NBER WORKING PAPER SERIES

CATEGORICAL REDISTRIBUTION IN WINNER-TAKE-ALL MARKETS

Roland G. Fryer, Jr.
Glenn C. Loury

Working Paper 10104
http://www.nber.org/papers/w10104

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2003

This paper was inspired by many stimulating discussions with Sherwin Rosen. We are grateful to Edward Glaeser, Jerry Green, and Bart Lipman for helpful comments and suggestions.The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Categorical Redistribution in Winner-Take-All Markets
Roland G. Fryer, Jr. and Glenn C. Loury
NBER Working Paper No. 10104
November 2003
JEL No. D3, D8, J4

ABSTRACT

This paper constructs a simple model of pair-wise tournament competition to investigate categorical redistribution in winner-take-all markets. We consider two forms of redistribution: category-sighted, where employers are allowed to use categorical information in pursuit of their redistributive goals; and category-blind, where they are not. It is shown that the equilibrium category-sighted redistribution scheme involves a constant handicap given to agents in the disadvantaged category. Equilibrium category-blind redistribution creates a unique semi-separating equilibrium in which a large pool of contestants exerts zero effort, and this pool is increasing in the aggressiveness of the redistribution goal.

Roland G. Fryer, Jr.
Department of Economics
Harvard University
Littauer Center, M-7
Cambridge, MA 02138
and NBER
rfryer@fas.harvard.edu

Glenn C. Loury
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
gloury@bu.edu

1 Introduction

Winner-take-all markets are pervasive in our society. Whom to hire, promote, admit, elect, or contract with are all determined in a tournament-like (winner-take-all) structure. In general, these markets tend to emerge when there is quality variation, but little price flexibility. As a result, prizes tend to be awarded on the basis of relative, not absolute, performance. Indeed, Best Sellers, World Cup champions, Harvard matriculants, Rhodes Scholars, first-round draft picks, clerks to a supreme court justice, cover girls, prime ministers, and Wimbeldon champions all have this feature in common.

Because of their structure, winner-take-all markets have the feature that small differences in quality can be associated with large differences in rewards, which makes it quite surprising that there has been no theoretical analysis of redistribution in these environments. Understanding the theoretical trade-offs involved in redistribution within winner-take-all markets is of great importance for public policy. For instance, affirmative action in college admissions is a form of categorical redistribution in a winner-take-all market.

In this paper, we analyze categorical redistribution in winner-take-all markets. A short synopsis of our approach is as follows. There are many employers and many more individuals seeking employment. Nature moves first and assigns a marginal cost of investment to each individual. Individuals observe their cost and choose a level of effort to exert in the contest. Nature, then, randomly assigns individuals to firms, where they are randomly matched to compete in pair-wise contests. Absent any redistributional goal, the individual in each match with the higher level of effort wins an exogenously determined prize.

We establish three main results. First, we solve explicitly for the equilibrium of a heterogeneous tournament model. In the unique equilibrium, an individual's behavioral strategy involves emitting effort as a function of the distribution of individuals with cost above his. This leads naturally to highly unequal outcomes between social groups endowed with different cost distributions and, hence, a demand for categorical (i.e., inter-group) redistribution. We go on to analyze category-sighted and category-blind redistribution policies. The distinction between these two forms of redistribution turns on how categorical information is used in the implementation of redistributive policies. We show that the unique equilibrium under category-sighted redistribution involves a constant handicap for agents in the disadvantaged category. This is similar, in particular, to employing a lower minimum standardized test threshold for admission into colleges and universities for applicants belonging to disadvantaged categories. When we impose the category-blind constraint, the nature of equilibrium handicapping changes drastically. Under this formulation, a non-trivial measure of individuals pool on a common low effort level and the prize is randomly distributed to members of the pool whenever they are matched against one another. We conclude by discussing the problem of optimal, centrally planned categorical redistribution, which draws on the impressive theory of optimal auction design.

This paper lies at the intersection of several literatures: tournament theory (Lazear and Rosen 1981, Green and Stokey 1983, Rosen 1986), optimal auction design (Myerson 1981), and income redistribution (Mirrlees 1971, Akerlof 1978). Our approach has little in common with the existing models in the tournament literature. These models were developed to investigate the economic efficiency of tournament play and to analyze tournaments as optimum labor contracts. In contrast, we focus on the implications of redistribution in tournament-like environments.

There is a strong relationship between tournaments and auctions, as many tournament models can be interpreted as all pay effort auctions. This insight is useful because the optimal auction design problem is well studied. The crucial difference lies in the distinction between centralized and decentralized planning. The optimal auction design literature studies the centralized problem: A planner designs auction procedures realizing that the structure of the auction will influence the behavior of the bidders. In a labor market, however, where many employers draw on a common pool of workers, it is more reasonable to think of de-centralized planning (i.e. each employer implements her policy while assuming that her decision will not affect applicants' incentives). Decentralized planning in the auction framework is an extremely unrealistic assumption, as it is rather like designing an auction mechanism while taking the distribution of bids as exogenous! So, while optimal categorical redistribution in a tournament model is closely related to the much studied optimal auction design problem, the problem of decentralized *equilibrium* redistribution (which is most pertinent in labor markets) has yet to be studied. This is the primary focus of our paper.

The paper is also related to the well studied problem of income redistribution. In the traditional optimal tax literature, individuals essentially pool their incomes and a central authority redistributes the pool back to individuals in an incentive compatible manner. This approach is quite different from our proposed framework. Categorical redistribution in winner-take-all contests imposes an additional restriction on redistribution by constraining the employer to redistribute utility (i.e. the probability of winning) using functionally irrelevant categorical identifiers in each match, but allowing employers to reach their redistribution goals by aggregating outcomes across all of their matches. In other words, when making a particular hiring decision, an employer must choose between a given slate of candidates, but she evaluates those candidates with an eye toward achieving sufficient categorical diversity across all of her hiring decisions.

The exposition proceeds as follows: section 2 presents and solves our pair-wise tournament model; section 3 introduces the notion of categorical redistribution and analyzes category-sighted and category-blind redistribution in a decentralized environment; section 4 draws parallels between optimal auctions and categorical redistribution in a centralized environment; section 5 concludes.

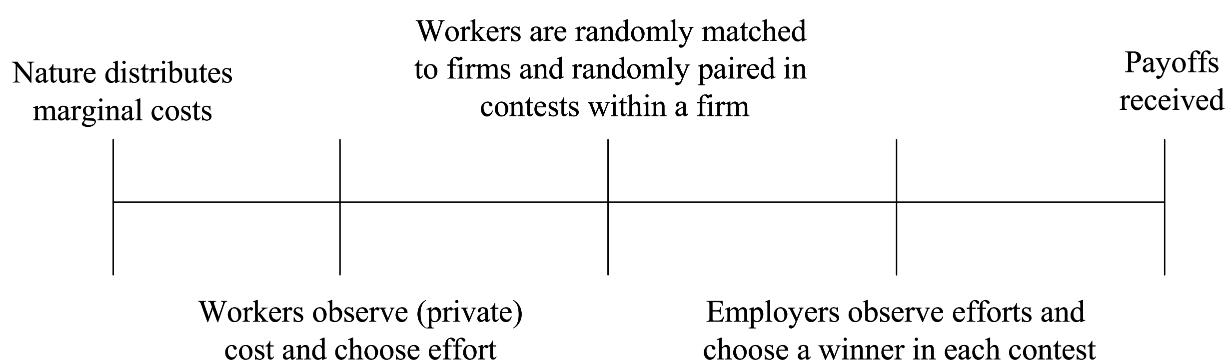

Figure 1: Sequence of Actions

2 Model: Pair-wise Tournament Competition

Consider a simple model of tournament competition. There is a continuum of workers with unit measure, and a large but finite number, N, of identical employers. Nature moves first and distributes a constant marginal cost of effort to each worker. This cost is distributed according to an atomless cumulative distribution function $F(c)$, where $c \in [c_{\min}, c^{\max}]$, $c_{\min} \equiv \inf\{c \mid F(c) > 0\} \geq 0$, and $c^{\max} \equiv \sup\{c \mid F(c) < 1\} \leq \infty$. Let $f : [c_{\min}, c^{\max}] \to \Re_+$ denote the associated probability density function.

