Recursive Contracts, Firm Longevity, and Rat Races: An Experimental Analysis☆

Peter Bardsley[a], Nisvan Erkal[a,*], Nikos Nikiforakis[a], Tom Wilkening[a]

[a]*Department of Economics, University of Melbourne, VIC 3010, Australia.*

Abstract

This paper reports the results from a laboratory experiment which investigates the structure of contracts that emerge in overlapping-generation firms where future ownership is a perquisite of employment. Workers in the young generation are offered employment contracts designed by the firms' owners who belong to the old generation. When old, employed workers are granted ownership rights as long as the firm continues to operate. In line with theoretical predictions, the results indicate that as firm longevity increases, the recursive nature of the contracts leads to a rat race characterized by low wages, high effort levels, and rent dissipation. These results have important implications for the optimal management of long-lived firms such as partnerships.

JEL Classifications: C91, D02, D21, D86, D92;

Keywords: Overlapping-generations models; Recursive contracts; Rat races; Experiments

1. Introduction

Many firms are ongoing organizations made up of individuals who join when young and leave when old. In such firms, young workers often move on to executive roles at a later stage of their career. As executive roles typically entitle individuals to a share of the firm's profits, a young worker's decision to join a firm may depend not only on the current rewards, but also the prospect of future rents that might arise through internal promotion (e.g., Baker, Gibbs and Holmstrom, 1994a, 1994b; Ferrall, 1996; Landers et al., 1996).

☆We thank the co-editor (Joerg Oechssler), two anonymous referees, Charles Noussair, Charles Plott, Katerina Sherstyuk, Larry Samuelson, Nori Tarui, conference participants at the 2009 North-American ESA Conference (Tuscon, AZ, USA), the 5th Annual Conference on Economic Growth and Development (ISI, Delhi, India), and seminar participants at the University of New South Wales and University of Tasmania for valuable comments. We gratefully acknowledge the financial support of the Australian Research Council.

*Corresponding author

Email addresses: p.bardsley@unimelb.edu.au (Peter Bardsley), n.erkal@unimelb.edu.au (Nisvan Erkal), n.nikiforakis@unimelb.edu.au (Nikos Nikiforakis), tom.wilkening@unimelb.edu.au (Tom Wilkening)

Contracts in such organizations are recursive. The contracts that young workers are willing to accept today depend on the expected value of owning the firm in the future, which in turn depends on the contract that they will be able to impose on the next generation of workers if promoted. The recursive nature of the contracts implies that factors that affect the expected value of owning the firm in the future, such as the expected longevity of the firm, may play a critical role in determining the contracts that will emerge in equilibrium.

This paper reports the results from a laboratory experiment investigating the structure and efficiency of contracts that may emerge in overlapping-generation (OLG) firms where future ownership is a perquisite of employment. A typical example of such organizations are partnerships where the ownership of non-fungible assets (e.g., corporate skills, reputation) is eventually passed vertically within the firm – from the older to the younger generation.[1] A few empirical studies have provided evidence that new workers in partnerships tend to work long hours for relatively low salaries (e.g., Ferrall, 1996; Landers et al., 1996). This finding suggests that recursive contracting can lead to rat races.[2]

The advantage of using a laboratory experiment to address our research question is that we are able to focus our attention on the role of recursive contracting and exclude other forces that may affect the efficiency of contracts in OLG firms. For example, one cannot rule out the possibility that the rat races reported in the aforementioned empirical studies are due to adverse selection or moral hazard.[3]

As a theoretical framework for our experiment, we present a model based on Bardsley and Sherstyuk (2006). In particular, we present a simplified version of their model with homogeneous agents and complete information, which is amenable for testing in the laboratory.[4] In the model presented in section 2, firm owners — members of the older generation — design the contracts offered to the young generation of workers. Workers who accept contracts will become owners of their firm as long as the firm continues operating. In this environment, the expected value of owning a firm depends on the firm's longevity, which is determined by the probability that the firm continues operating in the future.[5]

[1]Partnerships with such features are common in many fields including law, accounting, investment banking, management consulting, advertising, and some branches of medicine.

[2]In fact, all the empirical studies we found on rat races use data from partnerships. In economics, the term rat race is typically used to describe situations in which an individual is enticed to overwork, and where rents from this work are dissipated.

[3]Akerlof (1976) was the first to show that rat races may emerge if workers' abilities are unobservable to the employer. See also Stiglitz (1975), Miyazaki (1975), Landers, et al. (1996), and Andersson (2002) for models with adverse selection. In these models, workers may work inefficiently long hours to signal their higher ability. Models with moral hazard emphasize that promotion tournaments, used to provide workers with incentives not to shirk, may result in inefficiently high effort levels. See, for example, Gibbons and Murphy (1992), Ferrall (1996), Holmstrom (1999), and Baker, Choi and Gulati (2006).

[4]Bardsley and Sherstyuk (2006) develop a more general model than the one considered here with heterogeneous agents and adverse selection. The probability of being promoted is a choice variable in their model and one of the instruments that the firm can use to screen workers. In our framework, we assume that workers are promoted automatically as long as the firm survives to the next period. We treat firm longevity as an exogenous variable and focus on how it affects optimal contract design.

[5]As we explain in section 2, an increase in firm longevity can also be interpreted as an increase in the probability that a worker is promoted or a decrease in the discounting of future rents. We relate the discounting of future rents to firm longevity as it is more amenable to experimental testing.

The model reveals that when the expected value of owning the firm is high, the unique subgame-perfect equilibrium is a rat race characterized by inefficiently high levels of effort and the dissipation of firm rents.[6] The reason for this inefficiency is the inability of owners to write inter-generational contracts. Since the next (unborn) generation of workers is not privy to the original contract, the recursive structure implicitly includes a third party in the negotiation whose interests are not represented and whose actions are non-contractible. The additional rents offered from the next generation induce workers to accept lower wages or exert higher effort than they would in a single-period firm. Once wages are reduced to zero, owners offer contracts with inefficiently high levels of effort that workers are willing to accept. The model thus suggests a negative relationship between the expected longevity of a firm and efficiency.

We explore the relationship between firm longevity and efficiency using an OLG labor-market experiment. Subjects take on the roles of firm owners and potential workers. Owners make offers to potential workers specifying a wage and a level of effort that they must exert if they accept the contract. Effort is enforceable. Workers who accept an employment contract become owners in the next period if the firm continues to operate. The probability the firm continues to operate is our treatment variable. In particular, we consider two experimental treatments that differ in the expected longevity of firms. In the first treatment (*Efficient*), firm longevity is expected to be short and the model predicts an efficient outcome. In the second treatment (*Rat-Race*), firm longevity is expected to be longer and the theoretical prediction is a rat race.

