Repeated Tournaments*

Hugo Hopenhayn and Matthew Mitchell

November 13, 2009

Abstract

We study an environment where a principal and two agents enter into a long term relationship. In each of many period the agents provide costly inputs, valued by the planner, which influence which agent is chosen (by nature) as the winner. We study how the optimal contract leads to asymmetry and state dependence. We also study long run features of the optimal contract. We show how the model can be used in applications fo firm-employee dynamics, as well as to optimal reward structures in repeated racing environments.

*UCLA and University of Toronto. Very preliminary!

1 Introduction

In this paper we study a model of a repeated tournament. By a tournament we mean a situation with a single principal and multiple agents where incentives are obtained through a coarse signal of relative inputs of the agents, which can be thought of as a stochastic measure of relative performance. The use of such systems has been well studied, and applied to a variety of applications in labor economics and industrial organization. Here we extend the tournament structure to a dynamic setting, where agents repeatedly give effort and the principal receives a coarse signal of relative inputs in a given period. We study the optimal dynamic contract in this setting and show how characteristics of the optimal contract can help us understand issues in terms of employers and the rewards they pay employees, as well as the way in which races are optimally arranged.

Tournament structures, where agents are paid on the basis of relative success, has been justified in a variety of ways. When agents face a common, unobserved cost shock, it is useful to condition rewards not only on the absolute level of an agents output but also on the relative level of success, since relative levels provide insurance implicitly condition on the common shock (Lazear and Rozen (1981), Green and Stokey (1983)). Moreover, absolute levels may be difficult to contract upon compared to relative levels. If a firm employs workers to solve problems, for instance, the firm may be able to pay based on whether or not the solution of the agent is the one chosen to be implemented; however, it is more difficult for the principal to accurately report exactly how good a solution has been provided on an absolute scale.

We derive three sets of results. First, for the most general structure, we show the sense in which only a limited set of situations can lead to treating the agents equally (in the sense of promised payouts) in the long run. We describe the ways in which symmetric states are exited.

We then consider two applications. First we interpret the model as a repeated patent race, and the planner as a patent authority. Here the model builds on the classic literature on patent races (Lowry (1979), Reinganum (1982)) to allow the competitors in the patent race to race again, for a new prize, after each race ends. The planner has a costly reward that can be smoothed by using a dynamic contract with the firms.

For the patent example, we show a class of cost functions (with a static budget constraint) that lead to equal effort from the two agents in the long run. In a sense, heterogeneity, which the general results show is a feature of the repeated tournament, does not translate into long run differences in efforts; instead, it is simply that one agent is treated permanently better in all states in the long run.

Thinking about recurrent innovators has largely been ignored in the recent literature studying patent policy as an optimal mechanism. Among dynamic papers in that literature, Hopenhayn, et al (2006) consider the case where innovators innovate separately, and only once, never recurring. We add both racing for a given innovation and recurrent innovators to that structure. On the other hand, whereas they stress the constraint that arises when innovations are cumulative, from the innovations sharing the same market from which to draw rewards, we focus on a case where there is no dynamic budget constraint for the planner, as is assumed in the classic repeated moral hazard literature (for instance, Rogerson (1985)). In the patent context, this can be interpreted either as the rewards for successive innovations being patents in different markets, or that the rewards are in the form of prizes, which are raised through costly taxation.

Two papers that do consider recurrent innovators are Riis and Shi (2009) and Acemoglu and Akcigit (2009). The former takes the structure of Hopenhayn, et al (2006), but with only two firms which recur. Firms still innovate one at a time, in the spirit of "private ideas" models like Arrow (1962) and Nordhaus (1967). Acemoglu and Ackigit (2009) study the optimal stationary patent policy, from a particular set of policy tools, in a step-by-step model of innovation. Their computations suggest the value of making policy state dependent, which we echo in the optimal policy in our general structure. Our long run results for the patent example, however, show that even a world with strong state dependency motives may not settle to a place where efforts depend on history.

Our second example is for a specific structure that fits the notion of tournaments within a firm. We show that wage differentials are increasing as the firm ages. Since age and size are correlated, this provides a rationale for the increasing wage differentials within firms as they grow. We show that this occurs despite the fact that across firm wage differentials need not rise, consistent with evidence on differentials across firms.

