic: agent i's optimal bonus b_i^* increases with his cost c_i. At the same time, since the principal assigns more optimistic ranking beliefs to higher-cost agents relative to lower-cost agents, we also find that agent i's markup b_i^*/c_i decreases with his cost c_i, whereas both his bonus and markup increase with the cost c_j of other agents $j \neq i$.[22]

PROPOSITION 3: *Take symmetric production, and let b^* be the unique optimal limiting bonus profile. For each $i \in N$, agent i's bonus b_i^* is strictly increasing in his effort cost c_i, while his markup b_i^*/c_i is decreasing in c_i. Agent i's bonus b_i^* and markup b_i^*/c_i are both increasing in the effort cost c_j of any agent $j \neq i$.*

Figure 1 provides an illustration of the results presented in this section. We take the two-worker setting of Section I but allow for asymmetric costs of effort. The horizontal axis in each panel depicts the cost c_2 of worker 2, with the cost of worker 1 fixed at $c_1 = 1$. The vertical axis shows the workers' optimal ranking beliefs μ_{12} and μ_{22} in the top panel, their optimal bonuses b_1 and b_2 in the middle panel, and their resulting markups b_1/c_1 and b_2/c_2 in the bottom panel. We see that the workers are assigned uniform beliefs and offered equal bonuses and markups when they are symmetric. As worker 2's effort cost c_2 increases, the principal increases his bonus as well his ranking belief, thus lowering the ranking belief of worker 1. Consequently, the principal must also increase worker 1's bonus and thus his markup, whereas worker 2's markup goes down. As long as the difference in workers' costs remains small enough, inducing rank uncertainty remains optimal, with workers' beliefs interior and their bonuses increasing in c_2. It is only when c_2 becomes significantly higher than c_1 that the principal benefits from ranking the workers; further increases in c_2 then no longer affect ranking beliefs, so worker 2's bonus increases proportionally with c_2 while worker 1's bonus and both workers' markups stay constant.

Overall, our model predicts an organizational hierarchy determined by workers' effort costs or skill. Similarly skilled workers are assigned to the same hierarchical level, with uncertainty about their relative positions and with their pay changing with their relative skills. Instead, workers of sufficiently higher skill are publicly assigned to higher hierarchical levels and have their pay unaffected by those in lower levels. Throughout, the firm rewards higher-skilled workers with higher markups; these rents ensure that they are willing to work no matter what workers of lower skill do, and in turn that lower-skilled workers are also willing to work despite their lower markups.

[22] As illustrated in the example of Figure 1, b_i^*/c_i decreases strictly when c_i increases unless i is in a $\sim$ equivalence class by himself before and after the change, and b_i^* and b_i^*/c_i increase strictly when c_j increases unless i and j are in different $\sim$ equivalence classes before and after the change. Our proof of Proposition 3 in the online Appendix solves for the agents' optimal bonuses in closed form. The explicit formula that we derive, and the arguments supporting it, are adapted from those used in Rappoport (2020) to construct a receiver-optimal equilibrium in a verifiable disclosure setting.

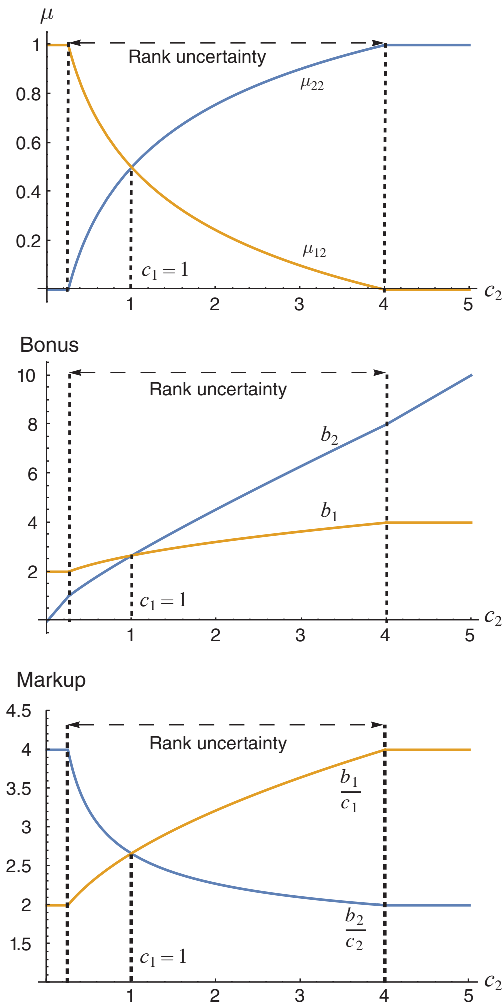

FIGURE 1. OPTIMAL RANKING BELIEFS, BONUSES, AND MARKUPS AS A FUNCTION OF EFFORT COSTS

Note: The figure is drawn for the two-worker setup of Section I with $p = 0.5$, $c_1 = 1$, and c_2 varying from 0 to 5.

V. Discussion

In this section, we briefly describe some ways in which our model can be enriched, point to how the analysis changes, and identify some open questions.

Information Sharing.—Our analysis, and the possibility of rank uncertainty and its benefits, relies on the principal's ability to withhold information from one agent about another agent's contract terms. A natural question, then, is how robust the principal's scheme is to agents being able to share this information. It turns out that the possibility of information sharing can completely undermine the principal's ability to leverage rank uncertainty.

To see this, let us restrict attention to the case of symmetric production and suppose that the agents have the option to verifiably disclose their type to other agents before effort choices are made. Given an incentive scheme $\sigma = \langle q, B \rangle$, we claim there is a perfect Bayesian equilibrium in which all agents disclose their types, and the lowest-effort equilibrium of the associated complete-information game (which exists and is in pure strategies) is played thereafter.[23] For any $t \in T^q$, let $J(t)$ be the set of agents who work in said equilibrium of the induced complete-information game. Additionally, for each $i \in N$ and type profile t_{-i} of agents other than i, let $t_i^*(t_{-i}) \in T_i^q$ be chosen to minimize the effort of such agents in the complete-information game for (t_i^*, t_{-i}); that is, let $t_i^*(t_{-i}) \in \operatorname{argmin}_{t_i \in T_i^q} |J(t) \setminus \{i\}|$. We can then construct an equilibrium in which, whenever an agent i fails to disclose his type, all other agents believe that agent i's type is $t_i^*(t_{-i})$, therefore punishing him for nondisclosure.

More precisely, consider the following strategy profile. Every agent discloses his type, and then exactly the agents in $J(t)$ work if type profile t was disclosed. If everybody but one agent i discloses his type, then all disclosing agents believe that agent i's type is $t_i^*(t_{-i})$, and exactly the agents in $J(t_i^*(t_{-i}), t_{-i})$ work. If any larger group of agents fail to disclose, let agents' beliefs about the types of non-disclosing agents and resulting continuation equilibrium be arbitrary. It is straightforward that this profile constitutes an equilibrium, and that this equilibrium generates a payoff for the principal that she could have attained by using public contracts.

In sum, we find that if agents can voluntarily, verifiably disclose their contracts to other agents, the principal cannot hope to outperform the public-contracts benchmark studied by Winter (2004). As noted in the introduction, this analysis may help explain why, in reality, firms insist that employees do not share their contractual terms with each other, often formally prohibiting them from discussing any salary information and punishing those who do so.

Strategic Substitutability.—While complementarities are a natural feature of team production, some substitutability might exist, either due to congestion costs or because some aspects of production may only require the efforts of a small number of workers. If such substitutability overwhelms any complementarity so that P is submodular, then our principal's problem is simple: making it an equilibrium for all agents to work necessitates making it dominant for all agents to work, because the case that all other agents are working is the worst-case belief for an agent's own marginal product. In particular, in this case, there is then no scope for beneficial rank uncertainty.

But what if, as seems more realistic when there are many agents, P is neither supermodular nor submodular? In this case, an analysis nearly identical to ours can be used to study the least-cost way to incentivize work from every agent as a unique (interim correlated) *rationalizable* choice. The main difference in the analysis is that an agent who knows that a subset J of other agents is working may suspect that even more other agents are working, and this can make him even more pessimistic about

[23] Although we cannot directly apply their result here, this constructive argument is similar to that in Theorem 1 of Hagenbach, Koessler, and Perez-Richet (2014) in the special case where all types share the same worst-case belief.

his own marginal product. To take this possibility into account, for each $J \subseteq N$ and $i \in N \backslash J$, let

$$\eta_i(J) := \min\left\{ P\left(\hat{J} \cup \{i\}\right) - P\left(\hat{J}\right) : J \subseteq \hat{J} \subseteq N \backslash \{i\} \right\}$$

denote the minimum marginal product that agent i can entertain given that all agents $j \in J$ work.[24] With this worst-case marginal product in hand, the rest of our analysis extends readily. Specifically, redefining the cost of providing incentives to agent i (given by (3) in our analysis) via

$$f_i(\mu_i) := \frac{c_i P(N)}{\sum_{\pi \in \Pi} \mu_i(\pi) \eta_i\left(\left\{ j \in N : \pi_j < \pi_i \right\}\right)},$$

all of our main results go through as stated.

