

Incentives and Discrimination¹

Eyal Winter²

ABSTRACT: Optimal incentive mechanisms may require that agents be rewarded differentially even when they are completely identical and induced to act the same. We demonstrate this point using a simple incentive model where agents' decisions about effort exertion are mapped into a probability that the project will succeed. We show that full discrimination across all agents is required if and only if the technology has increasing returns to scale, and discuss the role of hierarchies in generating optimal incentives. (*JEL C72, D23, D78, J31, M52*)

1. Introduction

The tension between equality and efficiency is often a key issue in debating incentive schemes in organizations. The notion that benefits should be assigned to individuals in a non-uniform manner that takes into account qualifications and performance is well established. Yet it is sometimes claimed that favoritism and internal politics often lead to differential rewards even when individuals hardly differ in their attributes, a phenomenon which is regarded as counter-efficient. The purpose of this paper is to argue that, from the point of view of optimal incentives, differential rewards may be unavoidable even when individuals are completely identical and when the mechanism aims at inducing all agents to exert effort.

The fact that optimal incentive mechanisms may require non-symmetric rewards even when agents are identical in their qualifications may not be surprising in some environments. If agents, for example, are asymmetrically informed about each other's exertion of effort, it may require different levels of incentives to induce each of them to exert effort. In this case optimal mechanisms are expected to yield non-symmetric rewards that depend on the information that agents have about each other (see Winter, (2001) for a full analysis of such a model). Another case in which non-symmetric

¹ Helpful comments and suggestions by Douglas Bernheim, three anonymous referees and seminar participants at Berkeley, CORE, EUI Florence, Cambridge, Essex, UCL, Jerusalem and Tel Aviv are gratefully acknowledged.

² The Hebrew University, Center for Rationality, Jerusalem 91904, Israel

mechanisms seem trivially unavoidable is when the principal's objective is to induce some but not all agents to exert effort (for example, when contributions by only a subset of agents are sufficient to guarantee the project's success). In this case it will be wasteful to reward all agents equally. Promising positive rewards (contingent on the success of the project) to some and zero to others seems to be the optimal incentive scheme.

What seems to be more surprising is that optimal mechanisms may have to be discriminatory even when all agents are completely identical (in terms of all their characteristics including information) and when the objective of the mechanism is to induce all of the agents to exert effort. We demonstrate our point with a simple model of organization similar to the one used in Winter (2001) in the context of optimal allocation of responsibility.

A project is managed by n agents each of which is responsible for a different task. If an agent exerts effort in performing his task he increases the probability that his task will end successfully from α to 1 at cost c , which is constant across agents. The overall project succeeds only when all tasks end successfully. Neither the principal nor the agents themselves can observe each other's effort. Therefore, the mechanism must reward agents only as a function of whether the project ends successfully or not. An optimal mechanism induces agents to exert effort as a unique Nash equilibrium, and it does so at a minimum total reward.

Our first observation in Proposition 1 is that any optimal mechanism must treat all agents differently. The intuition behind the result is quite simple. If agents' exertion of effort induces a positive externality on the effectiveness of other agents' effort, it is optimal to promise high rewards to some agents so as to make the others confidently believe that these highly paid agents will contribute, hence allowing the planner to save resources by offering other agents substantially less. Invoking this argument iteratively shows that when the project technology involves increasing returns to scale, no two agents should earn the same reward in spite of the fact that all agents are identical. Interestingly, the property of increasing returns to scale is not only a sufficient condition for full discrimination but also a necessary condition.

Optimal mechanisms in our framework give rise to an endogenous hierarchy of incentives. One agent is induced to exert effort regardless of his beliefs. Another agent is

paid enough to make him exert effort if he believes that a least one other agent does so as well, etc. We will use this property to argue that ranks and hierarchies in organizations where authority plays little role can be thought of instruments for generating optimal incentives.

Our results demonstrate the tension between efficiency and equity in the context of incentive mechanisms. Such tension also exists in other economic contexts — most prominently in the public finance literature. Atkinson and Stiglitz's (1976) seminal paper points out that optimal taxation might entail horizontal inequity; i.e., individuals who, for all relevant purposes, are identical may be treated differently. Such inequity is driven by the non-convexity of the set of allocations on which the social welfare function is maximized. Later Stiglitz (1982) argued that horizontal inequity may even be inconsistent with Pareto optimality and showed that random taxation may yield a Pareto improvement. Stiglitz demonstrated the phenomenon by recalling the well-known story of the two shipwrecked sailors with sufficient food only for one — with horizontal equity they are both doomed to death. As we shall see later the role of inequality in our framework is different as it relies on incentives and externalities.