We show that the model can also shed light on the supposed contradiction between tournament theory, where wage differentials are an incentive device which enhance firm value, and fair wage theory (Akerlof (1982), Levine (1990)), where unequal wages lead agents to take actions detrimental to the firm. In the optimal dynamic contract, it is perfectly consistent for a firm's value to be maximized when agents are treated equally, but still have it be the case that the optimal contract leads with certainty to a divergence from equal treatment.

2 Model

A principal (sometimes referred to as the planner) and two agents enter into a long term relationship. In each of a countable number of periods the agents provide costly inputs, valued by the planner, which also influence which agent is chosen (by nature) as the winner. The planner can commit to a contract that specifies payments to both players at every date as a function of the history of winners.

2.1 Period Payoffs

Each agent $i \in \{0, 1\}$ chooses an amount $x_i \geq 0$ to invest. The investment is not directly observed by the planner. Let $x = x_0 + x_1$. The planner values these investments according to the bounded, continuous and differentiable function $u(x_0, x_1)$,where $u(0, 0) = 0$.

Exactly one agent is declared the winner. Agent i is the "winner" with probability $p(x_i, x_j) \in [0, 1]$. Winning is observable and verifiable. We assume that p is increasing in the first argument and decreasing in the second. It is concave (in just the first argument may be sufficient), continuous and twice differentiable.

The principal pays a prize W_i to the winner and L_i to the loser if his identity is i. We assume that the agents value the prize linearly, with one interpretation that the prizes are denominated in utils. We assume that there is a lower bound (possibly minus infinity) of $\underline{P}$.. Interpretations of the bound include limited liability and outside options for the agents or utility that is bounded below. The principal would like to condition payments on his realized payoff $u(x_0, x_1)$, since it contains information about effort, but we assume that such information is non-verifiable, and therefore cannot be used in the contract.

We assume that the prize cannot be paid until the beginning of the following period, as it takes the term of the period to determine the winner, and therefore is discounted. Subsequent tournaments are discounted by $\beta(x) < 1$; the leading case is β constant, but we allow for the dependence on x to nest a common racing model we introduce later, where the expected time until the winner is determined by the sum of the investments. Except in numerical results, however, all of our results focus on the case where β is constant.[1] The payoff, viewed at the beginning of the period is $\beta(x)W_i$ for the agent who wins and $\beta(x)L_i$ for the agent who loses.

The cost to the planner of making payments W_i and L_j in the event that agent i is the winner is $c(W_i, L_j)$, and is discounted by the same $\beta(x)$. We assume that c is convex and symmetric with $c(0,0) = 0$ and $\lim_{W\to\infty} c(W, L) = \infty$; we will sometimes specialize to the case where $c(W_i, L_j) = c(W_i) + c(L_j)$ (util payouts with risk averse agents and constant cost of resources to the planner) or $c(W_i, L_j) = c(W_i + L_j)$ (monetary payouts from a costly pool).

2.2 Principal's Dynamic Program

We solve the dynamic contract using recursive techniques. At the start of any tournament each agent i has promised utility σ_i. Some of this is paid as prizes, and some as future utility in subsequent tournaments. Denote by σ_i^j the future payoff to player i if j wins. The planner solves

$$\begin{aligned}
V(\sigma_0, \sigma_1) &= \max_{\sigma_i^j, W_i, L_i, x_i} \left\{ u(x_0, x_1) + \beta(x) \sum_i p(x_i, x_j)(-c(W_i, L_j) + V(\sigma_0^i, \sigma_1^i)) \right\} \\
& \quad s.t. \\
x_i &\in \arg\max_{\hat{x}} \left\{ \beta(x) \left(p(\hat{x}_i, x_j)(W_i + \sigma_i^i) + p(x_j, \hat{x}_i)(L_i + \sigma_i^j) \right) - \hat{x}_i \right\} \\
\sigma_i &= \beta(x) \left(p(x_i, x_j)(W_i + \sigma_i^i) + p(x_j, x_i)(L_i + \sigma_i^j) \right) - x_i \\
\sigma_i^j &\geq \underline{P}/(1-\beta)
\end{aligned}$$

The first constraint is incentive compatibility for the agents; it guarantees that choosing x_i in period t is part of a subgame perfect Nash equilibrium strategy. The second constraint guarantees that the planner does, in fact, deliver on the promise of σ_i. The final constraint ensures that the planner does not promise a future payoff that cannot be delivered. We will denote optimal choices when the current state is (σ_0, σ_1), for instance for W_0, by $W_0(\sigma_0, \sigma_1)$

[1] When $\beta(x)$ is a constant, this discounting of payments is immaterial; it is equivalent to paying less (by a factor of β) at the start of the period.