At this stage, it remains an open question whether the principal can more cheaply incentivize all agents to work as a unique *equilibrium* instead of a unique rationalizable choice. The reason is that different types of a given agent might have their worst-case beliefs generated by different strategies for other agents. This means that the principal may be able to implement work as a unique equilibrium at a lower cost, but also that one cannot optimize agents' bonuses separably across their types, making the analysis more delicate.

Grades of Success.—In our model, agents work on a project which can only either succeed or fail. In reality, however, projects may yield different grades of success, and the principal may in turn want to pay the agents following different success outcomes. We find that a complete characterization of the principal's optimal incentive scheme can become more difficult in such a setting. The main reason is that the principal may now benefit from providing agents with private information about their realized ranking, unlike in our baseline model.

To illustrate, consider a simple setting with two symmetric agents. Suppose each agent chooses whether to work on an individual task as in our example of Section I, but now working results in partial completion of the task with probability $q \in (0,1)$ and full completion with probability $1-q$, whereas shirking results in partial completion with probability $p \in (0,q)$ and no completion with probability $1-p$. Assume the project outcome is failure if either one or both tasks are not completed, partial success if the two tasks are only partially completed, and full success in the remaining cases. Table 1 shows the resulting probabilities of partial and full success as a function of the number of agents working.

Observe that the probabilities of partial and full success are each strictly increasing and strictly supermodular in the set of agents who work. Hence, as in our baseline model, failure is indicative of shirking and will never merit a positive bonus from the principal. Furthermore, given that and for any subset of agent types that an agent is certain are working, the agent's worst-case belief is the one that takes all other agent types to be shirking. This means that it remains without loss of optimality for the

[24] Clearly, a minimizing $\hat{J}$ in the definition of $\eta_i(J)$ is simply equal to J itself (respectively, $N \backslash \{i\}$) in the case in which P is supermodular (respectively, submodular).

TABLE 1—PROBABILITIES OF PARTIAL AND FULL SUCCESS IN GRADES OF SUCCESS EXAMPLE

Number working	Partial	Full
2	q^2	$1-q^2$
1	pq	$p(1-q)$
0	p^2	0

principal to focus on ranking schemes, as in our baseline model, albeit with suitably modified bonus payments. Specifically, let $P_s(k)$ denote the probability of success level $s \in \{partial, full\}$ when k agents work, and denote by $\mu_i \in [0,1]$ an agent's worst-case belief that the other agent is working. Then we redefine

$$(8) \qquad f_i(\mu_i) = \frac{c}{\max_{s\in\{partial, full\}} \frac{\mu_i[P_s(2)-P_s(1)] + (1-\mu_i)[P_s(1)-P_s(0)]}{P_s(2)}}.$$

As shown by the denominator on the right-hand side, given an agent's ranking belief μ_i, the principal will benefit from rewarding the agent based on the success outcome for which he has the highest expected contribution to success scaled by its probability of occurring (in the uniquely implemented work equilibrium). Note that if an agent is certain that the other agent is working (i.e., $\mu_i = 1$), this success outcome is full success, since $[1-q^2-p(1-q)]/(1-q^2) > (q^2-pq)/q^2$. Assume instead that if an agent takes the other agent to be shirking (i.e., $\mu_i = 0$), then this success outcome is partial success; that is, assume $p(1-q)/(1-q^2) < (pq-p^2)/q^2$.[25] This means that there is a belief $\mu^* \in (0,1)$ such that the principal will want to reward a partial success for $\mu_i < \mu^*$ and a full success for $\mu_i > \mu^*$.

Figure 2 shows the resulting shape of the minimum payment function f_i. As in our baseline model, increasing an agent's expected contribution to project success allows the principal to reduce his bonus payment at a decreasing rate; that is, $f_i(\mu_i)$ is convex in the expected contribution induced by μ_i. However, when multiple success outcomes are rewarded, the agent's expected contribution to success can be nonlinear in the ranking belief. In this example, this occurs at μ^*, where the principal switches from paying for partial success to paying for full success. As a result, we obtain that $f_i(\mu_i)$ is convex in μ_i on each side of μ^* but exhibits a concave kink at μ^*.

The concave kink in the function f_i implies that Theorem 1 no longer applies. Instead, we can prove an analog to the result where we replace f_i with its convex envelope.[26] The interpretation is that the principal might now find it optimal to provide agents with private information about their ranks. As a consequence, the optimal incentive scheme becomes more complicated to characterize. Naturally, introducing rank uncertainty will remain strictly optimal for the principal under

[25] This inequality is equivalent to $q > p(1+q)$, which will clearly hold if p is small enough, that is, if effort is sufficiently important in achieving partial success.

[26] In our problem, what matters for the principal's cost of providing incentives is each agent's marginal distribution over ranking beliefs. This follows from additivity as in equation (4).

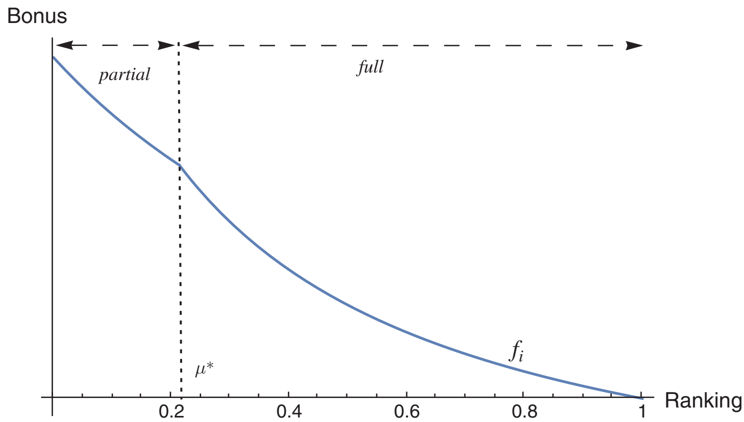

Figure 2. Nonconvex Minimum Payment Function in Grades of Success Example

Note: The figure depicts $f_i(\mu_i)$ in equation (8) with $c = 1$ and the probabilities in Table 1 for $q = 0.8$ and $p = 0.3$.

some parameter values. In fact, in this simple example, this will be the case unless the kink at μ^* is severe enough that the convex envelope of f_i is linear.

Interdependent Contracting.—In our model, an incentive scheme is required to make each agent's success-contingent bonus a function solely of his own type. This modeling choice reflects a realistic feature of employees' contracts, namely that they are perfectly informed of the terms of their own employment. However, in settings in which bonuses are of a discretionary nature, it is conceivable that agents face uncertainty about their own pay.[27] We study here how allowing each agent's bonus to depend on the entire profile of agent types changes our results.

Let an incentive scheme $\sigma = \langle q, B \rangle$ specify a finite-support prior $q \in \Delta(\mathbb{N}^N)$ and a bonus rule $B = (B_1, \ldots, B_N)$ as in our baseline model, but where the bonus for each agent $i \in N$ is now $B_i(t)$, i.e., a function not only of agent i's type $t_i \in T_i^q$ but of the entire type profile $t \in T^q$. We begin by noting that this distinction is immaterial if the principal only wants to implement every agent working as *some* equilibrium. Such an equilibrium requires that, conditional on all other agent types working, each agent prefer to work whatever his type than to shirk whatever his type: his average success-contingent bonus must be at least

$$\sum_{t \in T^q} q(t)\, B_i(t) \;\geq\; \frac{c_i}{P(N) - P(N \backslash \{i\})} \;=:\; b_i^L.$$

Clearly, publicly offering a bonus b_i^L to each agent yields every agent working as an equilibrium, so the possibility of making agents uncertain about other agents' or their own contract terms does not matter for this design problem.