A different nature of inequality — closer to ours — has been discussed in some IO papers dealing with trade contracts. The intuition behind the role of discrimination in our mechanism is related to the “divide and conquer” strategy that has been discussed in the literature on exclusionary contracts (see Rasmusen, Ramseyer and Wiley (1991), Innes and Sexton (1994), Segal and Whinston (2000)). A monopolist attempting to deter the entry of a potential rival can do so by signing exclusionary contracts with buyers. Since signing such contracts with some will make other buyers more willing to join, the monopolist can effectively discriminate among the buyers. Somewhat related is also the argument about introductory prices by a monopolist producing goods with positive consumption externalities such as in Bensaïd and Lesne (1996) and Cabral, Salant and Woroch (1999) (see also Farrell and Saloner (1985) and Katz and Shapiro (1986)). More recently Segal (2001) introduced a general model of trade contracts when the principal's trade with one agent generates externalities on others. While Segal's main interest is the volume of aggregate trade, he also points out that with increasing externalities the principal gains by using a divide and conquer strategy. Segal's model fits nicely into a

variety of IO applications (like takeovers, vertical contracting, exclusive dealing, and network externalities) but it does not apply to our model of organizational incentives. This is because agents' actions in our framework are not contractible. Players' payoffs depend only on the final outcome of the project (and not on the actual "trade" chosen by the agent). This is a crucial assumption in our framework. If the principal could monitor the agents' actions the optimal mechanism would instruct him to reimburse the cost of each contributing agent and there would be no role for discrimination.

There are several other differences between our framework and that of the literature cited above. In contrast to the latter we do not only demonstrate the gains from discrimination but also show their limits by identifying the precise conditions under which discrimination is optimal. We distinguish between discrimination and full discrimination (which separates all agents) and provide necessary and sufficient conditions for both. In particular discrimination in our framework applies under a conceptually weaker condition than in the above literature (see Proposition 2). The full characterization of discrimination is particularly important if one wants to draw testable implications from the model.

While the role of discrimination has been discussed in the contexts of optimal taxation and various IO applications its role in the context of organizational incentives has hardly been mentioned.³ This is a sharp contrast to the considerable attention that has been devoted to the issue of inequality within organizations. A large number of models in personnel economics (some of which will be surveyed later) establishes that unequal treatment of unequal agents may have major incentive advantages. The particular importance of demonstrating the optimality of treating *equals* unequally is that it potentially implies an additional gain for inequality in each of these models, offering a more comprehensive explanation for the various phenomena they address. We will expand on this point later.

We introduce the benchmark model in the following section. In Section 3 we provide a necessary and sufficient condition for discrimination within the general model. In Section 4 we show the equivalence between full discrimination and increasing returns

³ We are interested in incentive-based discrimination — and not market discrimination which is based on gender or race and on which extensive literature exists.

to scale and discuss the implication of this result. Section 5 demonstrates that if agents can coordinate their strategies, then the principal can sustain investment optimally without resorting to discrimination. This is done by focusing on the concepts of Strong Equilibria and Coalition-Proof Equilibria for our implementation. We conclude in Section 6.

2. The Model

The organizational project involves n tasks each performed by a different individual (agent). Each agent has to decide whether to exert effort towards the performance of his task or not. The cost of effort is c and is constant across all players. We henceforth use the term investment for the action of exerting effort.

If an agent invests, his task ends successfully with probability 1. If he doesn't invest, the probability of his task ending successfully is only α , which is again constant across all agents. Agents are not informed about each other's investment decisions. Thus these decisions will be modeled as if made simultaneously.

For our benchmark model we will assume that the project as a whole ends successfully if and only if all the tasks are performed successfully. We will later consider a general class of technologies that do not necessarily have this "O-Ring"⁴ property.

In keeping with Holmstrom (1982) agents' effort decisions are unobservable, and therefore agents' rewards in the mechanism depend only on whether the project ends successfully or not. Specifically, if the project fails all agents receive a zero reward, but if it succeeds they are paid the rewards $v = (v_1, \dots, v_n)$. We identify here the vector v with the incentive mechanism. Every mechanism v gives rise to a normal form game $G(v)$ in which each player i has two strategies: $d_i = 1$ for investment and $d_i = 0$ for non-investment. For a strategy combination $d = (d_1, \dots, d_n) \in \{0, 1\}^n$, the payoff function for player i is given as follows: $f_i(d) = v_i \alpha^{d(-i)} - c$ if $d_i = 1$ and $f_i(d) = v_i \alpha^{d(-i)+1}$ if $d_i = 0$, where $d(-i) = |\{j \neq i \mid d_j = 0\}|$ is the number of individuals (other than i) who choose to shirk.

We say that a mechanism v is incentive-inducing (INI) if v induces all players to invest in every equilibrium, i.e., $d = (1, 1, \dots, 1)$ is the only Nash equilibrium of the game

$G(v)$. We will say that a mechanism v is an optimal *INI* if it minimizes the total reward among all *INI* mechanisms.⁵

Proposition 1 claims that in any optimal *INI* mechanism no two players are rewarded equally. More specifically, one agent will be paid sufficiently well so as to make investment a dominant strategy for him. A second player can now be induced to invest with a smaller payment knowing that the first one will always invest. Continuing in this manner, the last agent is paid the amount that will make him invest if he believes that all other agents are investing as well.