After observing their cost, each worker chooses an effort level e. Each employer is then randomly matched with a continuum of workers of measure $\frac{1}{N}$, who in turn are randomly matched with one another to compete in a continuum of pair-wise contests for a measure $\frac{1}{2N}$ of positions. Thus, each employer faces a pool of workers that is a statistical replica of the overall worker population. Moreover, employers and workers anticipate that each worker will be paired for competition against an opponent drawn randomly from the overall population. We assume that each employer takes the workers' effort distribution as given, independent of her hiring policy, and chooses winners across these pair-wise contests so as to maximize the expected effort of those hired.[1] The worker hired from each pair receives an exogenously given wage ω, while the one not hired receives zero.

A strategy for workers is a function, $e : [c_{\min}, c^{\max}] \rightarrow \Re_{+}$, that maps their costs into an effort decision. A strategy for an employer is an assignment function in each pair-wise match, $A : \Re^{2}_{+} \rightarrow [0,1]$, that maps the effort levels she observes in each contest to a probability of winning for each worker in that contest. A worker presenting effort e wins against a worker presenting effort $\widehat{e}$ with a probability $A(e, \widehat{e})$. Thus, the payoff to a worker if he wins the contest is $\omega - ce$, while the payoff is $-ce$ if he loses.

A. Equilibrium

An equilibrium consists of functions $e^{*}(c)$ and $A^{*}(e, \widehat{e})$ such that each is a best response to the other.

Proposition 1 *The unique equilibrium consists of an assignment function satisfying*

$$A^{*}(e, \widehat{e}) = \begin{cases} 1 \text{ if } e > \widehat{e} \\ 0 \text{ if } e < \widehat{e} \end{cases}$$

and an effort supply function

$$e^{*}(c) = \omega \int_{c}^{c^{\max}} \frac{dF(y)}{y}, \text{ for all } c \in [c_{\min}, c^{\max}] \tag{1}$$

Proposition 1 provides a unique equilibrium behavioral strategy for workers that holds for general cost distributions, and (trivially) a unique equilibrium assignment function for employers.

[1] This, of course, means that employers simply hire the contestant with the greater effort, unless otherwise constrained. Section 4 discusses employer behavior under a categorical redistribution constraint.

In essence, the equilibrium effort supply for any worker depends (in the manner illustrated in equation 1) on the distribution of workers who have cost higher than his. Naturally, employers hire the worker in each match with the higher effort.

Let $G(e)$ denote the population cumulative distribution of effort in equilibrium, and let $v(c)$ denote the equilibrium net benefit to an agent with cost c. It follows directly that this equilibrium net benefit is given by

$$v(c) = \omega \int_{c}^{c^{\max}} \left[1 - \frac{c}{y}\right] dF(y); \tag{2}$$

and the equilibrium effort distribution satisfies

$$G(e) = 1 - F\left(e^{*^{-1}}(e)\right). \tag{3}$$

Notice that, taken together, equations (1) and (3) define a mapping from the exogenous cumulative distribution of costs, $F(c)$, to the cumulative distribution of effort in equilibrium, $G(e)$, in the natural way, which provides an explicit solution to the model.

To establish the proposition notice that, because workers choose effort to maximize their net benefit, we have the first-order condition: $\omega \frac{d}{de}[1 - G(e)]_{e=e^*(c)} = c$. A simple revealed preference argument establishes that $e^*(c)$ is non-increasing in c. Hence, equilibrium behavior implies that $[1 - G(e^*(c))] \equiv F(c)$. Putting these equations together, we have:

$$c = \omega \frac{d}{de}[G(e)]_{e=e^*(c)} = \omega \frac{d}{dc}[1 - F(c)] \cdot \frac{dc}{de} \,|_{\{(e,c)|e=e^*(c)\}} .$$

Therefore, $e^*(c)$ satisfies the differential equation $\frac{de^*}{dc} = -\left(\frac{\omega}{c}\right)\left(\frac{dF(c)}{dc}\right)$. Integrating this yields:

$$e^*(c) - e^*(c^{\max}) = \omega \int_{c}^{c^{\max}} \frac{dF(y)}{y}.$$

Finally, $e^*(c^{\max}) = 0$ since a worker with cost $c = c^{\max}$ loses with probability one. This establishes the result.

B. Introducing Categories

Suppose now that employers can divide the population of workers into two identifiable categories. Let $i \in \{1, 2\}$ index a worker's category, and let $\pi_i > 0$ denote the fraction of the worker population belonging to category i, where $\pi_1+\pi_2=1$. Hereafter, we use the subscript, i, to indicate categorical identity. Thus, $e_i(c)$ denotes the effort exerted by a member of category i with cost c, F_i denotes the cumulative distribution function of cost for category i workers, and G_i denotes the cumulative distribution of effort (in equilibrium) among category i workers. We will assume throughout that the cost distributions for the two categories have a common support.

Given our random matching assumptions, π_i is also the probability that any worker, regardless of his category, is paired to compete with a worker from category i. Since workers choose their effort prior to being paired against an opponent, and because they are assigned randomly to firms and then paired with each other at random for competition within firms, every worker faces the same distribution of opponents (i.e., a statistical replica of the overall worker population.) Thus, despite any asymmetry between categories that arises when $F_1(c) \neq F_2(c)$, the workers paired to compete with one another are playing a symmetric game. So, absent any category-based policy intervention, the equilibrium behavior of firms and of workers (whatever their category) will be as described in Proposition 1, with $F(c) \equiv \pi_1 F_1(c) + \pi_2 F_2(c)$.

Now, with two distinct categories, three types of matches are possible: a category 2 worker can be matched with a category 2 worker, which occurs with probability π_2^2; a category 1 worker can be matched with a category 1 worker, which occurs with probability π_1^2; and a mixed match (between category 1 and category 2 workers) occurs with probability $2\pi_1\pi_2$. We will assume that category 1 is "advantaged" relative to category 2, in that the category 1 cost distribution monotonically first-order stochastically dominates that of category 2.

Definition 1 *Category* 1 *is said to be advantaged relative to category* 2, *if*

$$\frac{f_1(c)}{f_2(c)} \text{ is a strictly decreasing function on } (c_{\min}, c^{\max}). \tag{4}$$

This definition implies $F_1(c) > F_2(c)$ for all $c \in (c_{\min}, c^{\max})$. Let γ_0 denote the probability that a category 1 worker wins when matched with a category 2 worker, in the laissez-faire equilibrium described in Proposition 1. We refer to γ_0 as the "natural win rate" of category 1 agents over category 2 agents. Given Definition 1, it is intuitively obvious and straightforward to show that

$\gamma_0 \equiv \int_{c_{\min}}^{c^{\max}} dF_1(c)\left[1 - F_2(c)\right] > \frac{1}{2}$.[2]

Because of the asymmetries in the cost distributions there will be a short-fall in the share of category 2 agents hired by the employer, relative to their share of the worker population. Specifically, the proportion of contests won by category 2 workers is

$$\pi_2^2 + 2\pi_2(1 - \pi_2)(1 - \gamma_0) < \pi_2,$$

in view of the fact that $\gamma_0 > \frac{1}{2}$. Thus, a demand for redistribution could naturally arise here. This is the subject of the next section.

3 Decentralized Categorical Redistribution

Decentralized redistribution involves equilibrium handicapping of workers in the disadvantaged category by employers who take the distribution of worker effort as exogenous when setting the handicap.[3] For instance, a set of universities designing their admissions policies to ensure sufficient categorical diversity, each of which thinks its policies are unlikely to affect the effort distribution in the pool of college bound seniors from which its applicants are drawn, is a case of decentralized categorical redistribution.