3 Symmetry, Asymmetry,and State Dependence

In this section we consider solely the case where $c()$ is strictly convex and $\beta(x) = \beta < 1$. In that case the incentive constraint can be replaced by the first order condition

$$\beta p_1(x_i, x_j)(W_i + \sigma_i^i - L_i - \sigma_i^j) = 1$$

We assume that, for any $W_i + \sigma_i^i - L_i - \sigma_i^j$, there is a unique solution to these two first order conditions.

We seek to understand the evolution of unequal promises to the agents, even starting from symmetry. Note that the agents are symmetric, and it could be, as in the standard form of p used in the racing literature, that it is "better," in an incentive sense, to have the agents closer to equal footing. Define

$$G_\sigma(\varepsilon) = V(\sigma + \varepsilon, \sigma - \varepsilon)$$

The following lemma shows that this function is either maximized at zero, or is differentiable and concave at zero.

Lemma 1 *Suppose for $i \in \{0, 1\}$ that $x_i(\sigma, \sigma) > 0$ and for each agent $W_i(\sigma, \sigma) > \underline{P}$. Then zero is either a local maximum of G_σ and G_σ is concave and differentiable at zero, or zero is a local minimum of G_σ.*

Proof. Since G is symmetric around zero, zero is either a local minimum or a local maximum. Suppose it is a local maximum. We need to show that G_σ is concave and differentiable at zero. Define the lower function $\hat{G}_\sigma(\varepsilon)$ by

$$\hat{G}_\sigma(\varepsilon) = u(x_0(\varepsilon), x_1(\varepsilon)) + \beta \sum_i p(x_i(\varepsilon), x_j(\varepsilon))(-c(W_i(\varepsilon), L_j(\sigma, \sigma)) + V(\sigma_0^i(\sigma, \sigma), \sigma_1^i(\sigma, \sigma)))$$

where $x_0(\varepsilon)$,$x_1(\varepsilon)$, $W_0(\varepsilon)$, and $W_1(\varepsilon)$ solve

$$\begin{aligned}
1 &= \beta p_1(x_i, x_j)(W_i + \sigma_i^i(\sigma, \sigma) - L_i(\sigma, \sigma) - \sigma_i^j(\sigma, \sigma)) \\
\sigma + \varepsilon &= \beta \left(p(x_0, x_1)(W_0 + \sigma_0^0(\sigma, \sigma)) + p(x_1, x_0)(L_0(\sigma, \sigma) + \sigma_0^1(\sigma, \sigma))\right) - x_0 \\
\sigma - \varepsilon &= \beta \left(p(x_1, x_0)(W_1 + \sigma_1^1(\sigma, \sigma)) + p(x_0, x_1)(L_1(\sigma, \sigma) + \sigma_1^0(\sigma, \sigma))\right) - x_1
\end{aligned}$$

Note that $\hat{G}$ is differentiable at zero by the implicit function theorem, and $\hat{G}_\sigma \le G_\sigma$, with equality at zero, since $\hat{G}_\sigma$ describes a feasible payoff for all ε. Now since $G(0) = \hat{G}(0)$ we can write

$$\frac{G(\varepsilon) - G(0)}{\varepsilon} = \frac{G(\varepsilon) - \hat{G}(\varepsilon)}{\varepsilon} - \frac{\hat{G}(\varepsilon) - \hat{G}(0)}{\varepsilon}$$

Consider ε in the open set where $G(\varepsilon) < G(0)$ (by $G(0)$ being a local max). Suppose $\varepsilon > 0$. Then the LHS is negative, and the first term of the RHS is positive (since $G(\varepsilon) > \hat{G}(\varepsilon)$). The second term can be made smaller than δ for any δ by taking ε near enough to zero, since the derivative of $\hat{G}(0)$ is zero. That means that the LHS must be smaller than δ in absolute value. An analogous argument for $\varepsilon < 0$ shows that, in fact, G is differentiable at zero. ■

An immediate consequence of the Lemma is that equal promised utility can only be delivered to the two agents if the agents are receiving the same current payout. The reason is that, if the current payout differs, there is a first order gain in costs from equalizing payouts. According to Lemma 1.there is no first order loss in offsetting the change by spreading continuation utilities when agent i wins.