For a permutation θ on the set of agents, and a vector x , we denote by $\theta(x)$ the vector with $\theta_i(x) = x_{\theta(i)}$.

Proposition 1: Let $v^* = (c/(1-\alpha), c/\alpha(1-\alpha), \dots, c/\alpha^{n-1}(1-\alpha))$. A mechanism v is an optimal *INI* mechanism if and only if $v = \theta(v^*)$ for some permutation θ .

Proof: We first note that $\theta(v^*)$ is an *INI* mechanism. Since all agents are symmetric we will denote by $v(k)$ the reward that would make an agent indifferent between investing and shirking given that he believes that exactly k other agents are investing, where $0 \leq k < n$. Note that if i chooses to invest his expected payoff is $v(k)\alpha^{n-k-1} - c$, whereas if he chooses not to invest the expected payoff is $v(k)\alpha^{n-k}$. Hence, $v(k)$ satisfies $v(k)\alpha^{n-k-1} - c = v(k)\alpha^{n-k}$, or $v(k) = c/\alpha^{n-k-1}(1-\alpha)$. We can therefore set $v^* = (v(n-1), v(n-2), \dots, v(0))$. Note also that $v(k) > v(k+1)$ and in particular $\theta_i(v^*) \geq v(n-1)$, which means that $d = (1, \dots, 1)$ is a Nash equilibrium of $G(\theta(v^*))$. Hence, to show that $\theta(v^*)$ is an *INI* mechanism, it is sufficient to show that no equilibrium exists in which some group of agents shirks if we increase rewards by an arbitrarily small amount. Consider a strategy combination in which exactly k agents choose to invest where $0 \leq k < n$. Consider the players who are assigned the rewards $v(0), \dots, v(k)$. Any arbitrarily small increase in these rewards will

⁴ See Kremer (1993)

⁵ We will also allow the reward-minimizing mechanism not to be *INI* in itself. The formal definition should be the following: v is an optimal *INI* mechanism if (1) there exists no *INI* mechanism with less total reward and (2) for any $\varepsilon > 0$, $\{v_i + \varepsilon\}_{i \in N}$ is an *INI* mechanism. This technical caveat is innocent and is needed because rewards take continuous values.

make each of these players better off investing if one believes that k other players are investing as well. Hence, by an arbitrarily small increase of rewards beyond $\theta(v^*)$ we get $d = (1, \dots, 1)$ as the unique equilibrium. We now have to show that $\theta(v^*)$ is optimal. We assume without loss of generality that θ is the identity permutation. Consider a mechanism u such that $u_i < v_i^*$ for some players and $u_j = v_j^*$ for the rest. Let r be the largest index for which $u_r < v_r^*$; then there is an equilibrium of $G(u)$ in which players $1, 2, \dots, r$ shirk and $r+1, \dots, n$ invest and this equilibrium survives for a sufficiently small increase of the rewards beyond u . This equilibrium will cease to exist only if we increase the reward of one of the players in $1, 2, \dots, r$ to become at least v_r^* . Hence, no *INI* mechanism can have a total reward, which is less than $\sum_i v_i^*$.

3. General Success Technologies

We now extend the model described in Section 2 to show that the non-existence of symmetric optimal mechanisms holds in a more general framework. In doing so we will establish necessary and sufficient conditions for discrimination. We will define the project's technology as a function p from the set investment strategy profiles $\{0, 1\}^N$ to $[0, 1]$ specifying the probability of success for any given profile. Since our interest lies with the case in which all agents are identical, we will consider symmetric technologies, i.e., $p: \{0, 1, 2, \dots, n\} \rightarrow [0, 1]$ specifying the probability of the project's success as a function of the number of agents who choose to invest. We assume that extra investment always raises the probability of success, i.e., p is strictly increasing. Finally, as before, we assume that the principal is interested in inducing all agents to invest. The definition of an optimal *INI* mechanism remains the same.

A key property of p , which will be used later, is that of increasing returns to scale (*IRS*). We say that the technology p has *IRS* if $D(k) = p(k+1) - p(k)$ ($k=0, \dots, n-1$) is increasing in k . Note that in our benchmark model with $p(k) = \alpha^{n-k}$ satisfies this property. It is also interesting to note⁶ that any success technology can be expressed as representing a model of individual tasks where the probability of success for each task given

investment is not constant but rather a function $\beta(k)$ of the number of investing agents. Increasing returns to scale of β implies that p has *IRS* as well.⁷ We defer a more formal treatment of this issue to the Appendix.