The results from our laboratory experiment are consistent with the theoretical predictions. In the Efficient treatment, contracts converge to the efficient effort level and the corresponding wage level. Wages act as the primary contracting instrument. In contrast, in the Rat-Race treatment, markets converge to a rat-race equilibrium with inefficiently high levels of effort, low wages, and high inequality in earnings between owners and workers. The convergence of the theoretical predictions in both treatments is notable given the complexity of the decision problem, the presence of high payoff inequality, and the fact that contracts in the early periods of the experiment are efficient and fair (i.e., they equalize the earnings of owners and workers) in both treatments.

To the best of our knowledge, our experiment is the first to document the emergence of rat races in a laboratory environment. In addition, the paper contributes to the literature on OLG experiments by being the first to study contract design in OLG firms. Previously, an OLG structure has mainly been used to test macroeconomic theories.[7] A few exceptions are Aliprantis and Plott (1992), Offerman, Potters and Verbon (2001), and Chermak and Krause (2002), who study competitive equilibria, cooperation, and environmental issues, respectively, using OLG experiments. We are also not aware of any empirical evidence linking rat races to firm longevity.

Our findings suggest that recursive contracts may help to understand behavior observed in many organizations. The closest example of the organizations in our design

[6] Rebitzer and Taylor (2007) also use an OLG framework to explore the organizational structure of firms. They model the emergence of "up-or-out" promotion contests in large law firms and argue that the "up-or-out" promotion structure enables firms to prevent attorneys from grabbing and leaving the firm with key client relationships. The existence of a steady state is asserted in their model while our paper contains a proof of the existence of a unique subgame-perfect equilibrium.

[7] For a survey, see Duffy (2008).

are legal firms and management-consulting companies where junior hires tend to work long hours for relatively low wages.[8] However, recursive contracts may lead to rat races in any environment where decision-making power accrues with promotion and rents can be appropriated from young workers. For example, rat races appear to be common in environments where skills are passed from old to young and contracts are recursive. Modern day doctors are an example of this case, with some medical specialties taking over ten years to complete training. The apprenticeship system of medieval and pre-modern Europe is another example. Apprentices actually paid substantial amounts to be apprenticed, and were bound to work for a subsistence wage for a number of years (Wallis, 2008). Recursive contracts may also help explain the long hours worked by assistant professors in academia where tenured faculty decides on the fate of their junior colleagues.

The paper proceeds as follows. In section 2 we present the theoretical framework for our experiment. We discuss our experimental design in section 3 and present the experimental results in section 4. Section 5 concludes.

2. Theoretical framework

This section provides a theoretical framework for our experiment. We present a model based on Bardsley and Sherstyuk (2006) with homogeneous agents and complete information.

2.1. The model

We consider a discrete-time model with short-lived agents and a single long-lived firm. At each time t, there are n individuals born into the economy. Each individual lives for two periods. Utility is separable between periods and there is no discounting. All individuals are homogeneous, risk neutral, rational, and act to maximize their monetary payoffs.

The single firm in the economy is composed of a worker from the young generation and an owner from the old generation. At the beginning of each period, the owner designs a contract specifying a non-negative wage w for the worker and an effort level e. She makes a take-it-or-leave-it offer of this contract to a single individual in the young generation.[9,10] The effort of the worker yields deterministic output to the owner which can be sold at gross profit κe. There is complete information, with no moral hazard or adverse selection, and the contract can be enforced costlessly.

If the contract is turned down, both the owner and the worker receive their reservation utility, which is normalized to zero in both periods. If the contract is accepted, the owner

[8] In both cases, junior workers are typically paired with a senior manager who closely monitors them. Monitoring in law firms is also conducted via the billing system. In Australia, for instance, lawyers fill in time sheets in 6-minute increments for billing purposes which are internally and externally audited.

[9] Note that the equilibrium prediction of the model is identical to a case where the owner posts the offer publicly and has $n > 1$ homogenous individuals in the young generation compete for the contract. Our model is also isomorphic to a model with one agent per generation and where the owner has full bargaining power. We model with n agents to make the relation between the theory and experiments clear.

[10] The assumption that there is a single worker is without loss of generality. If there are k identical workers, equally likely to inherit the firm, then the probability of inheritance is $\frac{1}{k}$ while the firm value increases k times. Thus, the expected value of future ownership is invariant to k.

exerts no effort and receives the surplus $V = \kappa e - w$. The agent exerts effort e and receives utility $w - c(e)$, where $c(e)$ is the cost of effort. We assume that $c(e)$ is smooth and convex, that $c(0) = 0$, and that $0 \leq c'(0) < \kappa$. These assumptions imply that there exists an effort level $e > 0$ such that $\kappa e > c(e)$, which ensures that first-best effort is positive. We also assume that effort is bounded above: there is an effort level $\bar{e} > 0$ such that $c(e) \to \infty$ as $e \to \bar{e}$.[11] As is usual in the Principal-Agent literature, we assume that if a worker is at the participation margin, then they will accept the contract.

The key benefit of employment in the model is the possibility of becoming the owner in the future. With probability α, the firm continues to operate and the worker in period t inherits the firm in period $t+1$. As owners have both the right to choose the contract for the next generation and the right to all residual claims, the model is recursive, with the promise of future residual rights being used to entice workers to participate today.

We interpret α as the firm's survival probability so that $\frac{1}{1-\alpha}$ is the firm's expected lifetime, but there are other interpretations which fit our environment. In the context of a long-lived partnership or firm, a useful interpretation of α is the probability of internal promotion in an up-or-out system. In thinking about the transfer of skills as in apprenticeships, $(1 - \alpha)$ may be the probability of external discovery, which eliminates the monopoly rents enjoyed by the industry. It may also, of course, represent pure time discounting.

We assume there is no lending or borrowing and that individuals are born with no assets other than their endowment of labor. As workers cannot borrow, these assumptions impose a non-negativity constraint on wages.[12] Individuals have rational expectations and anticipate the future equilibrium when choosing their actions.

Since all individuals are assumed to act to maximize their monetary payoff, the owner will choose $w \geq 0$ and $e \geq 0$ to maximize her rent

$$V = \kappa e - w, \tag{1}$$

subject to the individual rationality constraint and the non-negativity constraint on wage. Individual rationality requires that

$$w + \alpha V^{+} \geq c(e), \tag{2}$$

where we write V^{+} for the next period's rent. Noting that the individual rationality constraint will always bind and eliminating w from (1), it follows that

$$V = \alpha V^{+} + \kappa e - c(e). \tag{3}$$

Non-negativity requires $w \geq 0$ which, when combined with (1), implies that

$$V \leq \kappa e. \tag{4}$$

Combining equations (3) and (4), the rent obeys the Bellman equation

$$V = \max_{e \geq 0} \min \left[\kappa e, \alpha V^{+} + \kappa e - c(e) \right]. \tag{5}$$

[11] We make this assumption to exclude the possibility of Ponzi equilibria with ever increasing intergenerational transfers and arbitrarily large effort levels. The upper bound $\bar{e}$ must be finite, but it can be large.