We will also employ the distinction between *blind* and *sighted* redistribution. Blind versus sighted redistribution refers to what markers employers are allowed to use in the pursuit of her redistribution policies. Category-sighted handicapping allows employers to use category identification directly in achieving their redistributive target. Category-blind handicapping forbids the use of categorical information in achieving categorical redistributive goals.

[2] To make this transparent, notice that:

$$\int_{c_{\min}}^{c^{\max}} dF_1(c)\left[1 - F_2(c)\right] = \int_{c_{\min}}^{c^{\max}} dF_1(c)\left[1 - F_1(c)\right] + \int_{c_{\min}}^{c^{\max}} dF_1(c)\left[F_1(c) - F_2(c)\right] \tag{5}$$

and

$$\int_{c_{\min}}^{c^{\max}} dF_1(c)\left[1 - F_1(c)\right] = -\frac{1}{2}\int_{c_{\min}}^{c^{\max}} \frac{d}{dc}\{[1 - F_1(c)]^2\} = \frac{1}{2}. \tag{6}$$

Further, under the assumption that category 2 is disadvantaged, $\int_{c_{\min}}^{c^{\max}} dF_1(c)\left[F_1(c) - F_2(c)\right]$ is necessarily a positive number. Thus, we have the desired inequality.

[3] For concreteness, one can think about redistribution in this context as coming about because all employers are required by some external authority to increase their hiring rate for workers in the disadvantaged category, though each employer acts independently of the others to achieve this goal.

A. Category Sighted Decentralized Redistribution

Suppose a regulator wants to decrease the win rate of the advantaged category, and let $\gamma \in [\frac{1}{2}, \gamma_0)$ denote the target level of categorical diversity.[4] Notice that, as $\gamma \to \gamma_0$, the laissez-faire equilibrium described in Proposition 1 obtains, while as $\gamma \to \frac{1}{2}$, the employer is forced to achieve categorical parity.

It is straightforward to show that if an employer desires to maximize the expected effort of the contestants, subject to the constraint that agents from the advantaged category only win at the rate $\gamma \in \left[\frac{1}{2}, \gamma_0\right)$ when matched with an agent from the disadvantaged category, then the best way to do so is to give category 2 workers a constant effort handicap, $\lambda^*(\gamma)$. In particular, the employers' optimization problem implies the maximization of a Lagrangian form as follows:

$$\max_A \left\{ \int_0^\infty \int_0^\infty \left[A(e_1, e_2) e_1 + (1 - A(e_1, e_2)) e_2 + \lambda (\gamma - A(e_1, e_2))\right] dG_1(e_1) dG_2(e_2) \right\}$$

Obviously, the solution must takes the form:

$$A(e_1, e_2) = \begin{cases} 1 \text{ if } e_1^*(c) > e_2^*(c) + \lambda^*(\gamma) \\ 0 \text{ if } e_1^*(c) < e_2^*(c) + \lambda^*(\gamma) \end{cases},$$

where $\lambda^*(\gamma)$ is the Lagrangian multiplier on the redistribution constraint. (In effect, $\lambda^*(\gamma)$ is the "shadow price of diversity" when the redistribution target is γ.) Thus, the equilibrium handicap is independent of the effort levels of the contestants, and varies positively with the aggressiveness of the redistribution goal.

It follows that, if $e_i^*(c)$ denotes the equilibrium effort supply of workers in category $i \in \{1, 2\}$, then when an agent from category 1 is matched with an agent from category 2, the agent in category 1 wins the contests if

$$e_1^*(c) > e_2^*(c) + \lambda^*(\gamma).$$

Notice that a category 1 worker who exerts low effort, $e_1 \in (0, \lambda^*)$, must lose if matched with any category 2 worker.

[4] That is, γ denotes the target win rate of Category 1 workers when matched against opponents from Category 2. Notice that, under sighted redistribution, these are the only matches that a regulator would seek to influence.

We will now derive the equilibrium in this model under category-sighted redistribution with target $\gamma \in [\frac{1}{2}, \gamma_0)$. Suppose there is an effort cost threshold, $c^*(\gamma)$, such that category 1 agents with cost $c_1 \geq c^*(\gamma)$ lose when matched with any category 2 agents, and category 1 agents with cost $c_1 < c^*(\gamma)$, win when matched with any agent with higher cost.[5] Then, by the redistribution constraint:

$$c^*(\gamma) \text{ solves } \int_{c_{\min}}^{c^*} dF_1(c)\left[1 - F_2(c)\right] = \gamma. \tag{7}$$

Notice, $c^*(\gamma)$ is increasing in γ, and $c^*(\gamma)$ tends toward $c^{\max}$, as γ tends toward γ_0; the natural win rate of category1 agents over category 2's.

To solve the model for any desired level of redistribution, $\gamma \in \left[\frac{1}{2}, \gamma_0\right)$, we must solve for the equilibrium de-centralized handicap, λ, and the associated equilibrium effort levels. This is the subject of our next result.

Proposition 2 *Given $\gamma \in \left[\frac{1}{2}, \gamma_0\right)$ and $c^*(\gamma)$ defined in equation (7), the equilibrium de-centralized category-sighted handicap is given by*

$$\lambda^*(\gamma) = \omega\left[\pi_2 \int_{c^*(\gamma)}^{c^{\max}} \left[\frac{1}{c^*(\gamma)} - \frac{1}{y}\right] dF_2(y) + \pi_1 \int_{c^*(\gamma)}^{c^{\max}} \frac{dF_1(y)}{y}\right], \tag{8}$$

the associated effort levels are given by

$$e_1^*(c) = e_2^*(c) + \lambda^*(\gamma) = \omega\left[\pi_1 \int_c^{c^{\max}} \frac{dF_1(y)}{y} + \pi_2\left[\int_c^{c^*(\gamma)} \frac{dF_2(y)}{y} + \int_{c^*(\gamma)}^{c^{\max}} \frac{dF_2(y)}{c^*(\gamma)}\right]\right],$$

for all $c \in [c_{\min}, c^*(\gamma))$, and

$$e_1^*(c) = \omega\pi_1 \int_c^{c^{\max}} \frac{dF_1(y)}{y}; \qquad e_2^*(c) = \omega\pi_2 \int_c^{c^{\max}} \frac{dF_2(y)}{y}$$

for all $c \in [c^*(\gamma), c^{\max}]$.

Proposition 2 provides a solution to the decentralized category-sighted handicapping problem. The result depends critically on three factors: (1) constant marginal cost of effort; (2) de-centralized setting of handicaps by many independent employers facing regulation; and (3) random matching.

[5] We shall show momentarily that in the presence of constant effort handicapping for category 2 agents the equilibrium effort supply functions imply this property.

The solution implies a partition of agents into four classes: $\{\text{category 1 or 2}\}\times\{\text{high cost } (c \geq c^*(\gamma))$ or low cost $(c < c^*(\gamma))\}$. For convenience of exposition, let H_i (resp. L_i) denote the set of high (resp. low) cost types of category $i \in \{1,2\}$. The equilibrium behavior of the agents under decentralized category-sighted handicapping can be summarized in the following concise manner. $H_i's$ only compete at the margin against other $H_i's$ in their same category, and lose to $L_i's$ in either category. Further, $L_i's$ compete at the margin against anyone with whom they are matched, prevailing if and only if they encounter a contestant with higher cost. However, $L_2's$ receive the effort subsidy $\lambda^*(\gamma)$, such that $e_1^*(c) = e_2^*(c) + \lambda^*(\gamma)$ for all $c \in [c_{\min}, c^*(\gamma))$. Figure 2 provides a graphical illustration of Proposition 2.