While increasing returns to scale will play an important role later in our analysis, it turns out that the non-existence of symmetric optimal mechanisms holds under much weaker conditions. A mechanism v is symmetric if it assigns the same reward to all agents. Proposition 2 asserts that a necessary and sufficient condition for symmetry is that an agent's marginal contribution to the project's success is minimal when all other agents contribute as well. This condition implies a certain degree of substitution between the agents.

Proposition 2: A symmetric optimal *INI* mechanism exists if and only if $p(n) - p(n-1) \leq p(k+1) - p(k)$ for all $0 \leq k < n-1$.

Proof: Consider again the reward $v(k)$ for which an agent is indifferent between investing and shirking if he believes that exactly k other agents are investing. With a technology p a player's expected reward if he invests is $v(k)p(k+1) - c$, and with no investment it is $v(k)p(k)$. Thus $v(k)$ solves $p(k)v(k) = p(k+1)v(k) - c$ or $v(k) = c/[p(k+1)-p(k)]$. The condition in Proposition 2 implies that $v(n-1) \geq v(k)$ for all $k < n-1$. Consider now the mechanism $v = (v(n-1), \dots, v(n-1))$. For this mechanism $d = (1, \dots, 1)$ is a Nash equilibrium and for any arbitrary small increase of rewards, it is also the unique Nash equilibrium. Furthermore, if we decrease the reward to some agent, then $d = (1, \dots, 1)$ is no longer an equilibrium since this player is better off shirking. Hence, v is a symmetric optimal *INI* mechanism. We now show that the condition of the proposition is necessary: suppose by way of contradiction that $p(n-1) - p(n) > p(k+1) - p(k)$ for some $k < n-1$ and that a symmetric optimal *INI* exists in which $v_j \equiv u$. Clearly u must be one of the values $(v(0), v(1), v(2), \dots, v(n-1))$. Otherwise, there will be some slackness that will allow the

⁶ I'm grateful to a referee for pointing out this fact.

⁷ Interestingly, the monotonicity of β alone does not guarantee that p has increasing returns to scale (while constant β does imply that p has increasing returns to scale). If β is increasing but concave, then agents' incentives to invest may be declining with the number of contributors since the effect of my own investment on the success of *other* tasks (on which effort is exerted) is greater when the number of investors is small.

principal to reduce rewards. By the definition of $v(k)$ we have $v(n-1) < v(k)$ for some $k < n-1$. Let $k^* = \operatorname{argmax}_k v(k)$. We first assume that $u < v(k^*)$. Note that in order to sustain investment by all players we must have $u \geq v(n-1)$. Since $v(n-1) \leq u < v(k^*)$, there must exist some $k \leq n-2$ such that $v(k-1) \leq u < v(k)$. We now claim that under the mechanism u there exists a Nash equilibrium in which k agents invest and $n-k$ agents shirk. Indeed, no shirking player can profit by deviating to investment as investment requires a greater incentive given by the reward $v(k)$. Furthermore, no investing player will deviate by shirking because $v(k-1) \leq u$ and so any arbitrarily small increase of rewards results in agents preferring investment when $k-1$ others invest. Hence, u is not an *INI* mechanism as it yields equilibria in which some agents shirk. We now consider the case in which $u = v(k^*)$. Indeed, such a u is an *INI* mechanism. This is because $v(k^*)$ implies that $d = (1, \dots, 1)$ is an equilibrium and $v(k^*) \geq v(k)$ for all k implies that no other equilibrium exists. However, the mechanism $u^* = (v(n-1), v(k^*), \dots, v(k^*))$ is an *INI* mechanism as well. Since $v(k^*) > v(n-1)$ this mechanism involves a smaller reward for the first agent and the same for the rest. Hence, u is not an optimal *INI* and we obtain the desired contradiction.

4. Full Discrimination

In this section we characterize the technologies under which the optimal mechanism rewards all agents differently. An *INI* mechanism v is fully discriminating if $v_i \neq v_j$ for every pair of agents i, j . We will show the equivalence between the technology's *IRS* and the optimality of full discrimination, and will discuss the implication of this result.

Proposition 3: The technology p has increasing returns to scale if and only if all optimal *INI* mechanisms are fully discriminating.

Proof: If p has *IRS* then $v(k) = c/[p(k+1) - p(k)]$ and $v(k)$ is decreasing in k . Hence, using the same argument as in the proof of Proposition 1, we obtain that the optimal mechanisms are given by $\theta(v(0), \dots, v(n-1))$ where θ is some permutation of the set of agents. Hence, all mechanisms are fully discriminatory. We now show that if all optimal *INI* mechanisms are fully discriminating, then p must have *IRS*. As argued earlier, the

payoffs in an *INI* mechanism must involve only the values $v(0), \dots, v(n-1)$, so by assumption all these values are distinct. Consider an optimal *INI* mechanism and assume the following order of the values $v(k)$:

$v(k_0) > v(k_1) > \dots > v(k_{n-1})$ (where k_0, k_1, \dots, k_{n-1} is some permutation of $1, 2, \dots, n$). We first note that $v(k_{n-1}) = v(n-1)$. Otherwise, there is some k_j with $v(k_j) < v(n-1)$ and $d = (1, \dots, 1)$ cannot be a Nash equilibrium of the game because the player receiving $v(k_j)$ is better off deviating by shirking. We now establish by induction that $v(j) \geq v(k_j)$ for $j = 0, 1, \dots, n-2$. First, $v(k_0) \geq v(0)$ by the definition of $v(k_0)$ as the largest among the values.