[12] We discuss the implication of relaxing the non-negativity constraint at the end of the section.

This recursion determines the dynamics of the contract. It also reflects the basic dilemma of the worker who may be prepared to accept a harsh employment contract today if they expect that they can impose a similar contract on the next generation tomorrow.

2.2. Equilibrium contract

We show in the Appendix that there is a unique subgame-perfect equilibrium, that this equilibrium is stationary, and that the value function is given explicitly by

$$V = V(\alpha) = \max_{e \geq 0} \min \left[\kappa e, \frac{\kappa e - c(e)}{1 - \alpha} \right]. \tag{6}$$

In order to understand this equation, notice that it builds in two constraints. The first is the wage non-negativity constraint $V \leq \kappa e$, which requires that the rent that the owner can extract from the firm is bounded by current period production. The second is the value constraint $V \leq \frac{\kappa e - c(e)}{1-\alpha}$, which requires that the rent is bounded by the total value of the firm: the per-period surplus $\kappa e - c(e)$ multiplied by the expected lifetime $\frac{1}{1-\alpha}$.[13] Equation (6) states that these are the only constraints on rent extraction and that jointly they determine the entire dynamic equilibrium.

The interaction of these constraints in (e, V) space is shown in Figure 1. The wage constraint requires that the contract point (e, V) lie below the diagonal line $V = \kappa e$. This constraint does not depend on α. The value constraint requires that the contract point lie below the total value curve $V = \frac{\kappa e - c(e)}{1-\alpha}$, which bows upwards as α increases. We plot these curves and the equilibrium contract (e, V) for a range of values from $\alpha = 0$ (the single period firm) to $\alpha = 1$ (the infinitely long-lived firm). Since $w = \kappa e - V$, the wage w can also be read directly off this diagram as the vertical distance from the contract point (e, V) to the diagonal line $V = \kappa e$.

When $\alpha = 0$, we have a single period firm that will not survive into any future period. The owner chooses the first-best effort level e^{FB}, defined by $c'\left(e^{FB}\right) = \kappa$, and a wage w that is equal to the effort cost $c\left(e^{FB}\right)$. This is the smallest wage that will induce the worker to participate. Thus, the wage constraint does not bind. As can be seen from Figure 1, the contract point lies on the vertical line $e = e^{FB}$ at the highest point that satisfies the participation constraint.

As α increases, the expected longevity of the firm increases, and the value curve shifts up to include the benefits of future ownership. The owner adjusts the contract to extract this surplus. Since $c(e)$ is convex and $c'(e^{FB}) = \kappa$, the owner at first finds it in her best interest to lower wages rather than to increase effort above the first-best level. The contract point thus shifts vertically up the efficient effort line $e = e^{FB}$. This shift continues, as α increases further, until we reach the threshold value $\alpha^* = \frac{c(e^{FB})}{\kappa e^{FB}}$ where the wage constraint begins to bind.

As α increases beyond α^*, we move from the efficient region into the rat-race region. The value of future ownership continues to increase making participation more attractive to the worker. However, the owner cannot extract this surplus by further lowering the wage, which is constrained below at 0. In fact, the owner has no efficient instrument that

[13] This constraint aggregates the individual rationality constraints of the current and all future generations.

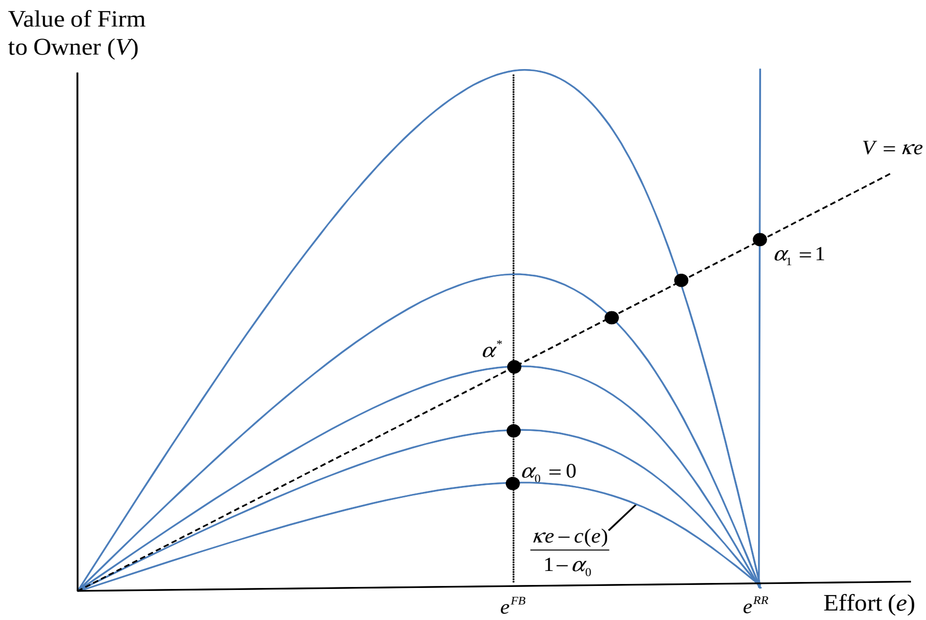

Figure 1: Equilibrium effort e and rent V for various survival probabilities between $\alpha = 0$ and $\alpha = 1$

can be used to extract this additional surplus. The best that she can do is to increase the effort level, even though this is beyond the efficient, first-best level. As can be seen from the Figure 1, the contract point moves out along the diagonal line into the rat-race region. As α approaches 1, effort approaches the fully dissipative rat-race effort e^{RR} defined by $\kappa e^{RR} = c\left(e^{RR}\right)$.