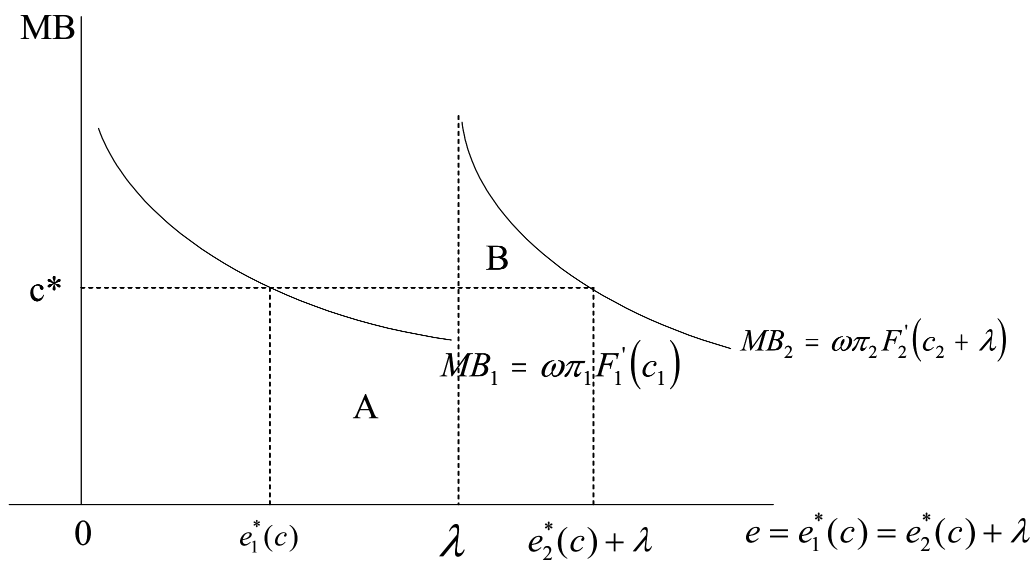

Figure 2: Decentralized Category-Sighted Equilibrium

To establish the proposition, let $e = e_1 = e_2 + \lambda$ denote effective effort supply of agents, when category 2 workers receive the effort handicap λ. Given that, in equilibrium, $e_i \geq 0$ and $e_i(c^{\max}) = 0, i = 1,2$, there exist a set of high costs agents in both categories (i.e., those with c "close to" $c^{\max}$) who supply relatively low effort. For category 1 agents this implies $e = e_1 \approx 0$, and for category 2 agents this implies that $e = e_2 + \lambda \approx \lambda > 0$. Thus, category 1 agents will face a non-convex decision problem, in equilibrium, since at very low effort levels $(e_1 < \lambda)$ they compete only against other high cost category 1 agents. At levels $e_1 > \lambda$, they discontinuously encounter additional benefits from marginal increases in effort, due to the presence of some high cost category 2 agents who are choosing $e_2 \approx 0$, so $e_2 + \lambda \approx \lambda$. Figure 2 captures this intuition. Because category 1's marginal benefit curve jumps upward at λ, it must be the case that the optimal effort function $e_1^*(c)$ is discontinuous. At cost $c = c^*$ the area (in Figure 2) $A \equiv [\lambda - e_1(c^*)] \cdot c^*$ is exactly equal (by construction) to the area $B \equiv \int_{\lambda}^{e_2(c^*)+\lambda} [MB_2(e) - c^*]\, de$. Therefore, $H_1's$ prefer $e_1^*(c) < \lambda$, along the MB_1 curve, whereas $L_1's$ prefer to "jump" past λ to some $e_1^*(c) > e_2(c^*) + \lambda$. In other words, all $H_1's$ choose to lose to all category 2 agents, should they end-up paired with one, and compete at the margin only against other $H_1's$; and all $L_i's$ lose only if they are matched with a worker who has lower cost. When λ has been set such that the constraint that workers in category 1 win against workers in category 2 with probability $\gamma \in \left[\frac{1}{2}, \gamma_0\right)$ holds, then it is straightforward to verify that we must have $c^* = c^*(\gamma)$ defined in equation (7).

Using equation (1) and noting that $H_i's$ only compete at the margin against other $H_i's$ in the same category, we deduce that the marginal benefit for category 1 and category 2 agents can be expressed as:

$$\begin{aligned} MB_1(e) &= c \text{ if and only if } e = \omega\pi_1 \int_c^{c^{\max}} \frac{dF_1(y)}{y} \equiv e_1^*(c) \\ MB_2(e) &= c \text{ if and only if } e = \omega\pi_2 \int_c^{c^{\max}} \frac{dF_2(y)}{y} \equiv e_2^*(c) \end{aligned}$$

Recall, however, $H_1's$ always have the option of boosting their effort to compete with category 2 workers. This is a break-even proposition when $c = c^*$ and strictly pays if $c < c^*$. Hence, for all $c < c^*$, the effort supply function looks much like that in Proposition 1. Both categories supply the same effective effort, given their cost, and $e^*(c)$ solves the differential equation:

$$\frac{de^*}{dc} = -\frac{\omega}{c}\frac{dF}{dc} = -\frac{\omega}{c}\left[\pi_1 F_1'(c) + \pi_2 F_2'(c)\right], \tag{9}$$

as in Proposition 1, but with the boundary condition: $e^*(c^*) = e_2^*(c^*) + \lambda$. Finally, λ must satisfy

$$c^*[\lambda - e_1^*(c^*)] = v_2(c^*)$$

where, using equation (2), $v_2(c) = \omega\pi_2 \int_c^{c^{\max}} \left[1 - \frac{c}{y}\right] dF_2(y)$. By integrating equation (9), using the relevant boundary conditions and definitions derived thus far, the conditions of the Proposition can be easily verified by algebraic manipulation.

B. Category-Blind Decentralized Redistribution

The employer's problem is more complicated when she is not allowed to use categorical information in the pursuit of her redistributive goals. Employers observe two effort levels in each match, e and $\widehat{e}$, but under the blindness assumption they do not know the categorical identity of the workers. Accordingly, to achieve their redistributive objectives, employers need to estimate the likelihood that each effort level was emitted from a worker in the disadvantaged category. Thus, we have a signalling model. This is a point worth further emphasis. Under category-sighted redistribution an employer is allowed to narrowly tailor her policies for agents in disadvantaged categories in order to achieve her categorical diversity goal – focusing exclusively on handicapping in mixed contests. When constrained to be category-blind, however, she has to implement her policy across all contests.

Let $\xi(e)$ denote an employer's belief about the probability that a worker with observed effort level e is from category 2. Then the employer's problem can be written as:

$$\max_A \left\{ \int_0^\infty \int_0^\infty \{A(e,\widehat{e})\, e + (1 - A(e,\widehat{e}))\, \widehat{e} + \lambda [\gamma - A(e,\widehat{e})\, \xi(e) - (1 - A(e,\widehat{e}))\xi(\widehat{e})]\} dG(e) dG(\widehat{e}) \right\}$$

The solution takes on the form:

$$A(e,\widehat{e}) = \begin{cases} 1 \text{ if } e + \lambda\xi(e) > \widehat{e} + \lambda\xi(\widehat{e}) \\ 0 \text{ if } e + \lambda\xi(e) < \widehat{e} + \lambda\xi(\widehat{e}) \end{cases},$$

where λ is a Lagrangian multiplier associated with the redistribution constraint – the shadow price on category 2 membership. Let $V(e) = e + \lambda\xi(e)$ denote the value to an employer of a worker with observed effort e. In any pair-wise contest, the employer hires the worker with the higher value. If two matched workers have equal value, the employer chooses either with probability 1/2. This value is comprised of two parts: the direct benefit to the employer of effort, and the expected benefit of diversity, which equals the product of the likelihood of an individual with effort e belonging to the disadvantaged category, times the shadow price of diversity.