Corollary 2 *Suppose the conditions of Lemma 1 are satisfied for some σ. Then if $\sigma_i^i = \sigma_j^i = \sigma$, $W_i = L_j$.*

This result will be used to place restrictions on the circumstances under which symmetry can be maintained. First we define the set of possible symmetric utilities.

Definition 3 *Let S be the set of symmetric continuation utilities, i.e. $S = \{\sigma_0, \sigma_1 > \underline{P}/(1-\beta) : \sigma_0 = \sigma_1)$.*

The next theorem shows that contracts that preserve symmetry can only occur for very special cases; it requires that agents do not both give effort in consecutive periods.

Theorem 4 *Suppose a subset of S is an ergodic set for the stochastic process generated by the race. If the continuation utilities are in S at some time $t-1$ and if $x_i > 0$ for all i at time t, $x_i = 0$ for all i at time $t-1$.*

Proof. Suppose a subset of S is an ergodic set and $x_i > 0$ for both agents with promised utilities in t, with $x_i > 0$ at $t-1$ for some i.

First we show that both agents have $W_i > 0$ at time t. If at time t an agent has $\sigma_i^j \geq \sigma_i^i$, then $W_i - L_i > \sigma_i^j - \sigma_i^i$ in order for the agent to give effort, and so $W_i - L_i > 0$ and so $W_i > \underline{P}$. At most one agent has $\sigma_i^j < \sigma_i^i$ (since the continuation utilities are symmetric). If that agent had $W_i = \underline{P}$, then he has L_i also at the boundary since $L_i \leq W_i$, which implies that his payments are worse in both the winning and losing states than agent j. But then it would be impossible for both agents to be receiving the same promised utility at t. Therefore it must be the case that both agents have $W_i > \underline{P}$.

Since the conditions of Lemma 1 are satisfied at the promised utilities σ, σ chosen at time t, both agents received the same payout at time $t-1$ by the Corollary to Lemma 1. In order for them to have received the same expected payout at $t-1$, since continuation utilities are symmetric and the payouts are identical in at least one state, the payouts must be identical in both states, and therefore neither agent gives effort at $t-1$. ■

Corollary 5 *No point in S can be a steady state of repeated tournament model with both players giving effort.*

Note that this is true even if the race "prefers" equal inputs by the agents, for instance in a Poisson patent race we compute below, or one where the planner values efforts in a complementary way. Even if $V(\sigma_0, \sigma_1)$ has a unique maximum at some σ, σ, the contract can not stay there; it must visit other states where the planner's payoff is lower.

The theorem shows that in many situations, starting from symmetry, you have to eventually favor someone. Intuitively, the best way to do this is state dependent, in that you favor winners, in order to get incentives. The following shows that it cannot be optimal to exit symmetry to a state that is unconditional of the winner's identity.

Proposition 6 *Suppose $(\sigma, \sigma) \in S$. It can't be that $\sigma_0^i(\sigma, \sigma) = \sigma_0'$ and $\sigma_1^i(\sigma, \sigma) = \sigma_1'$, $\sigma_0' \neq \sigma_1'$.*

Proof. Without loss, let $\sigma_0' > \sigma_1'$. Since the agents get the same payoff at σ, σ, the only possible orderings of total values are

$$W_1 + \sigma_1' \geq W_0 + \sigma_0' \geq L_0 + \sigma_0' \geq L_1 + \sigma_1'$$

or

$$W_0 + \sigma_0' \geq W_1 + \sigma_1' \geq L_1 + \sigma_1' \geq L_0 + \sigma_0'$$

Consider replacing σ_0^1 with σ_1' and σ_1^1 with σ_0'. Since V is symmetric, this leaves future payoff to the planner unchanged; he just interchanges the names and grants the higher continuation to the winner, instead of always granting it to agent 0.

To keep current effort and total payouts for both agents the same, replace W_1 with $\hat{W}_1 = W_1 + \sigma_1' - \sigma_0'$, and L_0 with $\hat{L}_0 = L_0 + \sigma_0' - \sigma_1'$. We will show that such a replacement is feasible and lowers current costs for the planner.