Figure 2 summarizes how the observable contract $(w(\alpha), e(\alpha))$ evolves as α increases and the firm's longevity increases. In the efficient region, when $\alpha \leq \alpha^*$, the effort is first best, but the wage declines as α increases. This is because the value of future ownership is increasing, and this future reward is a substitute for current compensation through the wage. When α increases beyond α^*, we move into the rat-race region where the wage constraint is binding, and rents generated by an increase in α can only be captured by requiring higher effort. This results in an increasingly severe rat race, with an inefficiently high effort level determined by the worker's individual rationality constraint. We summarize this characterization in the following proposition:

Proposition 1. *Let α be the firm's survival probability.*

1. *If $\alpha \leq \frac{c(e^{FB})}{\kappa e^{FB}}$, then the effort level $e(\alpha) = e^{FB}$ is efficient. The wage $w(\alpha)$ is non-negative and is decreasing in α.*
2. *If $\frac{c(e^{FB})}{\kappa e^{FB}} < \alpha < 1$, then the wage is zero and the effort level is inefficiently high with $e(\alpha) = \frac{c(e(\alpha))}{\alpha\kappa}$. Effort is increasing in α.*

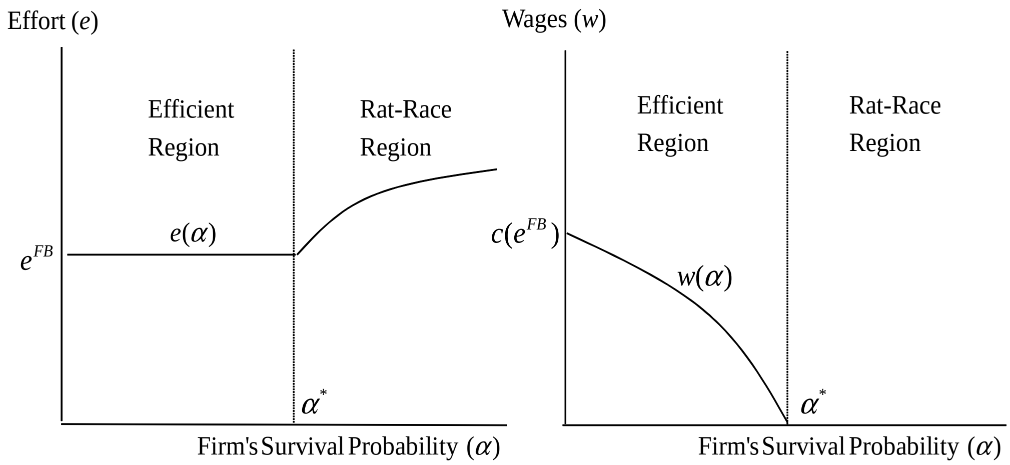

Figure 2: Wage w and effort e as functions of α

The value $V(\alpha)$ of owning the firm can be decomposed into the sum of the production surplus $\kappa e(\alpha) - c(e(\alpha))$ and a transfer $c(e(\alpha)) - w(\alpha)$ from the young generation to the old. As α increases and the firm becomes more long lived, the production surplus decreases (due to the inefficient effort level) while the transfer increases, eventually becoming the main contributor to owner profit. The value of owning the firm comes increasingly from the ability to extract a transfer from the young and intergenerational inequality increases as firm longevity increases.

As can be seen from Figure 1, although the value $V(\alpha)$ of the firm increases, it remains bounded, even though the expected lifetime $\frac{1}{1-\alpha}$ becomes infinite and there is no discounting. In fact, $V(\alpha) \to \kappa e^{RR} = c(e^{RR}) < \infty$ as $\alpha \to 1$. This is because as firm longevity increases, the firm is becoming less efficient. The per-period surplus $\kappa e - c(e)$ converges to zero more rapidly than the expected lifetime increases. The next proposition summarizes these results.

Proposition 2. *Let α be the firm's survival probability. As α increases, the value of owning the firm increases and the inequality between the old and the young generation increases. Moreover, as $\alpha \to 1$, the per-period surplus, $\kappa e - c(e)$, converges to zero and complete rent dissipation takes place.*

Note that this inefficiency would not occur if it were possible to write complete intergenerational contracts, which would guarantee efficient effort in each generation. It is precisely the inability to constrain rent seeking in the next generation which creates the inefficiency today.

Before we discuss the experimental design, we comment on the assumption of no borrowing. Following Becker (1962), it is typically argued that anti-servitude laws make it difficult (if not impossible) to borrow against human capital. Thus, when the valuation of the firm is primarily due to skill and reputation, there are likely to be limits to the amount of money that an individual can borrow against his future effort.[14]

Anti-servitude laws make the zero-wage constraint a useful benchmark for our model. Nevertheless, it is important to consider how the theoretical predictions of our model change when negative (but bounded) wages are allowed. This can happen, for example, if the firm's assets are valuable in an external market (due to a valuable client list), there is government intervention (student loans secured against a tax liability), or agents are simply born with some endowment of assets. As we will see, although allowing for negative wages improves efficiency, the exponential growth in the firm's value that occurs as the firm's longevity increases makes full efficiency unlikely in many settings.

Suppose wages can be negative up to a limit w_0 so that the owner chooses a wage amount $w \geq -w_0$. We can either assume that agents are born with an endowment w_0 or that they can borrow w_0, repaying w_0/α at the actuarially fair rate if they survive into the next period. Such borrowing relaxes the participation constraint and has the effect of shifting the diagonal line in Figure 1 up by w_0. As can be seen in the figure, this increases the range of α for which we expect to see an efficient outcome. However, as α increases, the level of debt required to support efficiency quickly becomes implausible. To see this, let $d(\alpha)$ be the level of borrowing that is sufficient to restore efficiency. We can solve for it using the individual rationality constraint since the owner sets the wage such that this constraint binds when the agent is exerting the efficient level of effort:

$$d(\alpha) = \frac{\alpha\kappa e - c(e)}{1-\alpha} \tag{7}$$

This shows that $d(\alpha)$ is increasing in α and goes to infinity as α approaches 1. Hence, supporting efficiency may mean, depending on the value of α, that the agent borrows an amount equal to many multiples of the annual profit of the firm.[15]

In summary, to the extent that agents can borrow to buy their way into the firm, the region of efficiency is increased and the rat race is ameliorated. Indeed, buying into a partnership or apprenticeship occurs in some professions, such as dentistry. We expect to see a rat race whenever the level of debt required to support efficiency exceeds the feasible level of borrowing for the agent. For the reasons discussed above, this feasible limit is likely to be small for agents endowed only with their human capital.

3. Experiment

To test the relationship between recursive contracts, firm longevity and rat races, we designed an OLG labor-market experiment. Evaluating the theoretical predictions is important as the acceptance of rat-race contracts depends on three assumptions which may be violated in practice.

[14] One can further consider debt contracts without collateral, relying on intertemporal incentives and other indirect mechanisms, but the limits to what can be done even in a world without adverse selection have been pointed out in the literature (see, e.g., Bulow and Rogoff, 1989). These limits are of course exacerbated by adverse selection and moral hazard.