But what if, as motivates our study, the principal wishes to implement every agent working as a *unique* equilibrium? Of course, she can do no better than having to

[27] Other work has highlighted potential benefits to subjecting a worker to opaque incentives. For example, the benefits of such opacity have been studied in the context of multidimensional decision making (Jehiel 2015; Ederer, Holden, and Meyer 2018).

pay only $P(N)\,b_i^L$ on average to each agent $i \in N$. We next show that when agents' bonuses can depend on the entire type profile, the principal can, in fact, approximate this theoretical bound arbitrarily well.

For given $\gamma \in (0,1)$, consider the following incentive scheme. Each agent $i \in N$ is assigned type $t_i = 1$ with probability γ and type $t_i = 2$ with probability $1-\gamma$, independent across agents. The type *profile*-contingent bonus rule is given by

$$B_i(t) := \begin{cases} \frac{c_i}{P\{i\} - P(\emptyset)}, & \text{if } t_i = 1; \\ \frac{1}{\gamma^{N-1}} \cdot b_i^L, & \text{if } t_i = 2 \text{ and } t_j = 1 \text{ for all } j \in N \backslash \{i\}; \\ 0, & \text{otherwise.} \end{cases}$$

Observe that this scheme uniquely implements work: a type-1 agent finds it dominant to work, and a type-2 agent wants to work given that all other agents work when assigned type 1 (since the probability that all others are type 1 is γ^{N-1}). Furthermore, the expected payment to each agent $i \in N$ is

$$P(N)\left[\gamma \cdot \frac{c_i}{P\{i\} - P(\emptyset)} + (1-\gamma)\gamma^{N-1} \cdot \frac{1}{\gamma^{N-1}} b_i^L\right],$$

which converges to $P(N)\,b_i^L$ as $\gamma \to 0$.

We thus find that if the principal is not required to inform agents of their own contractual terms, then she can implement work while excluding other equilibria at no extra cost. Allowing for contracts to be private in this (possibly unrealistically) stronger sense makes the unique implementation requirement essentially lose all of its bite.[28]

Appendix A. Proofs for Section III

This Appendix provides proofs for the results in Section III. See the online Appendix for proofs of the results in Section IV.

A. *Proof of Lemma 1*

We begin by proving the first part, namely that every ranking scheme uniquely implements work. Fix a ranking scheme q and an $\varepsilon > 0$. Consider an arbitrary Bayesian Nash equilibrium of the Bayesian game induced by $\langle q, B + \varepsilon \rangle$. By definition of a ranking scheme, every type t_i of every agent $i \in N$ has work as his unique best response given bonus

$$\frac{c_i}{\sum_{t_{-i}} q_i(t_{-i}|t_i)\left[P\{j \in N : t_j \leq t_i\} - P\{j \in N : t_j < t_i\}\right]} + \varepsilon$$

[28] The idea that complex enough contracts can allow one to exclude undesirable equilibria at no additional cost arises in other models of contracting with externalities. In particular, in settings in which agents' choices are contractible, this is true if contracts are not required to be bilateral but can be made contingent on third parties' choices; see Segal (1999, 2003).

if exactly all strictly lower types of other agents do, and so, by supermodularity of P, if at least all strictly lower types do. Therefore, it follows by induction on the type that each type t_i of every agent $i \in N$ works in the equilibrium.

We next prove the second part of the lemma. Fix an incentive scheme $\sigma = \langle q, B\rangle$ that uniquely implements work. We proceed in three steps.

Step 1: We assign each agent type t_i a finite index $k_i(t_i)$, which will represent the stage at which shirking is eliminated for said type under iterated deletion of weakly dominated strategies. Formally, $k_i: T_i^q \to \mathbb{N} \cup \{\infty\}$ (for $i \in N$) is defined as follows. Let $T_i^q(0) := \emptyset$, and then, recursively for $\kappa \in \mathbb{N}$, let $T_i^q(\kappa)$ be the set of all $t_i \in T_i^q$, such that

$$B_i(t_i)\sum_{t_{-i}} q_i(t_{-i}|t_i)\Big[P\Big(\big\{j \in N: t_j \in T_j^q(\kappa-1)\big\} \cup \{i\}\Big) - P\Big(\big\{j \in N: t_j \in T_j^q(\kappa-1)\big\} \backslash \{i\}\Big)\Big] \geq c_i.$$

Define $k_i(t_i) := \inf\big\{\kappa \in \mathbb{N}: t_i \in T_i^q(\kappa)\big\}$ for each $i \in N$ and $t_i \in T_i^q$.

Assume, for a contradiction, that some type of some agent has index equal to ∞. As T^q is finite, there is some $\overline{\kappa} \in \mathbb{N}$ strictly higher than everything in $\bigcup_{i \in N} k_i(T_i^q) \backslash \{\infty\}$. Now, for each agent $i \in N$ and each type $t_i \in T_i^q$ with $k_i(t_i) = \infty$, that $k_i(t_i) > \overline{\kappa}$ implies

$$\varepsilon_i(t_i) := c_i - B_i(t_i)\sum_{t_{-i}} q_i(t_{-i}|t_i)\Big[P\Big(\big\{j \in N: k_j(t_j) < \infty\big\} \cup \{i\}\Big) - P\Big(\big\{j \in N: k_j(t_j) < \infty\big\} \backslash \{i\}\Big)\Big] > 0.$$

Letting $\varepsilon := \min\big\{\varepsilon_i(t_i): i \in N, t_i \in T_i^q \text{ has } k_i(t_i) = \infty\big\} > 0$, observe that the Bayesian game induced by prior q and bonuses $B_i + \varepsilon$ has a Bayesian Nash equilibrium in which each agent i of type t_i works if $k_i(t_i) < \infty$ and shirks if $k_i(t_i) = \infty$. Indeed, the best response property for type t_i follows directly from the definition of $k_i(t_i)$ when $k_i(t_i) < \infty$, and it follows additionally from the definition of ε when $k_i(t_i) = \infty$. We have thus constructed a Bayesian Nash equilibrium of the Bayesian game induced by $\langle q, B + \varepsilon\rangle$ in which some type shirks, contradicting the hypothesis that σ uniquely implements work. Thus, every type is assigned a finite index.

Step 2: Using the indices from Step 1, we construct a ranking scheme $\sigma^* = \langle q^*, B^*\rangle$ from the original scheme $\sigma = \langle q, B\rangle$. First, we relabel types so that higher indices always correspond to higher types. That is, we construct a one-to-one function

$$\lambda: \bigcup_{i \in N} \big[\{i\} \times T_i^q\big] \to \mathbb{N}$$

such that, for any $i, j \in N$ and $t_i \in T_i^q$, $t_j \in T_j^\sigma$ with $k_i(t_i) > k_j(t_j)$, we have $\lambda(i, t_i) > \lambda(j, t_j)$; hence, a given type's perceived probability of other agents' types having a lower image under λ is at least as high as the probability of said agents having shirk eliminated sooner under iterated deletion.[29] Then, using this function, construct the incentive scheme $\sigma^* = \langle q^*, B^* \rangle$ as follows. Let the prior q^* be given by

$$q^*(t^*) = \begin{cases} q(t) & \text{if } t \in T^q \text{ has } \lambda(i,t_i) = t_i^* \text{ for all } i \in N; \\ 0, & \text{otherwise.} \end{cases}$$

Let the bonus rule $B^* = (B_i^*)_{i \in N}$ be given by $B_i^*(t_i^*) := (1/P(N)) f_i\big(\mu_i^{q^*}(\cdot|t_i^*)\big)$, where the ranking belief $\mu_i^{q^*}(\cdot|t_i^*)$ is as given in (2) and the minimum cost function f_i is as given by (3). Using these equations, it follows readily that

$$B_i^*(t_i^*) = \frac{c_i}{\sum_{t_{-i}^*} q_i^*(t_{-i}^*|t_i^*)\Big[P\big(\{j \in N : t_j^* < t_i^*\} \cup \{i\}\big) - P\{j \in N : t_j^* < t_i^*\}\Big]}.$$

Injectivity of λ ensures that, with q^*-distributed type profiles, no two agents ever have the same type realized. It follows directly that q^* is a ranking scheme.