To characterize equilibrium with category-blind redistribution, it is helpful to think about the qualitative properties of $V(e)$. As in the model under Laissez-faire with no categorical redistribution goal, a revealed preference argument establishes that the equilibrium effort supply of workers, $e^*(c)$, must be non-increasing. If this function is strictly decreasing on $[c_{\min}, c^{\max}]$ then we have a separating equilibrium: employers, observing a worker's effort, can invert the equilibrium effort supply function to learn the worker's cost, and to infer the likelihood that the worker belongs to category 2. If $e^*(c)$ is constant on some range of costs, then we have pooling in equilibrium. Now, let E denote the set of efforts reached by some worker in equilibrium: $E = \{e = e^*(c)$, for some $c \in [c_{\min}, c^{\max}]\}$. It is obvious that if $e^*(c)$ is an equilibrium effort schedule for workers, then $V(e)$ must be strictly increasing on E. For if there were two levels of effort, $e, \tilde{e} \in E$ with $e < \tilde{e}$ and $V(e) \geq V(\tilde{e})$, then any worker choosing $\tilde{e}$ could gain by reducing his level of effort to e, which lowers his cost incurred without lowering his chances of winning the contest. But, from this it follows that a separating equilibrium cannot obtain here. For, if $e^*(c)$ were strictly monotonic on $(c_{\min}, c^{\max})$, and if $V(e)$ were strictly increasing on E, then a worker would not be hired whenever matched against another worker with lower cost[6], in which case the redistribution constraint could not be satisfied. We conclude that there must be some pooling in equilibrium. That is, the equilibrium effort supply function, $e^*(c)$, must be constant over some non-empty interval(s) of costs. Moreover, a worker in the pool would be hired (not hired) with probability one when paired with a worker whose effort level is lower (higher) than that of the pool, and would be hired with probability 1/2 when paired with another worker in the pool. The size of the pool will increase with the aggressiveness of the redistribution goal.

This situation is captured in Figures 3, which show a pooling equilibrium where all worker types c in the closed interval $[\widehat{c}, \widehat{\widehat{c}}]$ select the common effort level, e_{pool}. Figure 3A depicts a worker's value to the employer as a function of his effort. Figure 3B shows the worker's best response to the employer strategy "hire that worker with the greater value," as a function of the worker's cost.

[6] Ties will occur with probability zero when $e^*(c)$ is strictly monotonic, since we have assumed the cost distribution to be non-atomic.

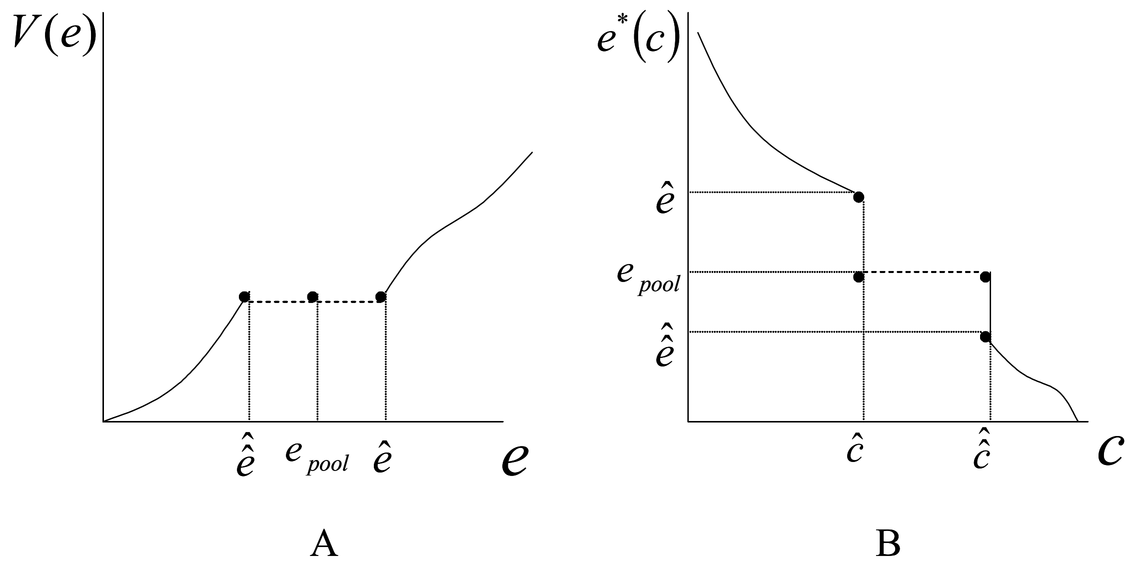

Figure 3: Decentralized Category-Blind Equilibrium

In any equilibrium under category-blind redistribution several conditions need to be satisfied, all of which can be illustrated in Figures 3.[7] First, the worker with marginal cost $\widehat{c}$ must be indifferent between leaving the pool by putting in effort $\widehat{e}$, which would imply that he wins for sure against all workers in the pool, and exerting the effort e_{pool}, which has him winning with probability $\frac{1}{2}$ when matched with anyone in the pool. Similarly, the worker with cost $\widehat{\widehat{c}}$ must be indifferent between staying with the pool, and reducing his effort to $\widehat{\widehat{e}}$. Also, firms must be indifferent between hiring from the pool, and hiring a worker with effort level $\widehat{e}$ (resp., $\widehat{\widehat{e}}$) when such a worker is known to have a cost [and associated probability of belonging to category 2] of $\widehat{c}$ (resp., $\widehat{\widehat{c}}$).[8] Finally, we must specify an employer's beliefs in the event that she were to observe an effort $e \in [\widehat{\widehat{e}}, e_{pool}) \cup (e_{pool}, \widehat{e}]$, which is off the equilibrium path. To support the candidate pooling equilibrium, employers' beliefs must be such that they would strictly prefer a worker in the pool when matched against a hypothetical worker with effort in (the interior of) this region. Here we will appeal to a large literature on

[7] These figures depict a single pooled effort level, whereas in principle there could be many pools in equilibrium. However, we will soon introduce a natural restriction on employers' out-of-equilibrium beliefs that implies the existence of a unique equilibrium in this model, with a single pool consisting of the highest cost worker types exerting the minimal effort level. Accordingly, the exposition proceeds from this point onward under the supposition that there is but one pool in equilibrium.

[8] This indifference condition for firms is required because, were it to fail, then for any plausible off-equilibrium-path beliefs that firms might hold, they would want to respond to some deviation from the pool in such a way as to make that deviation pay for some workers in the pool.

equilibrium refinements and select a natural one, the D1 refinement (Cho and Kreps, 1987).

Loosely speaking, the D1 refinement requires out-of-equilibrium actions by informed agents (workers) to be interpreted by uninformed agents (firms) as having been taken by the worker type who would gain most [lose least] from the deviation, relative to his payoff in the candidate equilibrium, when this gain [loss] is calculated under the supposition that firms, while adopting this interpretation, will respond to the deviant action in a manner that is best for themselves. (Even more loosely speaking, the D1 refinement requires firms to believe that the deviator is the type who gains most from taking the deviant action when he knows that firms will discover his type when he deviates, and then, based on knowing his type, will respond optimally to his deviant action.) Equilibria supported by such out-of-equilibrium beliefs are called *D1 equilibria.*[9] It is a straight-forward exercise to compute the (unique) D1 equilibrium in our model.