[15] The only reason the agent can borrow and in expectation repay an amount that can be many multiples of the annual profit of the firm is the expectation that future agents, far into the future horizon, can borrow similar amounts and use this borrowing to accept large negative wages. Hence, the borrowing equilibrium becomes more and more like a Ponzi scheme as α increases.

First, in order to accept a rat-race contract, a worker must expect future workers to be willing to accept a similar contract when she becomes an owner. Of course, the willingness of future workers to accept such contracts will depend on their expectations about the behavior of future generations, and so on. Previous experiments have illustrated the difficulty of establishing common expectations amongst subjects (e.g., Smith, Suchanek and Williams, 1988; Lei, Noussair and Plott, 2001; Hussam, Porter and Smith, 2008; Chaudhuri, Schotter, and Sopher, 2009). In our experiment, failure to establish common expectations about future contracts is likely to prevent convergence to a rat-race equilibrium.

Second, owners and workers must be able to calculate the workers' expected payoff from accepting a contract. This requires them to consider the expected payoff from becoming an owner and the short-run payoff from the contract being offered. Previous experiments (e.g., Johnson, Camerer, Sankar, and Rymon, 2002) have provided evidence that many participants in laboratory experiments do not look into the future when making decisions. Note that the uncertainty this creates about individual rationality may also hamper the formation of common expectations.

Third, rat-race contracts imply a substantial inequality in earnings between owners and workers. Moreover, even if workers do not care about the short-term inequality in earnings, they may still decline rat-race contracts if they dislike the risk they are being exposed to with the uncertainty regarding firm longevity. Evidence from laboratory experiments suggests that many individuals care about the distribution of earnings and are risk averse (see, e.g., Holt and Laury, 2002; Sobel, 2005).

3.1. Experimental design

The experiment consists of a sequence of labor markets which we refer to as "periods." In each period, 3 firms operate in the market. Every firm is owned by an individual (the owner) who needs to hire one of 6 workers to operate the firm. Hiring is conducted through a continuous-time one-sided market in which owners offer employment contracts. Workers can observe all contracts and accept one of them at any point in time. Employers can revise the terms in their contracts as many times as they wish until a worker accepts their contract or the period ends. A period ends when all owners have employed a worker or two minutes have passed.[16]

Employment contracts consist of a wage $w \in \{0, 1, \ldots, 150\}$ and a binding effort level $e \in \{0, 20, \ldots, 100\}$. We chose to have a discrete choice set for effort in order to simplify the subjects' decision problem.[17] Workers and owners who do not agree on a contract receive zero earnings. The owner's earnings from the period are given by $1.5e - w$ while the worker's earnings are given by $w - c(e)$, where $c(e)$ is given in Table 1. The efficient level of effort which maximizes the difference between $1.5e - c(e)$ is 40.

Experimental sessions consist of 21 subjects. At any given point in the experiment, there are 3 owners, 6 workers, and 12 observers. A (virtual) dice is rolled at the end of each period to determine whether the 3 firms continue to operate. If firms continue their

[16] In the first three periods of the experiment, the labor market operated for 4 minutes in order to give time to participants to familiarize themselves with the experiment and the computer interface. In practice, an average period lasted for 49 seconds.

[17] In order to facilitate learning, owners were also provided with a profit calculator which allowed them to calculate their earnings and that of the worker for any given contract.

Units of effort	0	20	40	60	80	100
Cost of effort	0	0	20	60	100	140

Table 1: Units and cost of effort

operation, employed workers become owners in the next period. 6 of the 12 observers are randomly chosen to become workers and all remaining participants become observers.[18] If the firms do not continue their operation, roles in the following period are randomly assigned.

Observers are included in our design to reduce incentives for strategic behavior which might arise from subjects being reborn immediately, and to create variation in the set of potential workers and firms. As their name suggests, observers are able to observe the decisions made in the labor market, but they do not directly participate in it.[19] All players have access to the complete history of trades including information about the earnings of owners and workers.[20]

We deliberately chose to have an excess supply of workers in the labor market. As explained in footnote 9, the equilibrium prediction of a model with an excess supply of workers and competition is the same as having the owner make a take-it-or-leave-it offer to a single member of the young generation. An excess supply of workers, which exists in most naturally occurring labor markets, also allowed us to reduce the impact of risk attitudes and social preferences on behavior. The existence of workers with heterogeneous preferences may play less of a role in real labor markets where individuals can select into firms. Self-selection implies that rat races may emerge even if many workers are risk averse and care about inequality in earnings.[21]

3.2. Experimental treatments

The experiment consists of two treatments that differ only with respect to the probability that a firm will continue to operate in the following period (α). The values for α were chosen such that in the Efficient treatment, the efficient level of effort would be exerted in equilibrium, whereas in the Rat-Race treatment, the equilibrium level of effort would be inefficiently high. We ran 6 sessions per treatment and 126 subjects participated in each treatment.

Table 2 presents the treatments we ran and the corresponding theoretical predictions based on Propositions 1 and 2. In the Rat-Race treatment, firms have a higher survival probability ($\alpha = 5/6$). In both treatments, wages are predicted to be 0 in equilibrium.

[18] In the case that a firm did not hire a worker in the previous period, a random observer is assigned the role of the owner. This rarely happened in practice with less than 0.2 percent of owners remaining unmatched (31 out of 2382 owners across all sessions).

[19] As owners in our model may expect to be owners again, reincarnation in our environment is expected to increase efficiency if owners are sufficiently forward looking. We thus view reincarnation as working against the rat-race equilibrium outcome and in the direction of efficiency.

[20] This information may facilitate learning as it allows individuals to observe strategies that are associated with higher payoffs (e.g., Apesteguia, Huck and Oechssler, 2007; Apesteguia, Huck, Oechssler, and Weidenholzer, in press).

[21] See, for example, Roth, Prasnikar, Okuno-Fujiwara, and Zamir (1991) where fairness is eliminated in an environment with competition.

Treatment	α	Wage	Effort	Owner's Per-Period Payoff	Worker's Per Period Payoff	Per-Period Contractual Surplus
Rat-Race	5/6	0	80	120	−100	20
Efficient	2/6	0	40	60	−20	40

Table 2: Experimental treatments and predictions

Hence, the participation constraint is binding. This implies that in the Rat-Race treatment, the employer can demand a higher effort level from the worker, since the worker's probability of becoming the owner of the firm in the next period is higher. While this increases the payoff of the employer, it reduces the per-period surplus created, $1.5e - c(e)$.

3.3. Experimental procedures

All of the experiments were run in the Experimental Economics Laboratory at the University of Melbourne in September and October of 2009. The experiments were conducted using z-Tree (Fischbacher, 2007). All of the 252 participants were undergraduate students at the University who were randomly invited from a pool of more than 2000 volunteers using ORSEE (Greiner, 2004).