Step 3: All that remains is to show that the bonus distributions induced by the ranking scheme q^* constructed in Step 2 are first-order stochastically dominated by those induced by the original scheme $\sigma = \langle q, B \rangle$. As q^* is constructed from q by relabeling types and leaving their respective probabilities unchanged, it suffices to check that every $i \in N$ and $t_i \in T_i^q$ has $B_i(t_i) \geq B_i^*\big(\lambda(i,t_i)\big)$. Letting $t_i^* := \lambda(i,t_i)$, we indeed verify that

$$\begin{aligned} B_i(t_i) & \\ &\geq \frac{c_i}{\sum_{t_{-i}} q_i(t_{-i}|t_i)\Big[P\big(\{j \in N : t_j \in T_j^q(k_i(t_i)-1)\} \cup \{i\}\big) - P\big(\{j \in N : t_j \in T_j^q(k_i(t_i)-1)\} \setminus \{i\}\big)\Big]} \\ &= \frac{c_i}{\sum_{t_{-i}} q_i(t_{-i}|t_i)\Big[P\big(\{j \in N : k_j(t_j) < k_i(t_i)\} \cup \{i\}\big) - P\{j \in N : k_j(t_j) < k_i(t_i)\}\Big]} \\ &\geq \frac{c_i}{\sum_{t_{-i}} q_i(t_{-i}|t_i)\Big[P\big(\{j \in N : \lambda(j,t_j) < \lambda(i,t_i)\} \cup \{i\}\big) - P\{j \in N : \lambda(j,t_j) < \lambda(i,t_i)\}\Big]} \\ &= \frac{c_i}{\sum_{t_{-i}^*} q_i^*(t_{-i}^*|t_i^*)\Big[P\big(\{j \in N : t_j^* < t_i^*\} \cup \{i\}\big) - P\{j \in N : t_j^* < t_i^*\}\Big]} = B_i^*(t_i^*), \end{aligned}$$

where the first inequality follows from the definition of $T_i^q(\kappa)$ for $\kappa = k_i(t_i)$, and the second inequality follows from supermodularity of P and that $\{j \in N : k_j(t_j) < k_i(t_i)\} \subseteq \{j \in N : \lambda(j,t_j) < \lambda(i,t_i)\}$ for every agent $i \in N$ and type $t_i \in T_i^q$. ■

B. *Proof of Theorem 1*

Note that a minimizer to $\sum_{i \in N} f_i$ exists and each f_i is bounded because the function is continuous over its compact domain.

[29] For instance, letting $L := 1 + n + \max_{i \in N, t_i \in T_i^q} t_i$, take $\lambda(i,t_i) := L^2 k_i(t_i) + Li + t_i$.

We first show that $\inf_{\sigma\,UIW} W(\sigma) \geq \min_{\mu\in\Delta\Pi}\sum_{i\in N} f_i(\mu)$. Given Lemma 1, it is sufficient to show that the principal's value for a ranking scheme q is no less than $\min_{\mu\in\Delta\Pi}\sum_{i\in N} f_i(\mu)$. Bayesian updating implies that a given agent i's ranking belief is, on average, equal to the true ranking distribution:

$$\sum_{t\in T^q} q(t)\,\mu_i^q(\cdot|t_i) = \sum_{t_i\in T_i^q} q_i(t_i)\,\mu_i^q(\cdot|t_i) = \mu^q.$$

Hence, since f_i is convex, Jensen's inequality implies

$$\sum_{t\in T^q} q(t) f_i\big(\mu_i^q(\cdot|t_i)\big) \geq f_i(\mu^q).$$

Summing over $i \in N$ yields $W(q) \geq \sum_{i\in N} f_i(\mu^q)$, proving the claim.

We next show that $\inf_{\sigma\,UIW} W(\sigma) \leq \min_{\mu\in\Delta\Pi}\sum_{i\in N} f_i(\mu)$. It suffices to find, given arbitrary $\mu \in \Delta\Pi$, a sequence $(q^m)_m$ of ranking schemes such that $W(q^m) \to \sum_{i\in N} f_i(\mu)$ as $m \to \infty$. For each $m \in \mathbb{N}$, define the finite-support prior $q^m \in \Delta(\mathbb{N}^N)$ as follows:

$$q^m(t) := \begin{cases} \frac{1}{m}\mu(\pi), & \text{if } \exists \ell \in \{0,\dots,m-1\} \text{ with } t_i = \ell + \pi_i \text{ for all } i \in N; \\ 0, & \text{otherwise.} \end{cases}$$

Let us now observe $W(q^m)$ converges (as $m \to \infty$) to the appropriate payoff bound. Indeed, by construction, for each $m \in \mathbb{N}$ and $i \in N$, every type $t_i \in T_i^{q^m}$ with $N \leq t_i \leq m$ has ranking belief $\mu_i^{q^m}(\cdot|t_i) = \mu$ and thus is offered bonus $\big(1/P(N)\big) f_i(\mu)$. Note that such types of agent i arise under q^m with probability $\max\{0, (m-N+1)/(m+N-1)\}$, which converges to 1 as $m \to \infty$. Hence, the expected bonus payment to each agent i converges to $f_i(\mu)$, proving the claim. ∎

C. *Proof of Theorem 2*

Before proving the theorem, we record the following lemma about the function $f_i : \Delta\Pi \to \mathbb{R}_{++}$ that describes the expected fees required to incentivize agent $i \in N$.

LEMMA 2: *For each $i \in N$, every $\tau_i \in \Delta\Delta\Pi$ has $\int f_i\, d\tau_i \geq f_i\big(\int \mu_i\, d\tau_i(\mu_i)\big)$, with a strict inequality if τ_i induces a nondegenerate distribution of f_i.*

PROOF:

The function $\mathbb{R}_{++} \to \mathbb{R}$ given by $x \mapsto 1/x$ is strictly convex and continuous, while the function $1/f_i$ is affine and continuous. Therefore, Jensen's inequality yields

$$\int f_i\, d\tau_i = \int \frac{1}{\left(\frac{1}{f_i}\right)}\, d\tau_i \geq \frac{1}{\int \frac{1}{f_i}\, d\tau_i} = f_i\Big(\int \mu_i\, d\tau_i(\mu_i)\Big),$$

where the inequality is strict if the τ_i-distribution of f_i (and thereby $1/f_i$) is nondegenerate. ∎

We now proceed to prove the theorem. Assume for a contradiction that the first claim is not true. Since $\sum_{i\in N} f_i(\mu)$ attains a minimum somewhere, there must then be two minimizers $\mu', \mu'' \in \Delta\Pi$ of $\sum_i f_i(\mu)$ with $(f_i(\mu'))_{i\in N} \neq (f_i(\mu''))_{i\in N}$. Applying Lemma 2 to a uniform mixture over $\{\mu', \mu''\}$ yields

$$\sum_{i\in N} f_i\Big(\frac{1}{2}\mu' + \frac{1}{2}\mu''\Big) \;<\; \sum_{i\in N}\Big(\frac{1}{2}f_i(\mu') + \frac{1}{2}f_i(\mu'')\Big) \;=\; \frac{1}{2}\sum_{i\in N} f_i(\mu') + \frac{1}{2}\sum_{i\in N} f_i(\mu''),$$

contradicting that $\mu', \mu'' \in \Delta\Pi$ are minimizers. Hence, there is a unique minimizer as claimed.

Consider now the second claim. The "if" direction of this claim follows immediately from Theorem 1. To prove the "only if" direction, consider an optimal sequence $(\sigma^m)_m$ of incentive schemes. Apply Lemma 1 to obtain, for each $m \in \mathbb{N}$, a ranking scheme $\hat{q}^m$ such that $\beta_i^{\sigma^m}$ first-order-stochastically dominates the bonus distribution $\beta_i^{\hat{q}^m}$ for each $i \in N$. Notice that $\beta_i^{\hat{q}^m}$ is supported on the compact set $[0, c_i/(P\{i\} - P(\emptyset))]$ by the definition of a ranking scheme. We therefore need only prove that the unique limit point of the sequence $(\beta_i^{\hat{q}^m})_m$ is degenerate on b_i^* for all $i \in N$. Indeed, because of first-order stochastic dominance and optimality of $(\sigma^m)_m$, it follows that $(\beta_i^{\sigma^m})_m$ converges to the same degenerate bonus distribution.