To see that it is feasible, note that clearly $\hat{L}_0 > L_0 \geq \underline{P}$, and that $\hat{W}_1 + x = W_1 + y \geq L_1 + y$, so $\hat{W}_1 \geq L_1 \geq \underline{P}$.

It also lowers costs, since it makes payouts closer together in the state where the change is made. To see this, note that

$$\hat{W}_1 - \hat{L}_0 = W_1 - L_0 - 2(\sigma_0' - \sigma_1')$$

If $\hat{W}_1 - \hat{L}_0 \geq 0$, it is immediate that payouts are closer from the fact that $W_1 - L_0 \geq 0$ under either of the orderings possible. If $\hat{W}_1 - \hat{L}_0 < 0$, then the payments could only be further apart if

$$2(\sigma_0' - \sigma_1') - (W_1 - L_0) > W_1 - L_0$$

which would imply that

$$\begin{aligned} \sigma_0' - \sigma_1' &> W_1 - L_0 \\ L_0 + \sigma_0' &> W_1 + \sigma_1' \end{aligned}$$

which is not possible under either of the orderings. ■

State dependence is a natural feature of the optimal contract. In the next section, however, we show that this need not lead to long run asymmetry between the effort levels of the agents.

4 Repeated Racing

A leading application of the repeated tournament model is as a contract to reward innovators, say firms, that work to develop innovations of value. Here the convexity of costs is more naturally modelled as coming through a convex cost of total payments awarded, because the principal has a cost of raising the funds required to reward the innovators. This might be a government's convex cost of raising funds through taxation, or a firms convex costs of raising liquidity to contract with other firms for the production of the innovation. Another interpretation is that the planner has convex costs of delivering profits through monopolization of the invention. Gilbert and Shapiro (1990) argue that when the strength of the patent can be chosen, many environments imply a convex relationship between the level of profits delivered and the social cost of the delivery. They use this result to argue that patents should be long and narrow. Here we focus on the impact of dynamic contracts that smooth this convexity when agents recur, requiring rewards in markets that will be monopolized under patents to come in the future.

In particular, in this section we assume that $c(W_i, L_j) = c(W_i + L_j)$. In this case it is essential that we also impose a lower bound on payments; otherwise the planner can simply spread both agent's payoffs indefinitely, keeping the total payment in each state constant. The result is arbitrarily large incentives at no additional cost. In this section we impose the limitation that $\underline{P} > -\infty$. Let $P^i = W_i + L_j$ be the total payment if i wins. Further, suppose that $c(P)$ diverges for some finite level of funds; one can imagine this as a particular kind of budget constraint that preserves differentiability.[2] We will denote this level by $\bar{P}$

4.1 Long Run Behavior

Consider a situation where the planner is paying one agent an amount above the lower bound, so $P_t^i > 2\underline{P}$. Consider paying the agent less today by an amount Δ, and more in each state in the following period by Δ/β. Since this is a feasible perturbation for $\Delta \geq 0$, it must be the case that

$$0 = \arg\min_{\Delta \geq 0} c(P_t^i - \Delta) + \beta p c(P_{t+1}^i + \Delta/\beta) + \beta(1-p)c(P_{t+1}^j + \Delta/\beta)$$

which implies that

$$c'(P_t^i) \leq pc'(P_{t+1}^i) + (1-p)c'(P_{t+1}^j)$$

[2] While this doesn't strictly fit the model described initially, it is easy to modify all of the previous results on symmetry and state dependence to fit this arrangement.

Intuitively, marginal costs can not be too high today, or the planner could deliver less in today, and deliver uniformly more in the following period, improving the planner's outcome while keeping the agent unchanged both in delivered utility and in incentives to give effort.

Clearly, if $P_t^i = 2\underline{P}$, marginal cost at t is as low as expected marginal cost could possibly be at $t+1$. Therefore the partial inverse Euler equation must hold, implying c' is a submartingale. Since c' is a submartingale, it must either converge almost surely (if its expectation does not diverge), or diverge on some paths. The following lemma shows that paths for which the marginal cost diverges are paths where the prize is converging to the upper bound.

Lemma 7 *The sequence of prizes P_t converges almost surely.*

Proof. All paths are such that either (1) $\liminf c'(P_t) = \infty$, (2) $\limsup c'(P_t) < \infty$, (3) $\limsup c'(P_t) = \infty$ but $\liminf c'(P_t) < \infty$. It is immediate That paths of type (1) have P_t converging to $\bar{P}$. We will show almost sure convergence for type (2) paths, and show that type (3) paths are not optimal.