Upon arrival to the laboratory, participants were asked to read the experimental instructions and answer some practice questions.[22] Their answers to the practice questions were checked by the experimenter. Once the answers of all participants were checked, the experimenter read aloud a summary of the instructions. The purpose of the summary was to ensure that the main features of the experiment were common knowledge amongst the participants.

Each session lasted at least 75 minutes. After the 75 minutes passed, the experiment ended when the existing set of firms ended. Hence, each session had a random ending time.[23]

Payments to the subjects were made in cash at the end of the experiment based on the earnings they accumulated throughout the experiment. Participants were paid at the rate of E$10 = 1 AUD. In addition, each subject received a show-up fee of $15. Since payoffs during the experiment could be negative, the subjects could use the show-up fee to prevent bankruptcy during the experiment.[24] The average payment at the end of the experiment was 46.17 AUD including the show-up fee. At the time of the experiment 1 AUD = 0.80 USD.

[22]A copy of the instructions and the practice questions can be downloaded from http://www.economics.unimelb.edu.au/staff/nisvan/instructions-rat-race.pdf.

[23]We scheduled the experiments for two and a half hours and set aside an additional 30 minutes of potential experiment time to finish the last round. As a typical period lasted 49 seconds, the probability that the last firm was expected to last longer than 30 minutes was 0.1%. In practice, all sessions ended within 5 minutes of the 75-minute cutoff time.

[24]We had four cases of bankruptcy across all experimental sessions. All four cases were in the Rat-Race treatment. Subjects were given additional money and allowed to continue in the experiment. It was explained that the additional money would be subtracted from their final payment.

4. Results

The analysis of the experimental data is divided into two subsections. In section 4.1, we concentrate our analysis on the theoretical predictions stated in Propositions 1 and 2. For each result, we provide support based on descriptive statistics, figures, and statistical tests. In section 4.2, we make some observations regarding the evolution of contracts over time.

4.1. Main Results

Result 1. *In line with Proposition 1, an increase in firm longevity leads to a rat race characterized by inefficiently high effort levels and low wages.*

Figure 3 presents the evolution of the mean effort level for the accepted contracts in each treatment. As can be seen in Panel A, effort converges to the efficient level ($e = 40$) in the Efficient treatment and remains at this level for the reminder of the experiment. Effort in the Efficient treatment appears thus to be remarkably well predicted by the model. The same applies for the Rat-Race treatment. As can be seen in Panel B, effort starts at a similar level as in the Efficient treatment, but increases rapidly towards the theoretical prediction ($e = 80$). This suggests that, while efficiency and equality may be salient consideration in the determination of the initial contracts, the recursive structure of the firm leads effort to inefficient levels, as predicted by the model.

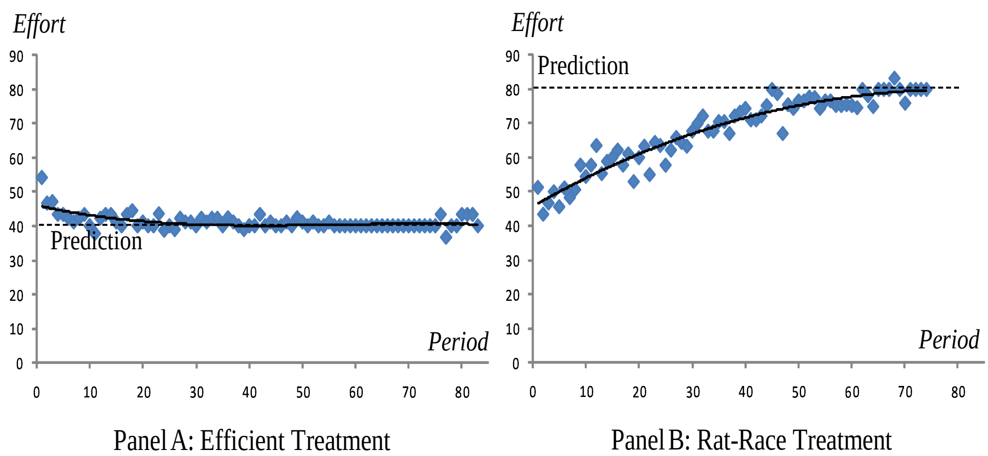

Figure 3: Average effort by period for the Efficient and Rat-Race Treatments.

Statistical support for these observations can be found in Table 3. Columns 1 to 4 present the mean effort, mean wage, and mean firm value, along with the theoretical predictions. To account for the fact that experimental sessions varied in length due to their stochastic ending, we use data from the last 30 periods in each session.

Columns 5 and 6 present the p-values from Wald tests examining whether there are significant deviations from the theoretical predictions in the Efficient and Rat-Race treatments, respectively.[25] Column 7 presents the p-values from Mann-Whitney tests examining whether there are significant differences across treatments using session averages (from the last 30 periods) as independent observations.

As can be seen in the first row of Table 3 (Columns 1-4), effort in both treatments is notably close to the theoretical prediction. Average effort in the Efficient treatment is 40.8 which is not significantly different from the predicted level of 40 (p-value = 0.80). Similarly, mean effort in the Rat-Race treatment is 74.6 which is not significantly different from the theoretical prediction of 80 (p-value = 0.08). Finally, in line with the theoretical predictions, the effort level in the Rat-Race treatment is approximately double that of the Efficient treatment, a treatment effect which is highly significant (p-value < 0.01).

	Efficient Treatment		Rat Race Treatment		Statistical tests (entries are p-values)		
					$H_0^{Efficient}$:	$H_0^{Rat\ Race}$:	H_0:
	Prediction	Mean	Prediction	Mean	Pred.=Mean	Pred.=Mean	Eff=Rat
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Effort	40	40.8 [34.7,46.8]	80	74.6 [68.6, 80.5]	0.80	0.08	< 0.01
Wage	0	15.8 [13.9, 17.8]	0	10.1 [8.1, 12.0]	< 0.01	< 0.01	< 0.01
Value of firm	40	39.5 [36.5, 42.5]	20	22.7 [19.6, 25.7]	0.75	0.08	< 0.01
Owner's profit	60	45.4 [35.6, 55.2]	120	101.8 [92.4, 111.7]	< 0.01	< 0.01	< 0.01
Worker's profit	-20	-5.9 [-18.6, 6.9]	-100	-79.0 [-92.0, -66.5]	0.03	< 0.01	< 0.01

The means and statistical tests are calculated using all observations from the last 30 periods of each experimental session. 95% Confidence Intervals for each variable are shown below each mean and are constructed from a random-effects regression specification, which includes only the treatment dummy as an explanatory variable and a session-level random effect, a worker random effect, and an owner random effect. Columns (5) and (6) present the p-values of Wald tests based on this random-effects regression. Column (7) presents the p-values of a Mann-Whitney non-parametric test with session averages as the unit of observation. All analysis conducted in Stata 12.1.