Under D1, if an employer observes an effort $e \in (e_{pool}, \widehat{e}]$, she believes that the deviator is the lowest cost type in the pool, $\widehat{c}$. (This type, compared to others either inside or outside of the pool, gains most [loses least] from such a deviation.) Hence (using Bayes's Rule), employer beliefs must satisfy:

$$\xi\left(e\right) = \frac{\pi_2 f_2(\widehat{c})}{f(\widehat{c})} \equiv \widehat{\xi}, \quad \text{for all } e \in (e_{pool}, \widehat{e}].$$

In light of the indifference conditions mentioned above, no workers inside or outside of the pool have an incentive to deviate by choosing e in this interval. On the other hand, if $e_{pool} > 0$, and if an employer observes an effort $e \in [\widehat{\widehat{e}}, e_{pool})$, then under D1 she must believe that the deviator is the highest cost type in the pool, $\widehat{\widehat{c}}$. (This type, compared to all of the others, gains most [loses least] from such a deviation.) Accordingly, under D1 employer beliefs must satisfy:

$$\xi\left(e\right) = \frac{\pi_2 f_2(\widehat{\widehat{c}})}{f(\widehat{\widehat{c}})} \equiv \widehat{\widehat{\xi}}, \quad \text{for all } e \in [\widehat{\widehat{e}}, e_{pool}).$$

But now, some workers will have an incentive to deviate. To see this, let ξ_{pool} be the probability that a worker belongs to category 2, conditional on the worker being in the pool. Then, in light

[9]It has been shown that the only D1 equilibrium to the canonical Spence job signaling model is the Riley (separating) equilibrium (Riley, 1979). This is the unique efficient, separating equilibrium defined by the initial condition wherein the infimal separating ability type adopts its complete information best educational investment level, while all higher types choose the lowest educational levels consistent with separation (which strictly exceed their respective complete information decisions.)

of the assumption that category 1 is advantaged, the monotone first-order stochastic dominance property implies:

$$\widehat{\xi} < \xi_{pool} \equiv \frac{\pi_2[F_2(\widehat{c}) - F_2(\widehat{\widehat{c}})]}{[F(\widehat{c}) - F(\widehat{\widehat{c}})]} < \widehat{\widehat{\xi}}.$$

Now, it was an implication of the firm's constrained maximization problem that the value of a worker to the firm is $V(e) = e + \lambda\xi(e)$, with the Lagrangian multiplier (the shadow price of diversity) $\lambda > 0$. Hence, a worker in the pool can anticipate that his value will increase if he deviates by lowering his effort slightly below e_{pool} (since this marginal reduction induces firms to believe he is strictly more likely than someone drawn from the pool to belong to category 2.) It follows that in all D1 pooling equilibria, e_{pool} must equal zero. But then, since $e^*(c)$ is non-increasing in any equilibrium, we must have that $\widehat{\widehat{c}} = c^{\max}$. Figures 4 replicate Figures 3 after the application of the D1 refinement, showing the unique D1 equilibrium in this model.

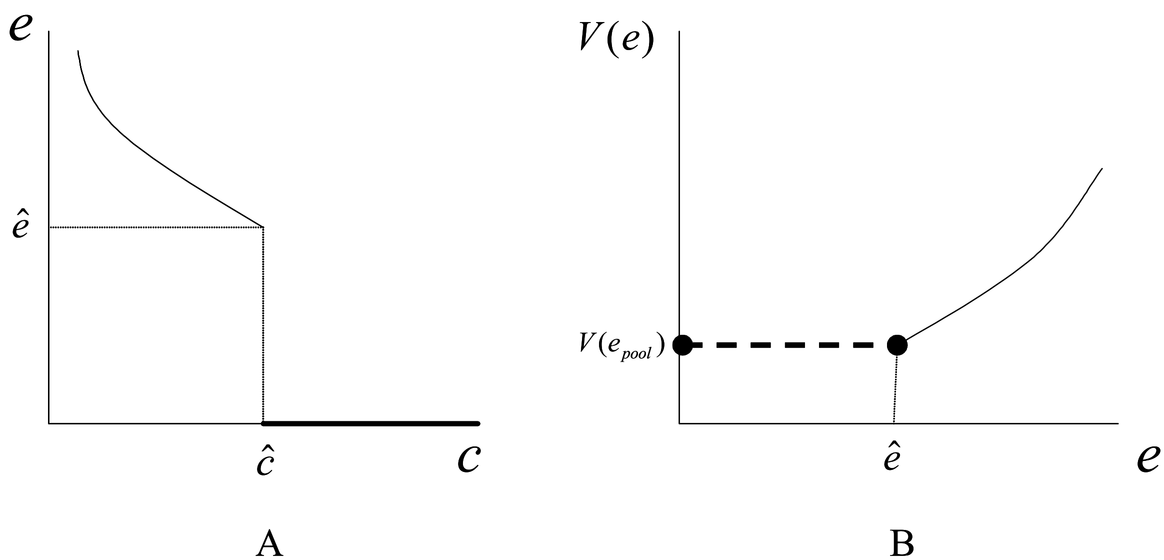

Figure 4: Decentralized Category-Blind Equilibrium After Application of D1

We can summarize the discussion to this point as follows:

Proposition 3 *Given* $\gamma \in \left[\frac{1}{2}, \gamma_0\right)$, *the equilibrium de-centralized category-blind handicap, after application of the D1 refinement, is given by*

$$\lambda = \frac{\widehat{e}}{\xi_{pool} - \widehat{\xi}}, \quad \textit{where } \widehat{e} = \frac{\frac{\omega}{2}\left[1 - F\left(\widehat{c}\right)\right]}{\widehat{c}}; \xi_{pool} = \frac{\pi_2[1 - F_2(\widehat{c})]}{[1 - F(\widehat{c})]}; \ \widehat{\xi} = \frac{\pi_2 f_2(\widehat{c})}{f(\widehat{c})};$$

and where $\widehat{c}$ *is the unique solution for*

$$\int_{c_{\min}}^{\widehat{c}} dF_1\left(c\right)\left[1 - F_2\left(c\right)\right] + \frac{1}{2}\left[1 - F_1\left(\widehat{c}\right)\right] \cdot \left[1 - F_2\left(\widehat{c}\right)\right] = \gamma.$$

The equilibrium effort supply function for both categories is given by

$$e^*\left(c\right) = \widehat{e} + \omega \int_{c}^{\widehat{c}} \frac{dF\left(y\right)}{y}, \ c \in \left[c_{\min}, \widehat{c}\right),$$

and

$$e^*\left(c\right) = 0, \ \textit{for all } c \in \left[\widehat{c}, c^{\max}\right].$$

Proposition 3 highlights an important feature of category-blind redistribution; there is a non-trivial measure of workers that supply zero effort. And, this pool increases with the aggressiveness of the redistributive goal. In the extreme case where the employer strives for categorical parity, the only *D1-equilibrium* involves all workers supplying zero effort and the employer picking at random between them![10]

To establish the proposition, consider figures 4. In order to derive the equilibrium, we must pin

[10] This result has a curious implication that warrants mention. If the cost distributions are identical for the two categories, then each category prevails in half the mixed contests, without any constraint on firm actions. So, the equilibrium effort schedule given in Proposition 1 (which is a positive, strictly decreasing function of effort cost) obtains in this case, automatically generating $\gamma = \frac{1}{2}$. Yet, with only the slightest (strict monotone likelihood ratio) difference in cost distributions favoring category 1, our characterization given above of the unique D1 equilibrium under category-blind redistribution with representation target $\gamma = \frac{1}{2}$ implies zero effort for all agents. This discontinuity of the equilibrium effort supply schedule, as a function of the cost distributions, when population proportionality is the affirmative action target is curious, and it is not an artifact of our having imposed the D1 refinement. For (per the argument just given) in any equilibrium, when firms see an effort level e they (in effect) place some value $V(e)$ on a worker with that effort level, hiring from any pair of workers the one whose value is greater. Moreover, $V(e)$ must be strictly increasing on the set of efforts observed by firms in equilibrium, and $e^*(c)$ must be non-decreasing, otherwise workers could not be best-responding. All of which implies $V(e^*(c))$ must be non-increasing in any equilibrium, which, in light of Definition 1 implies that the target $\gamma = \frac{1}{2}$ can only be met in equilibrium if the set E is a singleton. So, population proportionality as a target together with strict cost distribution differences between the groups, however small, requires a pooling equilibrium with all workers taking the same effort level. Imposing D1 merely forces that pooled effort level to be zero. Hence, the aforementioned discontinuity does not depend on imposing D1.