For each $i \in N$ and each $m \in \mathbb{N}$, the ranking scheme $\hat{q}^m$ generates a finite-support distribution $\tau_i^m \in \Delta\Delta\Pi$ over agent i's ranking belief, given by

$$\tau_i^m(\mu_i) \;\coloneqq\; \hat{q}_i^m\{t_i : \mu_i^{\hat{q}^m}(t_i) \;=\; \mu_i\} \quad \text{for all } \mu_i \in \Delta\Pi.$$

Note that any limit point β_i of $(\beta_i^{\hat{q}^m})_m$ exhibits some limit point of $(\tau_i^m)_m$ whose induced distribution of f_i is β_i. We thus focus on an arbitrary limit point of $(\tau_i^m)_m$, denoted τ_i^*. Since each agent has Bayes-consistent beliefs about the ranking distribution,

$$\int \mu_i \, d\tau_i^m(\mu_i) \;=\; \sum_{t_i} \hat{q}_i^m(t_i)\, \mu_i^{\hat{q}^m}(t_i) \;=\; \mu^{\hat{q}^m}$$

is independent of i. Therefore, taking limits, $\int \mu_i \, d\tau_i^*(\mu_i) \eqqcolon \mu^*$ is also independent of i. We next observe that τ_i^* must induce a degenerate distribution of f_i for every $i \in N$. Indeed, if not, then Lemma 2 would imply

$$\lim_{m\to\infty} W(\hat{q}^m) \;=\; \sum_{i\in N} \int f_i \, d\tau_i^* \;>\; \sum_{i\in N} f_i(\mu^*) \;\geq\; \min_{\mu\in\Delta\Pi} \sum_{i\in N} f_i(\mu),$$

which (given Theorem 1) would violate the optimality of $(\hat{q}^m)_m$.

Hence, we obtain that $(\tau_i^m)_m$ converges to a degenerate distribution on μ^* for every agent $i \in N$, implying that $\lim_{m\to\infty} W(\hat{q}^m) = \sum_{i\in N} f_i(\mu^*)$. By the first claim in this theorem and the result in Theorem 1, optimality then requires $f_i(\mu^*) = P(N)\, b_i^*$ for all $i \in N$. This completes the proof. ∎

REFERENCES

Abreu, Dilip, and Hitoshi Matsushima. 1992. "Virtual Implementation in Iteratively Undominated Strategies: Complete Information." *Econometrica* 60 (5): 993–1008.

Bennedsen, Morten, Elena Simintzi, Margarita Tsoutsoura, and Daniel Wolfenzon. Forthcoming. "Do Firms Respond to Gender Pay Gap Transparency?" *Journal of Finance.*

Bernstein, Shai, and Eyal Winter. 2012. “Contracting with Heterogeneous Externalities.” *American Economic Journal: Microeconomics* 4 (2): 50–76.
Carlsson, Hans, and Eric van Damme. 1993. “Global Games and Equilibrium Selection.” *Econometrica* 61 (5): 989–1018.
Cullen, Zoë B., and Bobak Pakzad-Hurson. 2020. “Equilibrium Effects of Pay Transparency in a Simple Labor Market.” Unpublished.
Cullen, Zoë B., and Ricardo Perez-Truglia. 2018. “The Salary Taboo: Privacy Norms and the Diffusion of Information.” NBER Working Paper 25145.
Doval, Laura, and Jeff Ely. 2020. “Sequential Information Design.” *Econometrica* 88 (6): 2575–608.
Ederer, Florian, Richard Holden, and Margaret Meyer. 2018. “Gaming and Strategic Opacity in Incentive Provision.” *RAND Journal of Economics* 49 (4): 819–54.
Edwards, Matthew A. 2005. “The Law and Social Norms of Pay Secrecy.” *Berkeley Journal of Employment and Labor Law* 26 (1): 41–63.
Eliaz, Kfir, and Ran Spiegler. 2015. “X-Games.” *Games and Economic Behavior* 89: 93–100.
Garicano, Luis, and Timothy Van Zandt. 2013. “Hierarchies and the Division of Labor.” In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 604–54. Princeton, NJ: Princeton University Press.
Gely, Rafael, and Leonard Bierman. 2003. “Pay Secrecy/Confidentiality Rules and the National Labor Relations Act.” *Journal of Labor and Employment Law* 6: 120–56.
Habibi, Amir. 2020. “Pay Transparency in Organisations.” Unpublished.
Hagenbach, Jeanne, Frédéric Koessler, and Eduardo Perez-Richet. 2014. “Certifiable Pre-Play Communication: Full Disclosure.” *Econometrica* 82 (3): 1093–131.
Halac, Marina, Ilan Kremer, and Eyal Winter. 2020. “Raising Capital from Heterogeneous Investors.” *American Economic Review* 110 (3): 889–921.
Hegewisch, Ariane, Claudia Williams, and Robert Drago. 2011. *Pay Secrecy and Wage Discrimination*. Washington, DC: Institute for Women’s Policy Research.
Holmström, Bengt. 1982. “Moral Hazard in Teams.” *Bell Journal of Economics* 13 (2): 324–40.
Hoshino, Tetsuya. 2020. “Multi-Agent Persuasion: Leveraging Strategic Uncertainty.” Unpublished.
Inostroza, Nicolas, and Alessandro Pavan. 2020. “Persuasion in Global Games with Application to Stress Testing.” Unpublished.
Jehiel, Philippe. 2015. “On Transparency in Organizations.” *Review of Economic Studies* 82 (2): 736–61.
Kajii, Atsushi, and Stephen Morris. 1997. “The Robustness of Equilibria to Incomplete Information.” *Econometrica* 65 (6): 1283–309.
Mathevet, Laurent, Jacopo Perego, and Ina Taneva. 2020. “On Information Design in Games.” *Journal of Political Economy* 128 (4): 1370–404.
Milgrom, Paul, and John Roberts. 1990. “Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities.” *Econometrica* 58 (6): 1255–77.
Mookherjee, Dilip. 2013. “Incentives in Hierarchies.” In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 764–98. Princeton, NJ: Princeton University Press.
Moriya, Fumitoshi, and Takuro Yamashita. 2020. “Asymmetric-Information Allocation to Avoid Coordination Failure.” *Journal of Economics and Management Strategy* 29 (1): 173–86.
Morris, Stephen, Daisuke Oyama, and Satoru Takahashi. 2020. “Implementation via Information Design in Binary-Action Supermodular Games.” Unpublished.
Oyama, Daisuke, and Satoru Takahashi. 2020. “Generalized Belief Operator and Robustness in Binary-Action Supermodular Games.” *Econometrica* 88 (2): 693–726.
Rappoport, Daniel. 2020. “Evidence and Skepticism in Verifiable Disclosure Games.” Unpublished.
Rosenfeld, Jake. 2017. “Don’t Ask or Tell: Pay Secrecy Policies in U.S. Workplaces.” *Social Science Research* 65: 1–16.
Rubinstein, Ariel. 1989. “The Electronic Mail Game: Strategic Behavior under ‘Almost Common Knowledge.’” *American Economic Review* 79 (3): 385–91.
Segal, Ilya. 1999. “Contracting with Externalities.” *Quarterly Journal of Economics* 114 (2): 337–88.
Segal, Ilya. 2003. “Coordination and Discrimination in Contracting with Externalities: Divide and Conquer?” *Journal of Economic Theory* 113 (2): 147–81.
Trotter, Richard G., Susan Rawson Zacur, and Lisa T. Stickney. 2017. “The New Age of Pay Transparency.” *Business Horizons* 60 (4): 529–39.
Winter, Eyal. 2004. “Incentives and Discrimination.” *American Economic Review* 94 (3): 764–73.
Winter, Eyal. 2006. “Optimal Incentives for Sequential Production Processes.” *RAND Journal of Economics* 37 (2): 376–90.
Winter, Eyal. 2010. “Transparency and Incentives among Peers.” *RAND Journal of Economics* 41 (3): 504–23.

Online Appendix for
"Rank Uncertainty in Organizations"

by Marina Halac, Elliot Lipnowski, and Daniel Rappoport

B. Proofs for Section 5

B.1. Proof of Proposition 1

For the present proof and those of later propositions, we define the following convenient notation for reasoning about the principal's problem.

Notation 1. For each $i, i', j, j' \in N$, let $D(i, i', j, j')$ be the matrix with its (i, j) and (i', j') entries taking value 1, its (i, j') and (i', j) entries taking value -1, and all other entries taking value zero.

Notation 2. Given a ranking matrix μ and for each $i \in N$, let $I_i(\mu)$ denote agent i's **incentive effect** defined as follows:

$$I_i(\mu) := \sum_{j=1}^{N} \mu_{ij} \left[P(j) - P(j-1)\right].$$

When not confusing, we will omit the dependence on μ and simply write I_i.

To begin our proof of Proposition 1, we fix an optimal ranking matrix μ^* with the minimum number of zero entries. As convex combinations of optimal ranking matrices are themselves optimal, this minimality property in fact implies that, for any optimal ranking matrix μ, every zero entry of μ^* corresponds to a zero entry of μ, so that μ^* in fact has the minimum set of zero entries. Toward constructing our order, we define $J^*(i) := \{j \in N : \ \mu^*_{ij} > 0\}$ for each $i \in N$.