Table 3: Summary statistics, theoretical predictions, and statistical tests

Figure 4 presents the evolution of the mean wage for the accepted contracts in each treatment. As can be seen, wages start at similar levels in both treatments (around E$35) and decline over time. In the Efficient treatment, wages stabilize after approximately 30 periods at roughly E$15, while wages in the Rat-Race treatment continue to decrease over the course of the experiment. The statistical results in Table 3 reveal that the

[25] The p-values are calculated from random-effect regressions which include only the treatment dummy as an explanatory variable and a session-level random effects, a worker random effect and an owner random effect. Wald tests are conducted against the linear restrictions implied by the theoretical model.

deviation from the theoretical prediction (E$0) is significant in both treatments (p-value < 0.01). Average wages are significantly higher in the Efficient treatment than in the Rat-Race treatment (p-value < 0.01), even though the actual difference is small (E$5.7).

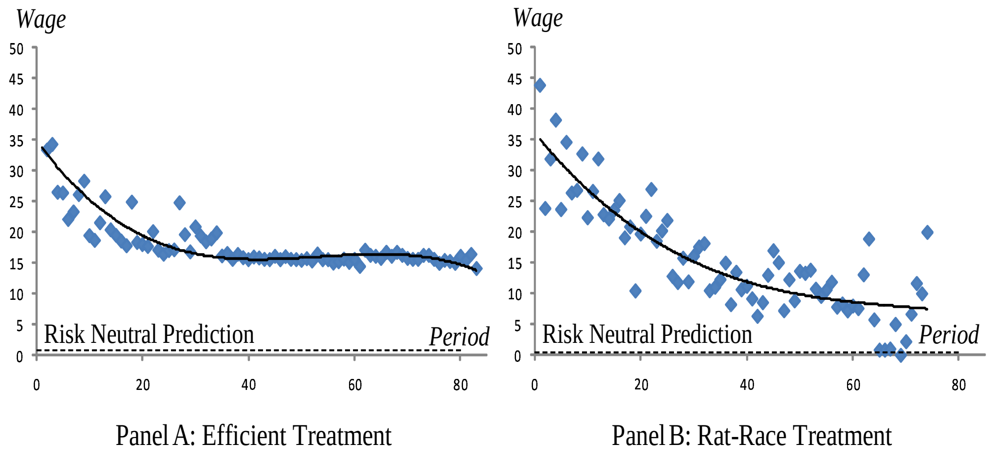

Figure 4: Average wage by period for the Efficient and Rat-Race Treatments.

One possible reason for the higher-than-predicted wages is the high degree of risk and inequality implied by the equilibrium contracts in both treatments. As discussed in Section 3, while we chose to have an excess supply of workers in order to mitigate the impact of these factors, it is not entirely surprising that workers require premia to accept a contract. In the Rat-Race treatment, the positive wages observed are consistent with workers exhibiting low degrees of risk (or loss) aversion.[26] Rationalizing the wages in the Efficient treatment in the same way, by contrast, requires an unreasonably high level of risk or loss aversion.[27] This suggests that other forces may be at play in this environment. For example, subjects may underweight the probability of the firm surviving in the Efficient treatment or may view zero profit as a norm for contracts.[28,29]

[26] Using a CRRA utility function of the form $\frac{x^{1-\sigma}}{1-\sigma}$ and assuming that an individual believes that she will always trade a contract that has the same wage and effort as the contract she is accepting in the current period, the data is rationalized with a $\sigma = .091$. Relative to most of the experimental literature which finds $\sigma \in [0, 2]$, this level of risk aversion is reasonable. Using the loss aversion model of Kahneman and Tversky (1979), where $u(x) = x$ for $x \geq 0$ and $u(x) = \lambda x$ for $x < 0$, the data can be rationalized with a loss aversion parameter of $\lambda = 1.11$. Estimates of $\lambda \in [1, 3]$ are common in the literature, again suggesting that this deviation from the prediction can be explained by loss aversion.

[27] Using a CRRA utility function, the data is rationalized with a $\sigma = 4.5$. Using the Kahneman and Tversky (1979) loss aversion model, the data can be rationalized with a loss aversion parameter of $\lambda = 4$. Both estimates lie outside of the typical range found in the experimental literature.

[28] Another potential explanation for the difference between the two treatments may be participation bias since observers spent time in the observer pool and this time is expected to be somewhat longer in the Rat-Race treatment (Lei, Noussair and Plott, 2001). However, we find no evidence that would support this hypothesis.

[29] Our instructions include two examples with a wage of 20, two examples with a wage of 40 and one example with a wage of 110. As there were no examples with a wage of zero, it is possible that subjects formed expectations based on these examples and that the instructions contributed to the higher-than-predicted wages.

Having established the formation of rat races, we next study some of their consequences.

Result 2. *In line with Proposition 2, the formation of rat races leads to a dissipation of rents. Inequality in payoffs between owners and workers increases in α.*

Recall that the per-period value of the firm is $1.5e - c(e)$. This value is not affected by wages which represent a net transfer. Given the high degree of accuracy in the effort predictions, it is unsurprising that the value of the firm is also well-predicted. As can be seen in Column 2 in Table 3, the mean per-period value of the firm in the Efficient treatment is E$39.5, which is remarkably close to the predicted E$40 and not significantly different from it (p-value = 0.75). Similarly, the mean per-period value of the firm in the Rat-Race treatment is E$22.7, which is close to the predicted E$20 and not significantly different from it (p-value = 0.08). In contrast, the per-period value of the firm is significantly different across the two treatments (p-value < 0.01). Thus, the main prediction of the theoretical model — that an increase in firm longevity can lead to a dissipation of rents — is borne out in our data.

With regards to the inequality in owner-worker payoffs mentioned in Result 2, it is useful to examine the data from the 12 individual sessions. This will also allow us to better understand how the two treatments differ in their convergence patterns. In the top panel of Figure 5, the profit of workers and owners is shown for each of the 6 sessions of the Efficient treatment. As can be seen, there is remarkably little difference in behavior across sessions with workers having a mean loss of E$6 and owners a mean profit of roughly E$45. As shown in Table 3, profits do not fully converge to the theoretical predictions (p-value ≤ 0.03). This is because the wages are higher than those predicted by the theoretical model, as described above.