For the stochastic process implied by the optimal path, consider a modified stochastic process that is bounded by replacing every sequence of prizes with a constant sequence as soon as the sequence goes over the bound. This modified process is a bounded submartingale, and therefore converges almost surely; since it also coincides with the original stochastic process for all paths as in (2) that never surpass the bound, those paths must converge almost surely.

For type (3) paths, take some period t with P_t such that $c'(P_t)$ is large. Then the benefit of lowering payments in that period, and raising them in a subsequent period where the payment is near $\liminf c'(P_t)$ is arbitrarily large. Even though this may result in lost effort, the benefit of effort is bounded, so this must be an improvement for $c'(P_t)$ large enough, implying a better strategy. Therefore such paths cannot exist. ■

The following lemma shows that convergence of P_t implies that the outcome must converge to equal effort levels

Lemma 8 *Suppose P_t converges to some $\hat{P}$. Then the agents efforts are converging to a constant, $x_i = x_j = x/2$.*

Proof. Let G_i be the total discounted payoff (payment plus continuation value) if the agent wins, and B_i be the amount if he loses. For large t,since P_t

is converging to some $\hat{P}$, G_i+B_j and G_j+B_i are both converging to $\hat{P}/(1-\beta)$. But then it is immediate that $G_i - B_i$ is converging to $G_j - B_j$, since the difference between them is the same as the difference between $G_i + B_j$ and $G_j + B_i$. Therefore both agents choose the same level of efforts according to the FOC in (). ■

We see that, even though a certain sort of state dependency is inevitable in the optimal contract, and it comes hand in hand with asymmetry of utility levels, it none the less leads to equal efforts by the two agents in the long run.

5 Repeated Tournaments in Firms

In this section we interpret the agents as risk averse employees and the principal as a firm that employs them. We assume that $c(W_i, L_j)$ can be written as the sum of strictly convex functions, i.e.$c(W_i) + c(L_j)$.

Tournament theory has traditionally been applied to labor market interactions of this sort. A lesson of that literature is that asymmetry is a force generating incentives, and therefore may support good outcomes. The same is true here. A parallel literature has theorized about the possibility that asymmetry in wages may be bad for firm performance, as agents feel mistreated, and take counterproductive actions (citations).

5.1 Inverse Euler equations and Wage Differentials

In this section we introduce some notation to describe agents utility in a concise way. Suppose that some agent, say 0 receives u_0 in a given period. If he wins the next tournament he receives u_{00}; if he loses he is gets u_{01}. By contrast, agent 1 receives u_1, followed by u_{11} if he wins and u_{10} if he loses.

Fixing the efforts in the next period, all other wages, and the total utility received, it must be the case that $u_{00} - u_{01}$ and $u_{11} - u_{10}$ are both constant (to keep effort next period constant). We can perturb u_0 by $\beta\Delta$ so long as both the next period wages are perturbed (in the opposite direction) by Δ in both states. With all these things held constant, the probability of agent 0 winning can be fixed at p.

Suppose wages are interior, for instance if utility is unbounded below.

The planner's problem requires that these wages are chosen such that

$$\begin{aligned}0 &= \arg\min_{\Delta} c(u_0 + \beta\Delta) + \beta p c(u_{00} - \Delta) + \beta(1-p)c(u_{01} - \Delta) \\ &= \arg\min_{\Delta} c(u_1 + \beta\Delta) + \beta p c(u_{10} - \Delta) + \beta(1-p)c(u_{11} - \Delta)\end{aligned}$$

the first order conditions imply that

$$\begin{aligned}c'(u_0) &= pc'(u_{00}) + (1-p)c'(u_{01}) \\ c'(u_1) &= pc'(u_{10}) + (1-p)c'(u_{11})\end{aligned}$$

Note that these are the usual inverse Euler equations, if you take the view that c is the inverse function of utility. As such, fixing any wages earned outside the job, the function $c(u)$ measures the wages paid to the employee.