We first establish, via a perturbation argument, the following structural claim that we will use several times.

Claim 1. *Suppose $i_1, i_2, j_1, j_2, j_3 \in N$ are such that $j_1 < j_2 < j_3$, that $j_1, j_3 \in J^*(i_1)$, and that $j_2 \in J^*(i_2)$. Then $j_1, j_2, j_3 \in J^*(i_1) \cap J^*(i_2)$.*

Proof. Given $\varepsilon, \varepsilon' > 0$, define the $N \times N$ matrix

$$\mu := \mu^* + \varepsilon D(i_1, i_2, j_2, j_1) + \varepsilon' D(i_1, i_2, j_2, j_3).$$

As $\mu^*_{i_1 j_1}, \mu^*_{i_1 j_3}, \mu^*_{i_2 j_2} > 0$, the matrix μ is doubly stochastic with strictly positive entries wherever μ^* has strictly positive entries, as long as $\max\{\varepsilon, \varepsilon'\}$ is small enough. Define now the ratio

$$\rho := \frac{[P(j_2) - P(j_2 - 1)] - [P(j_1) - P(j_1 - 1)]}{[P(j_3) - P(j_3 - 1)] - [P(j_2) - P(j_2 - 1)]},$$

which is strictly positive because P is strictly supermodular. Upon choosing $\varepsilon, \varepsilon'$ to further satisfy $\frac{\varepsilon'}{\varepsilon} = \rho$, direct computation shows that μ generates the exact same incentive effects $(I_i)_{i \in N}$ as μ^* does;[29] μ is therefore also optimal. That $\mu_{i_1 j_2}, \mu_{i_2 j_1}, \mu_{i_2 j_3} > 0$ then implies $\mu^*_{i_1 j_2}, \mu^*_{i_2 j_1}, \mu^*_{i_2 j_3} > 0$ as well by definition of μ^*. *Q.E.D.*

Next, using the above structural claim, we derive more detail on the set of nonzero entries over the next three claims. First, in Claim 2, we show that the set of nonzero entries of each row is an interval. Then, in Claim 3, we show that, if two rows have distinct sets of columns in which they are nonzero, then these two column sets can overlap at most once. Finally, in Claim 4, we strengthen the latter to show that any two rows must have nonzero entries in either the exact same set of columns or in disjoint sets of columns.

Claim 2. *Suppose $i_1, j_1, j_2, j_3 \in N$ are such that $j_1 < j_2 < j_3$ and $j_1, j_3 \in J^*(i_1)$. Then $j_2 \in J^*(i_1)$.*

[29] For example, for i_1, we obtain

$$I_{i_1}(\mu) = I_{i_1}(\mu^*) + \varepsilon \left\{ \begin{array}{c} [P(j_2) - P(j_2 - 1)] \\ -[P(j_1) - P(j_1 - 1)] \end{array} \right\} + \rho\varepsilon \left\{ \begin{array}{c} [P(j_2) - P(j_2 - 1)] \\ -[P(j_3) - P(j_3 - 1)] \end{array} \right\} = I_{i_1}(\mu^*).$$

Proof. Assume for a contradiction that the claim is false. So there is some $j_2 \in N$ such that $j_1 < j_2 < j_3$ and $\mu^*_{i_1 j_2} = 0$. As the j_2 column of μ^* sums to 1, there is some $i_2 \in N$ such that $\mu^*_{i_2 j_2} > 0$. But then Claim 1 implies $\mu^*_{i_1 j_2} > 0$, a contradiction. *Q.E.D.*

Claim 3. *Suppose $i_1, i_2 \in N$ have $J^*(i_1) \neq J^*(i_2)$. Then $|J^*(i_1) \cap J^*(i_2)| \leq 1$.*

Proof. Assume for a contradiction that $J^*(i_1) \neq J^*(i_2)$ and $|J^*(i_1) \cap J^*(i_2)| > 1$. Without loss, say $J^*(i_1) \not\subseteq J^*(i_2)$. Then let $j_1 := \min J^*(i_1)$ and $j_3 := \max J^*(i_1)$. That $|J^*(i_1) \cap J^*(i_2)| > 1$ implies that $j_3 \neq \min J^*(i_2)$ and $j_1 \neq \max J^*(i_2)$. There therefore exists $j_2 \in J^*(i_2)$ such that $j_1 < j_2 < j_3$. But Claim 1 then says $j_1, j_3 \in J^*(i_2)$, so that Claim 2 implies $J^*(i_1) \subseteq J^*(i_2)$, a contradiction. *Q.E.D.*

Claim 4. *Suppose $i_1, i_2 \in N$ have $J^*(i_1) \neq J^*(i_2)$. Then $J^*(i_1) \cap J^*(i_2) = \emptyset$.*

Proof. Assume for a contradiction that $j^* \in J^*(i_1) \cap J^*(i_2)$. By Claim 2 and Claim 3, $S_+ := \{i \in N : \max J^*(i) > j^*\}$ cannot contain both i_1 and i_2; nor can $S_- := \{i \in N : \min J^*(i) < j^*\}$ contain both i_1 and i_2. But $J^*(i_1)$ and $J^*(i_2)$ cannot both be $\{j^*\}$, so that at least one of S_+, S_- contains one of i_1, i_2. We will derive a contradiction from $i_1 \in S_-$, the other three cases being completely analogous.

In this case, $i_2 \notin S_-$, so that $\min J^*(i_2) = j^*$. Moreover, by Claim 3, every $i \in S_-$ has $|J^*(i) \cap J^*(i_2)| \leq 1$, so that (since both $J^*(i), J^*(i_2)$ are intervals by Claim 2) $\max J^*(i) \leq j^*$.

Finally, observe that

$$\begin{aligned}
|S_-| - (j^* - 1) &= \left(\sum_{i \in S_-} \sum_{j \in N} \mu^*_{ij} \right) - \left(\sum_{j=1}^{j^*-1} \sum_{i \in N} \mu^*_{ij} \right) \\
&= \left(\sum_{i \in S_-} \sum_{j=1}^{j^*} \mu^*_{ij} \right) - \left(\sum_{j=1}^{j^*-1} \sum_{i \in S_-} \mu^*_{ij} \right) \\
&= \sum_{i \in S_-} \mu^*_{ij^*} \in [\mu^*_{i_1 j^*},\ 1 - \mu^*_{i_2 j^*}],
\end{aligned}$$

where the second equality follows from the dropped entries of μ^* being zero. Hence, it follows that $|S_-| - (j^* - 1) \subseteq (0, 1)$, which contradicts $|S_-| - (j^* - 1)$ being an integer. *Q.E.D.*

Now, we can define a weak order $\succsim$ on N by saying $i \sim i'$ if and only if $J^*(i) = J^*(i')$, and $i \succ i'$ if and only if $\max J^*(i) < \min J^*(i')$. The relation is obviously transitive, and it is complete by Claim 2 and Claim 4. By construction of $\succsim$, under μ^* there is complete rank uncertainty over any $\sim$ equivalence class (because for any $i \sim j$ we have $\mu^*_{ij} > 0$), and i is ranked above i' in any optimal ranking matrix if $i \succ i'$.

All that remains is to check that $c_i < c_{i'}$ whenever $i \succ i'$. Assume otherwise for a contradiction: so $i \succ i'$ but $c_i \geq c_{i'}$. Observe that the incentive effects generated by μ^* satisfy $I_i < I_{i'}$. As $c_i \geq c_{i'}$, switching the i and i' rows from μ^* would weakly improve the principal's objective (given that $f_i(\mu)$ is a submodular function of I_i and c_i), preserving optimality. But the new ranking matrix would have a nonzero entry where μ^* does not, a contradiction.

B.2. Proof of Proposition 2

In what follows, let $\succsim$ be as given by Proposition 1. The proof of Proposition 2 proceeds in three steps. First, we provide two preliminary results that we will use in our arguments. Second, we explicitly characterize the form of the order $\succsim$. Finally, we specialize this characterization to understand when it is perfectly coarse or perfectly fine.

Step 1. We provide two preliminary results. First, in the next claim, we show that the ordering induced by $\left\{\frac{I_i}{\sqrt{c_i}}\right\}_{i \in N}$ at an optimal ranking matrix respects the order $\succsim$; this is an expression of the principal's first-order conditions.