down four parameters: (i) $\widehat{c}$; (ii) $\widehat{e}$; (iii) λ; and (iv) $e^*(c)$ for all $c \in [c_{\min}, \widehat{c}]$. Under category-blind redistribution, the effort incentives at the margin (for both categories) of those not in the pool are identical to the marginal incentives facing workers in the laissez-faire equilibrium. Using Proposition 1, it follows directly that $e^*(c) - \widehat{e} = \omega \int_c^{\widehat{c}} \frac{dF(y)}{y}$ for all $c \in [c_{\min}, \widehat{c})$. Thus, we are left with three equations (the workers' and firms' indifference conditions and the representation constraint,) and three unknowns $(\widehat{c}, \widehat{e}, \lambda)$. Consider, first, the representation constraint (which requires that the probability a category 1 worker wins when matched against a category 2 worker just equals γ.) In the *D1-equilibrium* being asserted here, this amounts to:

$$\int_{c_{\min}}^{\widehat{c}} dF_1(c)\left[1 - F_2(c)\right] + \frac{1}{2}\left[1 - F_1(\widehat{c})\right] \cdot \left[1 - F_2(\widehat{c})\right] = \gamma \tag{10}$$

Hence, the cost cut-off $\widehat{c}$ solves equation 10, which pins down (i).

Now, consider the workers' indifference condition. The worker with cost $\widehat{c}$ must be indifferent between exerting effort $\widehat{e}$ and effort 0. If he exerts $\widehat{e}$ he beats all workers in the pool and incurs the cost $\widehat{c}\widehat{e}$; if he invests 0, he ties all workers in the pool – winning with a probability of $\frac{1}{2}$ when matched against any one of them, but paying zero effort costs. The indifference condition implies that

$$\widehat{e} = \frac{\frac{\omega}{2}\left[1 - F(\widehat{c})\right]}{\widehat{c}},$$

which pins down (ii).

Finally, to establish the equilibrium shadow price of diversity, λ, consider the firm's indifference condition:

$$\widehat{e} + \lambda\widehat{\xi} = \lambda\xi_{pool},$$

where ξ_{pool} is the probability of a randomly drawn worker in the pool being disadvantaged. Obviously, this implies,

$$\lambda = \frac{\widehat{e}}{\xi_{pool} - \widehat{\xi}},$$

which establishes the desired result.

4 Centralized Categorical Redistribution: A Discussion

There is an impressive literature in economics involving optimal auction design. Insights from this literature can be directly applied to the analysis of tournaments in labor economics. In this section, we provide a brief discussion of the setup and solution of the optimal auction design problem and how it can be used to inform centralized categorical redistribution.[11]

Centralized redistribution is a mechanism design problem in which an employer selects the optimal (incentive compatible and individually rational) handicapping scheme – realizing that his employment policies will affect the workers' investment behavior. For instance, if a substantial fraction of college bound high school seniors attend in-state universities, then affirmative action programs designed by the state university system's Board of Regents could provide an example of centralized categorical redistribution. In our model, with the assumptions of constant marginal cost of effort and random pair-wise matching of competitors, the solution to the optimal tournament design problem is isomorphic to the optimal auction design problem with two (asymmetric) bidders. Indeed, optimal handicapping in tournaments can be interpreted as the optimal auction design problem with the additional constraint that a certain category of bidders win the auction at no less than some minimal rate.

A. Setting Up the Optimal Tournament Design Problem

Let there be one employer and two workers competing for employment. Assume that this competition will be replicated many times over, as pairs of competing workers are drawn from a large population. The employer is interested in designing a tournament that maximizes the expected effort of the winner, perhaps subject to some categorical representation constraint. The employer's problem derives from the fact that he does not know the effort costs of the workers.

In a direct revelation mechanism, the contestants simultaneously and confidentially announce their cost of effort. The employer then determines who wins the employment tournament and how much effort each contestant must exert, as some functions of the reported costs $c = (c_1, c_2)$.[12] A direct revelation mechanism is described by a pair of functions (A, e) (of the form $A : [c_{\min}, c^{\max}]^2 \rightarrow$

[11] For the details of the optimal auction design problem, we refer the reader to Myerson (1981).

[12] Here, by taking c to be the cost vector for a pair of contestants rather than the cost type of an arbitrary contestant, we slightly shift our notational convention from that of the previous section.

$[0,1]$ and $e : [c_{\min}, c^{\max}]^2 \to \Re^2_+)$ such that if c is the vector of announced costs, $A_i(c)$ is the probability that i wins the tournament and $e_i(c)$ denotes the effort i must exert.[13]

The employer and the contestants are assumed to be risk neutral, and contestants have utility functions that are additively separable in wages and effort cost. If contestant i knows that his effort cost is c_i, then his expected utility from the tournament mechanism described by (A, e) is

$$U_i(A, e, c_i) = \int_{c_{\min}}^{c^{\max}} [A_i(c)\,\omega - e_i(c)\,c_i]\, f_{-i}(c_{-i})\, dc_{-i} \tag{11}$$

Similarly, the expected utility for employers from this tournament mechanism is

$$U_E(A, e) = \int_{c_{\min}}^{c^{\max}} \int_{c_{\min}}^{c^{\max}} \left(\sum_{i=1}^{2} A_i(c)\, e_i(c) \right) dF_1(c_1)\, dF_2(c_2) \tag{12}$$

There are three types of constraints that must be imposed on the pair of functions (A, e) to ensure their feasibility. First, since each employer only has one position to fill, the function A must satisfy the following probability conditions:

$$\sum_{i=1}^{2} A_i(c) \leq 1 \text{ and } A_i(c) \geq 0, \text{ for all } i \in \{1, 2\} \text{ and } c \in [c_{\min}, c^{\max}]^2. \tag{13}$$

Second, we assume that the employer cannot force a worker to participate in a tournament. Thus, the following individual rationality conditions must be satisfied:

$$U_i(A, e, c_i) \geq 0, \text{ for all } c_i,\ i \in \{1, 2\}\ . \tag{14}$$

Third, we assume that the employer can not prevent any worker from lying about his effort costs if he can gain from lying. If worker i were to claim that $\widehat{c}_i$ were his effort cost when it really was c_i, then his expected utility would be

$$\int_{c_{\min}}^{c^{\max}} [A_i(\widehat{c}_i, c_{-i})\,\omega - e_i(\widehat{c}_i, c_{-i})\,c_i]\, f_{-i}(c_{-i})\, dc_{-i}$$

Therefore, to guarantee that no worker has an incentive to lie about his effort cost, the following

[13] Here the subscript i denotes contestants, not categories. Later, when we restrict attention to categorically heterogeneous matches, this distinction will be irrelevant.

incentive compatibility conditions must be satisfied:

$$U_i(A,e,c_i) \geq \int_{c_{\min}}^{c^{\max}} \left[A_i(\widehat{c}_i, c_{-i})\,\omega - e_i(\widehat{c}_i, c_{-i})\,c_i\right] f_{-i}(c_{-i})\,dc_{-i} \text{ for all } i \in \{1,2\} \text{ and } \widehat{c}_i \in [c_{\min}, c^{\max}] \tag{15}$$

We say that a tournament mechanism (A,e) is feasible if and only if (13), (14), and (15) are all satisfied. So, (A,e) represents an optimal tournament if and only if it maximizes $U_E(A,e)$ subject to $(13)-(15)$.[14]

B. A Solution to the Optimal Tournament Design Problem

There is little difference, analytically, between our "tournaments" and the "auctions" in Myerson (1981). To move from our current framework to that analyzed in Myerson (1981), one needs to multiply each agent's utility by a scaling factor that depends only on his own type. This amounts to a simple change of variables. It is straightforward to show that by defining bidder i's "value" in the optimal auction, t_i, design problem such that $t_i = \frac{\omega}{c_i}$, one arrives at a formulation that is isomorphic to the tournament problem when the wage is ω and the i^{th} contestant's cost is c_i. So, to study the optimal centralized tournament, we can simply carry over Myerson's results, translated into our notation, where his bidder valuation t_i is related to our worker effort cost, c_i, via that equation: $t_i = \frac{\omega}{c_i}$.