Suppose utility is log, so c is exponential. Then

$$\begin{aligned}c(u_0) &= pc(u_{00}) + (1-p)c(u_{01}) \\ c(u_1) &= pc(u_{10}) + (1-p)c(u_{11})\end{aligned}$$

so

$$c(u_0) - c(u_1) = p(c(u_{00}) - c(u_{10})) + (1-p)(c(u_{01}) - c(u_{11}))$$

Without loss of generality let $u_0 \geq u_1$. It is easy to show from the inverse euler equations that $u_{00} > u_{10}$. We have that

$$|c(u_0) - c(u_1)| \leq p|c(u_{00}) - c(u_{10})| + (1-p)|c(u_{01}) - c(u_{11})|$$

so the resource differential is increasing.

Note that if we have $CRRA$ utility with coefficient of relative risk aversion smaller than one, then $c()$ follows a submartingale and therefore the inequality continues to hold. However, if the coefficient is big enough, the result does not hold, since the inequality in $c()$ is reversed.

5.2 Numerical Results

We compute a numerical example with $\beta = .96$, $u(x) = 10\sqrt{x}$, $c(x) = e^x$, and a lower bound of zero utility for agents. We use the common racing form of $p(x_i, x_j) = (x_i + \kappa)/(x + 2\kappa)$, for small κ. We choose the discount factor to set the length of a period to a year. We choose the constant multiplying the square root of x to get a reasonable level of innovation. We choose $c()$

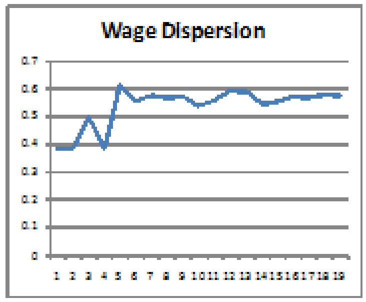

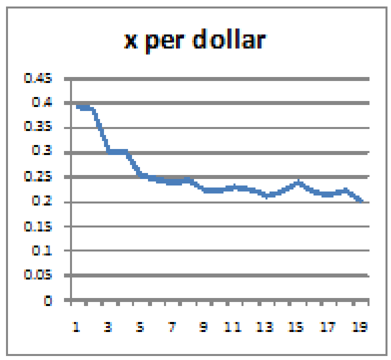

to match log utility, and include a lower bound on utility to show that our results from the inverse Euler equation can generalize to cases without the possibility of grim outcomes for the agents. We compute three values, as a function of the cumulative number of innovations. The pictures represent averages over 100 simulations.

Wage dispersion in the absolute value of the difference between the wages paid to the two agents:

And the level of innovation per dollar spent:

Notice that, as firms age (and grow larger, if one assumes that the successful innovations raise the size of the firm) they pay more dispersed wages, but are less innovative, even per dollar spent. The notion that larger firms have more wage dispersion is documented in Davis and Haltiwanger (1991). That they are less innovative per dollar spent on research has been documented many places, perhaps most notably Bound et al (1982).

Note that Davis and Haltiwanger (1991) stress that wage dispersion between plants actually falls with size. Here, wage dispersion is zero for the smallest (youngest) plants, but again falls in the long run as all firms converge to the same point. The model does not require rising wage dispersion across firms as firms age (except due to the fact that the smallest firms are all alike by assumption), and therefore rationalizes both elements that they report.

This example has the feature that the value function of the firm is maximized at symmetry, since the form of p implies that the planner gets incentives more inexpensively when agents give similar levels of efforts.

5.3 Implications of Symmetry and Asymmetry for Firms

The notion that symmetry might be better than asymmetry can be part of this model for some functional forms, as we introduced numerically above. Intuitively, imagine a case where the principal values inputs from multiple team members in a complementary way. When promised utilities diverge, getting incentives on both agents may become difficult. Note that this force toward good outcomes at symmetry is despite the fact that the model has no force by which workers feel explicitly "mistreated" by inequality.

Further, note that the general result shows that even if a firm begins at symmetry, and even if the planner gets its highest value at symmetry, it won't stay there, if the planner demands effort in every period. This is true for the example computed above; the intuition is simply that the peak at symmetry makes the first order cost of getting incentives low. If one observes many firms operating such contracts, the firms in relatively symmetric states have higher values than ones in less symmetric states, but still the only force at work is a classic tournament force, where inequality is serving as an incentive device. In a sense, tournament theory alone is generating a sense in which the firm does better when the contract calls for it to pay relatively symmetric wages.

6 Conclusions

(TBA)

References

[1] (To be added)