Now assume by induction that $v(j) \leq v(k_j)$ for all $j \leq r-1 < n-2$ and consider $j = r$. Suppose by way of contradiction that $v(r) > v(k_r)$. We argue that the mechanism admits a Nash equilibrium in which r agents invest and $n-r$ shirk. Indeed, consider the set of agents whose payoffs are $v(k_0), \dots, v(k_{r-1})$. Call this set $R = \{1, \dots, r\}$. Consider the strategy combination in which agents in R invest and agents in $N \setminus R$ shirk. Since the agent receiving $v(k_{r-1})$ is indifferent between investing and shirking and $v(k_{r-1}) < v(k_j)$ for $j < r$ all agents in R are best responding (by investing) in the putative equilibrium. It is therefore sufficient to argue that when the agents in R invest, no other agent can increase his payoff by shifting from shirking to investing. But this follows from the fact that for every player j in $N \setminus R$ the reward v_j satisfies $v_j \leq v(k_r) < v(r)$. We thus obtained that $v(j) \leq v(k_j)$ for $0 \leq j \leq n-1$. But since the sets $\{v(k_j); 0 \leq j \leq n-1\}$ and $\{v(j); 0 \leq j \leq n-1\}$ are identical — these inequalities must imply equality. Hence, we have $v(0) > v(1) > v(2) > \dots > v(n-1)$. But as we argued at the beginning of the proof, the $v(j)$'s are the inverse of the $D(j) = p(j+1) - p(j)$, which is therefore increasing. Hence, p has *IRS*.

We have established that full discrimination is optimal when the organizational technology has increasing returns to scale, but the implication of this result goes far beyond the technology specifications. The only thing we need to establish full discrimination is that an agent's incentive to exert effort increases with the number of other agents who do so. This can be either a result of the technology's properties or simply a consequence of the psychology of peer pressure. We submit that this is a typical feature of workers' externalities in organizations. There are several interesting pieces of

empirical evidence to this effect. Ichino and Maggi (2000) show that prevalence of shirking within large Italian banks can be explained by the effect of peer pressure. Sacerdote (2001) finds that Dartmouth students study harder if their (randomly assigned) roommates do so as well. Recently, Falk and Ichino (2003) have even produced experimental evidence to this effect showing that subjects who were assigned to stuff letters in envelopes with remuneration independent of output worked harder if they believed their peers were working hard as well.

Our results imply that in any environment that exhibits externalities of the kind described above it should be optimal to provide agents with differential incentives. But just as these externalities may take different forms so does discrimination. We believe that in many organizations these differential incentives appear in the form of a hierarchy or in other types of rankings when authority plays little role and when agents of different levels deal with similar tasks. For our optimal mechanism to work it is not enough that agents are rewarded differentially, it is also necessary that this fact is commonly known. Appointing someone as “Project Head” of a firm or “Team Captain” of a sports team facilitates the common knowledge that there is at least one individual in the team whose stake in the mission is so great (in terms of either benefits or esteem) that she would choose to exert effort no matter what the rest are doing.⁸ More refined hierarchies can also lend themselves as instruments for generating different stakes in the organization’s success. Law and consulting firms distribute titles like, associate, senior associate, partner and principal. However, the tasks performed by individuals do not differ much and there is no clear authority structure that can justify the different titles. The same applies to various civil servants’ pay structures that assign grades and steps to workers who may often do the same job.

Extensive literature in personnel economics discusses the role of hierarchies and ranks in organizations. Notable are Calvo and Wellisz (1979) who show how internal supervision across hierarchy layers shapes wage scales within organizations, and Lazear and Rosen (1981) who view organizations as tournaments in which agents compete for higher ranks and thereby exert effort. The gains from hierarchies as established by this

⁸ In contrast to Hermalin (1998) the leadership role described here does not rely on asymmetric information between the leader and the rest of the team.

literature almost always rely on either authority relations or at least the ability to monitor individuals' performance. Our model assumes neither.⁹ We view ranks and hierarchies as instruments for generating differential incentives, i.e., agents who are assigned higher incentives are those placed at higher ranks.¹⁰ This does not diminish the validity of the others explanations offered by the literature discussed above, but it suggests an additional gain for hierarchies and indeed one that relies on minimal assumptions about the organizational structure and objectives.