Thus, let $H_i(t_i)$ be the CDF of bidder $i's$ valuation in Myerson's context, and let $h_i(t_i)$ be the associated density function. Then, implementing the change of variables, we have:

$$H_i(t_i) = 1 - F_i(\frac{\omega}{t_i}) \text{ and } h_i(t_i) = (\frac{\omega}{t_i^2}) f_i(\frac{\omega}{t_i}) \tag{16}$$

Now, let $\psi(t_i) = t_i - \frac{1-H_i(t_i)}{h_i(t_i)}$ denote what Myerson (1981) called the *virtual valuation* of bidder i. Myerson proved that the solution to the optimal auction design problem assigns the object to the bidder with the highest virtual valuation, and requires that bidder to pay an amount equal to the minimum actual valuation he could have had that would still have allowed him to prevail in the auction. To translate this idea into our tournament model, first note that our analogue to

[14]It may appear, superficially, that we have changed the structure of the game by assuming that workers exert effort *before* being matched with employers in our analysis of decentralized redistribution and *after* being matched with employers in the centralized case. But, these two formulatoins are equivalent so long as workers exert effort before they know the effort of the worker with whom they are matched.

Myerson's virtual valuation (expressed in terms of our worker's effort cost) is given by:

$$\widehat{\psi}(c_i) = \omega[\frac{1}{c_i} - \frac{F_i(c_i)}{c_i^2 f_i(c_i)}]. \tag{17}$$

For any costs c_j, $j \neq i$, let

$$\mu_i(c_j) = \sup\left\{c_i \mid \widehat{\psi}(c_i) \geq \widehat{\psi}(c_j)\right\}.$$

So, $\mu_i(c_j)$ is the highest cost report for worker i such that his virtual valuation is greater than that of worker j. Employing a similar derivation to that found in Myerson (1981, Lemmas 2 and 3), one can verify that the solution to the optimal tournament consists of functions:

$$\begin{aligned} A_i(c) &= \begin{cases} 1 \text{ if } \widehat{\psi}(c_i) > \mu_i(c_j) \\ \quad 0 \text{ otherwise} \end{cases} \\ e_i(c) &= \begin{cases} \omega \int\limits_{\mu_i(c_j)}^{c^{\max}} \frac{dF_i(y)}{y}, \quad \text{if } A_i(c) = 1 \\ \qquad 0 \text{ otherwise} \end{cases} . \end{aligned}$$

We conclude that the optimal tournament gives the position to the contestant with the higher virtual valuation, as expressed in equation 17 above, and only requires effort from those that are hired. The amount of effort that they exert is the minimum consistent with incentive compatibility that is needed to distinguish themselves from a marginally higher effort cost type.

Let us now adopt the interpretation that the two contestants in the tournament are members of two distinct categories of the worker populations. (In effect, we restrict attention to the design problem for categorically heterogeneous contests.) The interesting aspect of the optimal tournament from our perspective is that, even in the absence of any categorical representation constraint, the expected effort-maximizing tournament will generally involve some (possibly mild) degree of handicapping. To see this, notice that the virtual valuation is simply a scalar times the reciprocal of each contestant's effort cost, minus what is, in effect, an inverse hazard ratio. Thus, suppose the reciprocal of a worker's effort costs were distributed exponentially in each category. In this special case, the virtual valuation in equation 17 is simply the reciprocal of a contestant's effort cost, minus the mean of the reciprocal of the effort cost for their respective category. So, i beats j if the difference in their efforts is larger than the difference in the means of the reciprocal of effort cost for their respective categories! It follows that a desire for small amounts of categorical redistribution in a centralized planning framework need not conflict with efficiency.[15]

If we want more aggressive redistribution, however, this constraint may indeed be binding. Suppose a tournament designer wants to have a level of categorical diversity beyond that which naturally occurs in the tournament design problem. Setting up the employers maximization problem is virtually identical to our analysis of the first best mechanism, with an added representation constraint. That is, if a function $A : [c_{\min}, c^{\max}] \to \Re^2$ maximizes

$$\int_{c_{\min}}^{c^{\max}} \int_{c_{\min}}^{c^{\max}} \left(\sum_{i=1}^{2} A_i(c) \widehat{\psi}(c_i) + \lambda [\gamma - A_2(c_1, c_2)] \right) dF_1(c_1) \, dF_2(c_2)$$

subject to (13), (14), and (15). Then, (A, e) represents an optimal tournament, conditional on achieving representation, γ.

The solution to this constrained optimal tournament is straightforward. If contestant i belongs to the advantaged category and contestant j belongs to the disadvantaged category; i beats j, if and only if

$$\widehat{\psi}(c_i) > \widehat{\psi}(c_j) + \lambda.$$

The category-blind centralized problem is solved in a similar fashion. The optimal tournament design problem for an employer who desires redistribution and is forced to be category-blind shows that in all close matches, the contestant with the lower virtual valuation wins. The intuition is that the employer realizes that the disadvantaged category will, on average, finish second when matched with the advantage category. Therefore, in close matches, she allows the lower valuation contestant to win–assuming that enough of the winners will be disadvantaged to achieve the desired representation goal. She does this in all matches, because she is constrained to be category blind.

[15] Our point here is similar to the classical observation that a price-setting monopsonist buying from sellers drawn from two identifiably distinct population groups with different cost distributions will genereally price-discriminate in such a way as to favor the group with the more price-elastic supply curve. Here the firm is "buying" effort from workers in the two categories, while paying them with a probability of winning the contest. Moreover, the "elasticity of supply" of workers able to profitably exert a given level of effort is greater for the disadvantaged category. This interpretation centralized handicapping in the optimal tournament is an exact analogue of that offered in Bulow and Roberts (1989) for the optimal auction.

5 Conclusion

Tournament competition is an economic phenomenon that arises in many venues. The existing literature has been centered around the economic efficiency of tournament play and analyzing tournaments as optimum labor contracts. This paper opens new directions in the study of tournament theory by deriving the equilibrium handicapping strategy of employers in a decentralized environment. This, together with the contributions in the optimal auction design literature, take us a considerable way in understanding categorical redistribution in winner-take-all environments.

References

[1] Akerlof, George. 1978. "The Economics of Tagging as Applied to Optimal Income Tax, Welfare Programs, and Manpower Planning." *American Economic Review*, 68: 1, 8-19.

[2] Bulow, J. and J. Roberts. 1989. "The Simple Economics of Optimal Auctions," *Journal of Political Economy,* 97(5):1060-90.

[3] Cho, I-K. and D. Kreps. 1987. "Signaling Games and Stable Equilibria." *Quarterly Journal Economics,* May: 179-221.

[4] Fryer, Roland., and Loury, Glenn. 2003. "The Economics of Categorical Redistribution." Working paper. Harvard University.

[5] Green, Jerry R. and Stokey, Nancy L., "A Comparison of Tournaments and Contracts," *Journal of Political Economy*, June 1983, 91, 349-65.

[6] Lazear, Edward P. and Rosen, Sherwin, "Rank Order Tournaments as Optimum Labor Contracts," *Journal of Political Economy*, October 1981, 89, 861-64.

[7] Mirrlees, James. A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies*, vol. 38 (2), 175-208.

[8] Myerson, Roger. "Optimal Auction Design." *Mathematics of Operations Research.* February 1981, 6 (1), pp. 58-73.

[9] Riley, J G. 1979. "Informational Equilibrium," *Econometrica*, 47, 331-360.

[10] Rosen, Sherwin, "Prizes and Incentives in Elimination Tournaments" *American Economic Review*, September 1986, 76, 701-15.