Claim 5. *If $i \succsim i'$, then any optimal ranking matrix induces $\frac{I_i}{\sqrt{c_i}} \geq \frac{I_{i'}}{\sqrt{c_{i'}}}$.*

Proof. Let μ be an optimal ranking matrix, and take $i \succsim i'$ with $i \neq i'$. Given the unique-bonuses result in Theorem 2, the incentive effects $\{I_{\tilde{i}}\}_{\tilde{i} \in N}$ remain unchanged if we replace μ with any other optimal ranking matrix. Therefore, by Proposition 1, we may assume without loss that μ exhibits complete rank uncertainty over every $\sim$ equivalence class. A consequence is that $\mu_{ii}, \mu_{i'i'} > 0$, implying that $\mu + \varepsilon D(i, i', i', i)$ is a ranking matrix for small enough $\varepsilon > 0$. Optimality of μ then requires that the directional derivative of the principal's objective in direction $D(i, i', i', i)$ be nonnegative. By direct computation, this derivative is equal to

$$P(N)\left\{[P(i') - P(i'-1)] - [P(i) - P(i-1)]\right\}\left[\frac{c_{i'}}{I_{i'}(\mu)^2} - \frac{c_i}{I_i(\mu)^2}\right],$$

which (by strict supermodularity of P and because $i < i'$) has the same sign as $\frac{c_{i'}}{I_{i'}(\mu)^2} - \frac{c_i}{I_i(\mu)^2}$. Therefore, $\frac{c_{i'}}{I_{i'}(\mu)^2} \geq \frac{c_i}{I_i(\mu)^2}$, as desired. *Q.E.D.*

Second, we establish the following result comparing the incentive effects given by two ranking matrices.

Claim 6. *Take two ranking matrices μ', μ'' and $k_1, k_2, k_3 \in N$ with $k_1 \leq k_2 < k_3$. Suppose that:*

1. $\forall i \in \{k_1, \ldots, k_3\},\ j \notin \{k_1, \ldots, k_3\}$ *we have* $\mu''_{ij} = 0$;
2. $\forall i \in \{k_1, \ldots, k_2\},\ j \notin \{k_1, \ldots, k_2\}$ *we have* $\mu'_{ij} = 0$;
3. $\exists i \in \{k_1, \ldots, k_2\},\ j \in \{k_2, \ldots, k_3\}$ *such that* $\mu''_{ij} > 0$.

Then, $\displaystyle\sum_{i=k_1}^{k_2} I_i(\mu') < \sum_{i=k_1}^{k_2} I_i(\mu'')$.

Proof. It is straightforward to construct a new ranking matrix $\tilde{\mu}$, also with property 2, such that $\tilde{\mu}_{ij} \geq \mu''_{ij}\ \forall i, j \in \{k_1, \ldots, k_2\}$. Property 2, and the fact that both $\tilde{\mu}$ and μ' are doubly stochastic, implies $\displaystyle\sum_{i=k_1}^{k_2} I_i(\mu') = \sum_{i=k_1}^{k_2} I_i(\tilde{\mu})$. Notice that $\tilde{\mu}_i$ is weakly first-order-stochastically dominated by μ''_i for each $i \in \{k_1, \ldots, k_2\}$ by construction and property 1. Also notice that at least one such dominance relationship holds strictly because of property 3. Since $P(j) - P(j-1)$ is strictly increasing in $j \in N$, we have that $I_i(\tilde{\mu}) \leq I_i(\mu'')\ \forall i \in \{k_1, \ldots, k_2\}$ with at least one strict inequality. This means that $\displaystyle\sum_{i=k_1}^{k_2} I_i(\tilde{\mu}) < \sum_{i=k_1}^{k_2} I_i(\mu'')$. *Q.E.D.*

Step 2. Having shown that the principal's first-order conditions take the simple form in Claim 5, we can convert this result into a complete characterization of the order $\succsim$. To achieve this characterization, Claim 7 and Claim 8 below derive concrete algebraic conditions on $\succsim$ that follow from these first-order conditions, and then Claim 9 shows that no two distinct orders can satisfy the same concrete conditions. We provide the complete characterization of $\succsim$ in Claim 10.

In what follows, the following function on sets of agents will be of use.

Notation 3. Given any nonempty set $S \subseteq N$, let

$$\varphi(S) := \frac{\sum_{j \in S}[P(j) - P(j-1)]}{\sum_{i \in S} \sqrt{c_i}}.$$

Remark 2. An important property φ satisfies, which is easy to establish given its "fractional sum" form, is a (strict) betweenness property. Specifically, any collection $\mathcal{S} \subseteq 2^N$ of pairwise disjoint, nonempty sets has

$$\min_{S \in \mathcal{S}} \varphi(S) \leq \varphi\left(\bigcup \mathcal{S}\right) \leq \max_{S \in \mathcal{S}} \varphi(S),$$

with both inequalities strict if $\{\varphi(S)\}_{S \in \mathcal{S}}$ are not all the same. We take this property for granted throughout the proof.

Claim 7. *Let S be some $\sim$ equivalence class. Any optimal ranking matrix μ has $I_i(\mu) = \varphi(S)\sqrt{c_i}$ for every $i \in S$.*

Proof. By Claim 5, there is some $\overline{I} \in \mathbb{R}$ such that $I_i(\mu) = \overline{I}\sqrt{c_i}$ for every $i \in S$. But, by Proposition 1, we know that $\mu_{ii'} = \mu_{i'i} = 0$ for every $i \in S$ and $i' \in N \setminus S$. Therefore,

$$\overline{I} \sum_{i \in S} \sqrt{c_i} = \sum_{i \in S} I_i(\mu) = \sum_{j \in S}[P(j) - P(j-1)] = \varphi(S) \sum_{i \in S} \sqrt{c_i},$$

implying $\overline{I} = \varphi(S)$. *Q.E.D.*

Claim 8. *Let S be some $\sim$ equivalence class, and $S' := [1, i'] \cap S$ for some $i' \in S$ with $i' < \max S$. Then $\varphi(S') < \varphi(S)$.*

Proof. Proposition 1 yields some optimal ranking matrix μ such that $\mu_{ij} > 0$ for every $i, j \in S$. Claim 7 says each $i \in S$ has $\varphi(S)\sqrt{c_i} = I_i(\mu)$, so that

$$\frac{\varphi(S')}{\varphi(S)} = \frac{\sum_{j\in S'}[P(j) - P(j-1)]}{\varphi(S)\sum_{i\in S'}\sqrt{c_i}} = \frac{\sum_{i\in S'} I_i(\delta)}{\sum_{i\in S'} I_i(\mu)},$$

where δ is the identity matrix. By Claim 6, this ratio is strictly below 1. *Q.E.D.*

Claim 9. *Suppose weak orders $\succsim_1, \succsim_2$ are such that, for both $k \in \{1, 2\}$:*

1. $1 \succsim_k \cdots \succsim_k N$*;*
2. *Every pair S, S' of $\sim_k$ equivalence classes with $S \ll S'$ have $\varphi(S) \geq \varphi(S')$;*
3. *Every $\sim_k$ equivalence class S and $i' \in S\backslash\{\max S\}$ have $\varphi(S\cap[1, i']) < \varphi(S)$.*

Then $\succsim_1$ and $\succsim_2$ are identical.

Proof. Assume for a contradiction that the claim fails. Given the first condition above, both $\succsim_1, \succsim_2$ are fully determined by their equivalence classes. So let $i_0 \in N$ be the lowest-labeled agent such that $S_1 := \{i \in N : \ i \sim_1 i_0\} \neq \{i \in N : i \sim_2 i_0\} =: S_2$. By construction, $i_0 = \min S_1 = \min S_2$, and the first condition implies that both S_1 and S_2 are intervals in N. Without loss, say $S_2 \not\subseteq S_1$, so that $\max S_2 > \max S_1$. Now, define $\mathcal{S}_1$ to be the set of all $\sim_1$ equivalence classes contained in $[i_0, \max S_2)$, and let $S_1' := \{i \in N : \ i \sim_1 \max S_2\}$.