5. Coordination

The lack of symmetric optimal mechanisms under increasing returns to scale is a consequence of mis-coordination between the agents. The mechanism that pays agents a uniform reward of $c/[p(n) - p(n-1)]$ (the smallest reward in the optimal discriminatory mechanism) admits one Nash equilibrium in which all agents exert effort. But there is another equilibrium in which no one exerts effort. If the principal has no way to coordinate agents to play the equilibrium that he prefers the most, he has to incur the extra cost of paying some agents more.

In more cooperative environments, where agents can coordinate their effort strategies, the principal can implement investment at smaller expense, and he can even achieve it with a symmetric mechanism. This is due to the fact that the possibility of agents coordinating joint deviations filters out the “bad” equilibria in which some agents shirk. We analyze this framework by adopting the very same model but assuming different solution concepts. We will implement investment via *strong equilibria* and

⁹ This is also where our model departs from standard principal-agent models (such as Holmstrom (1982) and Nalebuff and Stiglitz (1983)) which assume that the principal has access to some informative signals about individuals' output.

¹⁰ Adopting this interpretation of hierarchies, increasing returns to scale implies that each agent should occupy a different layer of the hierarchy. However, our model allows also for more realistic structures, where a single layer accommodates several agents. In particular, if the technology is such that agents' marginal contributions are constant for $k < n-1$ and increase at $k = n-1$, then the optimal structure is of two layers — the top consisting of a single agent (the “Project Head”) and the second accommodating all other agents.

¹¹ This of course in addition to other roles of hierarchies like authority (see Aghion and Tirole (1997)) or leadership (see Hermalin (1998)).

coalition-proof equilibria (see Bernheim, Peleg and Whinston (1987)). As will be shown both concepts lead to the same (symmetric) optimal mechanism.

Strong equilibrium requires that no subgroup of player can coordinate a joint deviation that will make all its members better off. More specifically, a strategy profile σ is a strong equilibrium if there exists no coalition of players S and a strategy profile $\sigma_S' = \{\sigma_i'\}_{i \in S}$ for that coalition such that all players in S are made better off by deviating to σ_S' assuming that players in $N \setminus S$ are still playing σ .

The coalition-proof concept sets weaker equilibrium conditions by imposing a stronger condition on permissible deviations. It requires that the agreement by the deviating coalition be itself self-enforcing.¹² The formal definition is recursive and is omitted here. However, for our purposes we need only note that to rule out a strategy profile as a coalition proof equilibrium it is enough to show one profitable (joint) deviation that constitutes a strong equilibrium in the game reduced to the set of deviating players (holding fixed the strategies of outsiders).

A mechanism v is an incentive-inducing mechanism with respect to strong (resp. coalition-proof) equilibria if $d = (I, \dots, I)$ is the unique strong (resp. coalition-proof) equilibrium of the investment game. Optimality is now defined in the same way as before.

We point out that the choice of solution concept for implementation should reflect the principal's assessment regarding the nature of interaction that takes place within the organizational environment. If the principal believes that the environment is sufficiently open and conducive to cooperation (which he can affect by offering recreational and social activities within the working place) then the concepts of strong equilibrium or coalition-proof equilibrium may be appropriate for implementation. Otherwise, only the standard Nash implementation is viable, requiring the extra burden of higher rewards and the necessity of discrimination.

Proposition 4: If the technology p is increasing,¹³ then the (symmetric) mechanism v

¹² I.e., the agreement constitutes a coalition-proof equilibrium in the reduced game among the coalition members, keeping the strategies of outsiders fixed.

¹³ Not necessarily increasing returns to scale.

with $v_j = c/[p(n) - p(n-1)]$ for all j is the unique optimal incentive-inducing mechanism with respect to both strong equilibria and coalition-proof equilibria.

Proof: We show the following: (1) with $v_j \equiv v = c/[p(n) - p(n-1)]$ the strategy profile $d = (1, \dots, 1)$ is a strong equilibrium (and therefore also a coalition-proof equilibrium). In this stipulated equilibrium a player earns $vp(n) - c$. If a group of agents S of size s chooses to shirk, each of its members will earn $vp(n-s)$, which is less than $vp(n) - c$ since $p(n) - p(n-s) > p(n) - p(n-1)$ (with equality for $s=1$). We now show that (2) $d = (1, \dots, 1)$ is the unique coalition-proof equilibrium (and thus the unique strong equilibrium). The inequalities we used in (1) show that any strategy profile in which s agents shirk has a profitable joint deviation for these players, which prescribes them all to invest.¹⁴ It is thus sufficient to argue that this joint deviation is self-enforcing. For this consider the reduced games involving the deviating players keeping the profile of outsiders fixed (all of them invest). In this reduced game the profile prescribing all agents to invest is a strong equilibrium (which again follows from the inequalities in (1)). Hence, the joint deviation is self-enforcing and $d = (1, \dots, 1)$ is the unique coalition-proof equilibrium. (1) and (2) show that v is investment-inducing with respect to both coalition-proof and strong equilibria. Finally, to show that v as specified above is optimal note that if some player's reward is less than $v = c/[p(n) - p(n-1)]$, then $d = (1, \dots, 1)$ is not even a Nash equilibrium (and therefore neither strong nor coalition-proof) since this player will choose to shirk when the rest invest getting $vp(n-1)$ instead of $vp(n) - c$.

6. Conclusion

Optimal incentive mechanisms may have to reward agents differentially even when they are identical and are induced to take the same action. This is shown by introducing a simple model of organization in which agents' effort decisions are mapped into a probability of the project's success. Furthermore, the technology's property of increasing returns to scale is both a sufficient and necessary condition for full discrimination. We have argued that ranks and hierarchies in organizations where authority plays little role

¹⁴ If $s = 1$, then we have indifference, but if rewards increase by an arbitrarily small increment, this indifference is broken also for $s=1$ (see footnote 2).

can be thought of as instruments for generating incentive-optimal discrimination. We restricted ourselves to symmetric environments in order to demonstrate the role of discrimination. However, to further support our intuition about hierarchies, it is instructive to consider the non-symmetric case. Our model admits two types of non-symmetric extension. One involves differential effort costs, and the other assumes differential probabilities of tasks' success (in the individual tasks model).¹⁵ The structure of the optimal mechanism here is similar to that of the symmetric case. However, in contrast to the symmetric case the assignment of incentives to agents is not arbitrary anymore. To understand how this assignment is set up, we will say that agent i is provided higher incentives than agent j if i finds it optimal to invest under any belief in which it is optimal for agent j to invest.

We can now summarize the observations in the non-symmetric framework as follows:¹⁶

(1) In both the differential cost model and the differential probability model the optimal mechanism is unique and negligible differences in attributes result with substantial differences in rewards.¹⁷

(2) With differential costs agents whose costs are lower (i.e., those who have higher skills) are assigned higher incentives. With differential probabilities, agents whose α_i 's are lower (i.e., those who are dealing with more important tasks) are assigned higher incentives.

Observation 1 assures us that the discrimination result in the symmetric case is generic and is not a consequence of some knife-edge phenomenon. But the implication of observation 2 is more substantial. If we associate higher incentives with higher ranks in organizations, as we did earlier in Section 3, then observation 2 tells us that the optimal mechanism assigns agents and tasks to different ranks in the “right” manner, i.e., agents of higher skills and tasks of higher importance are placed at higher ranks. Interestingly, Calvo and Wellisz (1979) reach similar conclusions using a completely different model. Observation 2, therefore, suggests that this descriptive pattern of organizational

¹⁵ In the differential cost model we assume that the technology p is symmetric and in the differential probabilities case we assume that effort costs are identical across agents.

¹⁶ The formal statements of these observations and their proofs are with the author.

¹⁷ A completely different model that gives rise to a similar feature is Rosen's (1981) model of superstars.

structures can emerge from the externalities among agents without assuming that individuals' outputs or efforts are observable.

We conclude with the following testable implication of our results: Controlled for other differences, organizations with strong peer effect (or where tasks are compliments) tend to be more hierarchical than ones where peer effect is small (or where tasks are substitutes). This interesting “task” is for a separate empirical study.

Appendix

Consider a more general model than the benchmark (individual task) model by assuming that a task is performed successfully with probability α if the agent shirks, and with probability $\beta(k)$ if he invests, where $\beta(k)$ is a function of the number of agents who invest.

Proposition 5: (1) Let $p: \{0, 1, 2, \dots, n\} \rightarrow [0, 1]$ be an increasing success technology. Then there exist probabilities $\beta(k)$, and α with $\beta(k)$ increasing in k and $\alpha < \beta(k)$ such that p is equivalent to the corresponding model of individual tasks. Furthermore any $\beta(k)$ and α satisfying these conditions give rise to an increasing success technology. (2) If p has *IRS*, then α and $\beta(k)$ also satisfy that $[\beta(k)/\alpha]^k$ has increasing return to scale. Furthermore, any $\alpha, \beta(k)$ satisfying these conditions give rise to an *IRS* technology p .