Note that as S_2 is the disjoint union of $\mathcal{S}_1 \cup \{S_1' \cap S_2\}$, the betweenness property says

$$\begin{aligned}
\varphi(S_2) &\leq \max\{\varphi(S) : \ S \in \mathcal{S}_1 \text{ or } S = S_1' \cap S_2\} \\
&\leq \max\{\varphi(S) : \ S \in \mathcal{S}_1 \text{ or } S = S_1'\} \\
&= \varphi(S_1).
\end{aligned}$$

The second inequality follows from applying the third property to $\succsim_1$ because $S_1' \cap S_2$ is an initial segment of S_1'. Moreover, the equality follows from noting that $S \ll S_1'$ for every $S \in \mathcal{S}_1$ and applying the second property to $\succsim_1$. But applying the third property to $\succsim_2$ delivers $\varphi(S_1) < \varphi(S_2)$, a contradiction. *Q.E.D.*

Claim 10. *The weak order $\succsim$ is the unique transitive complete relation satisfying:*

1. $1 \succsim \cdots \succsim N$;
2. *Every pair S, S' of $\sim$ equivalence classes with $S \ll S'$ have $\varphi(S) \geq \varphi(S')$;*
3. *Every $\sim$ equivalence class S and $i' \in S \setminus \{\max S\}$ have $\varphi(S \cap [1, i']) < \varphi(S)$.*

Proof. By definition, $\succsim$ satisfies the first property. It satisfies the second property by Claim 5 and Claim 7. It satisfies the third property by Claim 8. But then, no other order can satisfy these three properties by Claim 9. *Q.E.D.*

Step 3. With the above claims in hand, the proposition is easy to establish. To see the second part of the proposition, observe that there is an optimal ranking matrix exhibiting complete rank uncertainty over the whole set N of agents if and only if $1 \sim \cdots \sim N$, which Claim 10 shows holds if and only if $\varphi(\{1, \ldots, n\}) < \varphi(N)$ for every $n \in \{1, \ldots, N-1\}$.

To see the first part of the proposition, observe that the identity (ranking) matrix $\delta = [\delta_{ij}]_{i,j \in N}$ is the unique optimal ranking matrix if and only if $1 \succ \cdots \succ N$, which Claim 10 shows holds if and only if $\varphi(\{1\}) \geq \cdots \geq \varphi(\{N\})$. All that remains, then, is to establish that δ is an optimal ranking matrix if and only if it is the uniquely optimal ranking matrix. But this result follows directly from the following claim (taking $\succsim'$ to satisfy $1 \succ' \cdots \succ' N$).

Claim 11. *Suppose $\succsim'$ is a weak order on N such that some optimal ranking matrix has i ranked higher than i' for every $i \succ' i'$. Then $\succsim$ is a (weak) refinement of $\succsim'$, i.e., $i \succ i'$ for any $i \succ' i'$. Therefore, $\succsim$ is the finest order on N with this property.*[30]

Proof. Let μ' be an optimal ranking matrix that has i ranked higher than i' for every $i \succ' i'$, and (appealing to Proposition 1) let μ be an optimal ranking matrix whose nonzero entries are exactly $\{(i,j) \in N^2 : \ i \sim j\}$. Fix $i^* \in N$ and let $S := \{i \in N : \ i \succsim' i^*\}$. It suffices to show that $i \succ i'$ for any $i \in S$ and $i' \in N \setminus S$. Assume otherwise, for a contradiction.

[30] That $\succsim$ as delivered by Proposition 1 is the finest order with this property implies it can be inferred directly from the set of nonzero entries of any optimal ranking matrix. This observation plays no role in the proof of the current proposition.

Observe that all nonzero $S \times N$ entries of μ' are in $S \times \{1, \ldots, |S|\}$, but that μ has at least one strictly positive entry in $S \times \{|S|+1, \ldots, N\}$. Claim 6 then implies $\sum_{i \in S} I_i(\mu') < \sum_{i \in S} I_i(\mu)$, contradicting the unique-bonuses result in Theorem 2. *Q.E.D.*

B.3. Proof of Proposition 3

In what follows, define $\varphi^* : N \to \mathbb{R}$ by letting $\varphi^*(i) := \varphi(\{i' \in N : i' \sim i\})$, where $\succsim$ is as given by Proposition 1. By Claim 7, the optimal bonus paid to each agent $i \in N$ is exactly $b_i^* = \frac{\sqrt{c_i}}{\varphi^*(i)}$, and so his markup is exactly $\frac{b_i^*}{c_i} = \frac{1}{\varphi^*(i)\sqrt{c_i}}$. It therefore suffices to show that $\varphi^*(i)$ weakly decreases as any agent's marginal cost increases, and that $\varphi^*(i)\sqrt{c_i}$ weakly increases as c_i increases.

To see the above, first observe that any set $S \subseteq N$ which contains agent i has

$$\varphi(S) = \frac{1}{\sum_{i' \in S} \sqrt{c_{i'}}} \sum_{j \in S} [P(j) - P(j-1)],$$

which weakly decreases with the vector of marginal costs, and

$$\varphi(S)\sqrt{c_i} = \frac{\sqrt{c_i}}{\sum_{i' \in S} \sqrt{c_{i'}}} \sum_{j \in S} [P(j) - P(j-1)],$$

which weakly increases with agent i's marginal cost.

Therefore, the proposition will follow directly if $\varphi^*(i)$ is an increasing function of the vector $(\varphi(S))_{S \subseteq N : i \in S}$. But this fact follows directly from the following claim.

Claim 12. *Each agent $i \in N$ has* $\varphi^*(i) = \max_{i_1 \in \{i, \ldots, N\}} \min_{i_0 \in \{1, \ldots, i\}} \varphi(\{i_0, \ldots, i_1\})$.

Proof. As $\max_{i_1 \in \{i, \ldots, N\}} \min_{i_0 \in \{1, \ldots, i\}} \varphi(\{i_0, \ldots, i_1\})$ is always weakly below $\min_{i_0 \in \{1, \ldots, i\}} \max_{i_1 \in \{i, \ldots, N\}} \varphi(\{i_0, \ldots, i_1\})$, it suffices to show that

$$\min_{i_0 \in \{1, \ldots, i\}} \max_{i_1 \in \{i, \ldots, N\}} \varphi(\{i_0, \ldots, i_1\}) \leq \varphi^*(i) \leq \max_{i_1 \in \{i, \ldots, N\}} \min_{i_0 \in \{1, \ldots, i\}} \varphi(\{i_0, \ldots, i_1\}).$$

Let us establish that $\min_{i_0 \in \{1, \ldots, i\}} \max_{i_1 \in \{i, \ldots, N\}} \varphi(\{i_0, \ldots, i_1\}) \leq \varphi^*(i)$, the other inequality following by a symmetric argument.[31] Toward establishing the inequality, let $i_0^* := \min\{i_0 \in N:\ i_0 \sim i\}$ and $i_1^* := \max\{i_1 \in N:\ i_1 \sim i\}$, and take an arbitrary $i_1 \in \{i, \ldots, N\}$. We aim to show that $\varphi^*(i) \geq \varphi\left(\{i_0^*, \ldots, i_1\}\right)$, i.e., that $\varphi\left(\{i_0^*, \ldots, i_1^*\}\right) \geq \varphi\left(\{i_0^*, \ldots, i_1\}\right)$.

To accomplish this, we provide an alternative characterization of $\succsim$. Let $\ell_0 := 0$ and, working recursively for $k \in N$, define

$$\ell_k := \begin{cases} N & : \ \ell_{k-1} = N, \\ \min \arg\max_{i' \in \{\ell_{k-1}+1, \ldots, N\}} \varphi\left(\{\ell_{k-1}+1, \ldots, i'\}\right) & : \ \ell_{k-1} < N. \end{cases}$$

Letting $S^* := \{\ell_k\}_{k \in N}$, we can then define the order $\succsim'$ on N by letting $i' \succsim' i''$ if and only if $\min\left(\{i', \ldots, N\} \cap S^*\right) \leq \min\left(\{i'', \ldots, N\} \cap S^*\right)$. It is easy to see that this order satisfies the three properties listed in Claim 10: the first and third are immediate, and the second follows from applying betweenness of φ to the union of any two adjacent $\sim'$ equivalence classes. Claim 10 then implies that $\succsim'$ is exactly $\succsim$. But then, $i_0^* - 1 = \ell_{k-1}$ for some $k \in N$. It follows by construction that $i_1^* = \ell_k$ maximizes $\varphi(\{i_0^*, \ldots, i'\})$ over $i' \in \{i_0^*, \ldots, N\}$, delivering the desired inequality. *Q.E.D.*

[31] To see that symmetry obtains, note that the only conditions on which we base our arguments below—betweenness of φ, and the conditions established in Claim 10—would apply directly if we were to replace $\succsim$ with $\precsim$ and φ with $-\varphi$. In particular, by strict betweenness, the third condition in Claim 10 is equivalent to requiring that every $\sim$ equivalence class S and $i' \in S \setminus \{\min S\}$ have $\varphi(S \cap [i', N]) > \varphi(S)$.