Proof: For a technology $p: \{0, 1, 2, \dots, n\} \rightarrow [0, 1]$ define $\alpha = [p(0)]^{1/n}$ and $\beta(k) = [p(k)/p(0)]^{1/k} [p(0)]^{1/n}$. Simple algebra yields $p(k) = [\beta(k)]^k \alpha^{n-k}$. Note that the RHS is precisely the probability that the project succeeds if k agents invest in the underlying individual task model. Now since $p(k)$ is increasing so is $p(k)/p(0)$ and therefore also $[p(k)/p(0)]^{1/k}$. Since $[p(0)]^{1/n}$ is constant we have $\beta(k)$ increasing. Also $\alpha < \beta(k)$ follows immediately from the fact that $p(k) > p(0)$ for all k . For the second part of (1) note that with $\beta(k)$ increasing $[\beta(k)/\alpha]^k$ is increasing and so also $p(k) = \alpha^n [\beta(k)/\alpha]^k$. To show 2 use the same definition for α and $\beta(k)$ and note that $\beta(k)/\alpha = [p(k)/p(0)]^{1/k}$ or $[\beta(k)/\alpha]^k = [p(k)/p(0)]$. Since $p(k)$ has *IRS* so has $[p(k)/p(0)]$. Finally, if α and $\beta(k)$ satisfy the conditions in 2, then the probability of project success given by $p(k) = [\beta(k)]^k \alpha^{n-k}$ has *IRS*.

Remark: It is worthwhile noting that while $\beta(k)$ being constant implies that $p(k)$ has *IRS*, $\beta(k)$ strictly increasing will not necessarily render $p(k)$ to have *IRS*. This follows from the fact that for an increasing function $f(x)$, the function $[f(x)]^x$ may not be convex. However, if $\beta(k)$ has *IRS*, then $p(k)$ *IRS* as well, since $[f(x)]^x$ is convex whenever $f(x)$ is convex (see also footnote 9).

References

Aghion, P. and J. Tirole (1997) "Formal and Real Authority in Organizations," *Journal of Political Economy* 105, 1-29.

Atkinson, A.B. and J. E. Stiglitz (1976) "The Design of Tax Structure: Direct versus Indirect Taxation," *Journal of Public Economics* 6, 55-75.

Bensaid, B., Lesne, J. P. (1996) "Dynamic Monopoly Pricing with Network Externalities," *International Journal of Industrial Organization* 14, (6) 837-855.

Bernheim, B. D., B. Peleg and M. D. Whinston (1987) "Coalition Proof Nash Equilibria I. Concepts," *Journal of Economic Theory* 42, 1-12.

Cabral, L.M., Salant, D.J. and Woroch, G. A. (1999) "Monopoly Pricing with Network Externalities," *International Journal of Industrial Organization* 17(2), 199-214

Calvo, G. and S. Wellisz (1979) "Hierarchy, Ability and Income Distribution," *Journal of Political Economy* 8, (5) 991-1010

Falk, A. and A. Ichino (2003) "Clean Evidence on Peer Pressure," Mimeo, The European University Institute.

Farrell, J. and G. Saloner (1985) "Standardization, Compatibility and Innovation," *Rand Journal of Economics* 16, 70-83.

Hermalin, B. H. (1998) "Toward an Economic Theory of Leadership: Leading by Example," *American Economic Review* 88, 1188-1206.

Holmstrom, B. (1982) "Moral Hazard in Teams," *Bell Journal of Economics* 13(2) 324-340.

Ichino, A. and G. Maggi (2000) "Work Environment and Individual Background Explaining Shirking Differentials in Large Italian Firms," *Quarterly Journal of Economics* 115, 1057-1090.

Innes, R. and R.J. Sexton (1994) "Strategic Buyers and Exclusionary Contracts," *American Economic Review* 84(3), 566-84.

Katz, M. and C. Shapiro (1986) "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy* 94, 822-841.

Kremer, M. (1993) "The O-Ring Theory of Economic Development," *Quarterly Journal of Economics* 108, 551-575.

Lazear, E. P., and S. Rosen (1981) "Rank-Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy* 89, 841-864

Nalebuff, B. and J. E. Stiglitz (1983) "Prizes and Incentives: Towards a General Theory of Compensation and Competition," *Bell Journal of Economics* 13, 21-43.

Rasmusen, E. B., J. M. Ramseyer, J.S. Wiley (1991) "Naked Exclusion," *American Economic Review* 81(5), 1137-45.

Rosen, S. (1981) "The Economics of Superstars," *American Economic Review* 71(5), 845-858

Sacerdote, B. (2001) "Peer Effects with Random Assignment: Results for Dartmouth Roommates," *Quarterly Journal of Economics* 116(2) 681-704

Segal, I and M. Whinston (2000) "Naked Exclusion: Comment," *American Economic Review* 90(1), 296-309

Segal, I. (2001) "Coordination and Discrimination with Externalities: Divide and Conquer?," forthcoming in the *Journal of Economic Theory*.

Stiglitz, J. E. (1982) "Utilitarianism and Horizontal Equity—The case of Random Taxation" *Journal of Public Economics* 18(1), 1-33.

Winter, E. (2001) "Scapegoats and Optimal Allocation of Responsibility," Discussion Paper #265, The Center for Rationality, The Hebrew University of Jerusalem.