As predicted by the model, the inequality between the payoffs of owners and workers is even greater in the Rat-Race treatment, with owners enjoying a mean profit of E$102 and workers having a mean loss of E$79 ($p$-value < 0.01). As shown in the bottom panel of Figure 5, behavior in all sessions seems to have followed a similar pattern with payoff inequality increasing over time.

4.2. The evolution of contracts

In section 4.1, we saw that, while contracts converge towards the equilibrium, there also appear to be some interesting dynamics. Most notably, contracts are efficient in the early periods of the Rat-Race treatment, and slowly move towards becoming inefficient. In this section, we investigate the evolution of contracts in the Rat-Race treatment.

To simplify the discussion, we concentrate our attention mainly on a single session.[30] Figure 6 presents the evolution of accepted contracts in the second session of the Rat-Race treatment. The horizontal axis indicates the period in which the contract was offered and accepted, while the vertical axis indicates the wage specified in the contract. We distinguish between the contracts based on the effort level specified: triangle trades are those with an effort of 40, diamond trades are those with an effort of 60, circle trades are those with an effort of 80, and rectangle trades are those with an effort of 100. The

[30] Behavior in this session is similar to that in other sessions in the Rat-Race treatment.

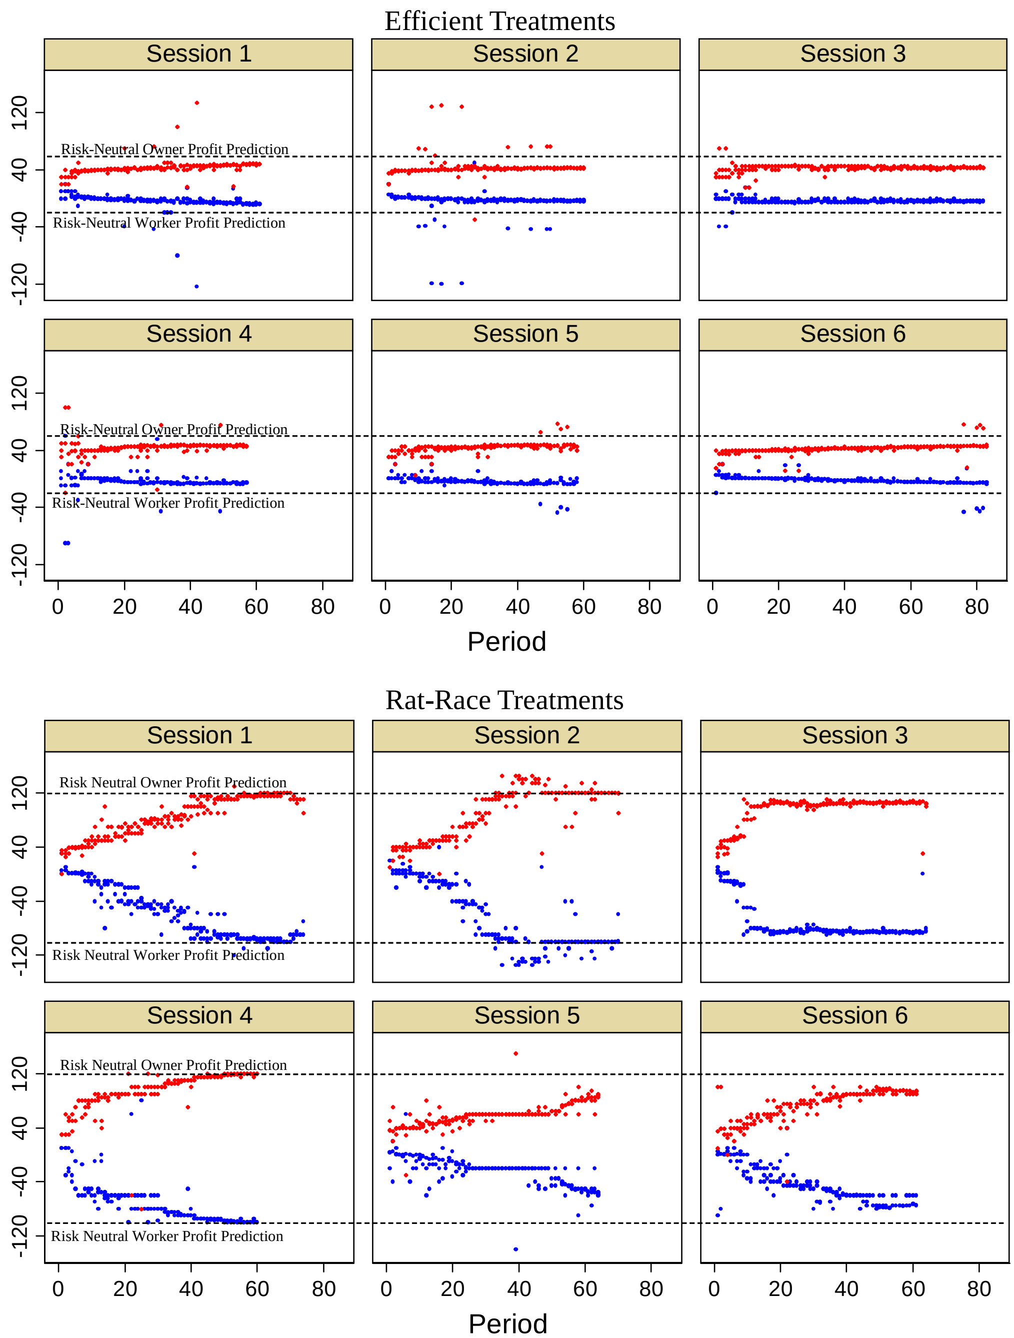

Figure 5: Profits of the owner and worker for each of the twelve treatments.

time series is truncated at period 65 as there is little variation in trades beyond this point.

As can be seen, up to period 10, most workers accept contracts that specify an efficient level of effort and a wage of 20. Unlike the equilibrium contract, this contract does not generate losses for the worker even in cases where the firm does not survive. The fact that this contract is observed in the early periods of all sessions across both treatments suggests that contracts in which wages perfectly offset the cost of effort are salient in this market.

Between periods 10 and 20, contracts specifying an effort of 40 continue to be the ones most commonly accepted. During this time period, however, we observe the wages specified in these contracts to be falling. By period 20, most accepted contracts involve wages close to zero. This indicates that, as assumed in our model, owners are willing (and able) to extract rents from the workers. Further, consistent with the model, the effort specified in the contracts continues to increase over time: between periods 20 and 30, most contracts specify an effort of 60 — an inefficient level of effort. Interestingly, wages do not remain at zero, but increase back up to 20. Thereafter, a new sequence of contracts are accepted with inefficient levels of effort and wages again falling toward zero until the process begins again.

The process of decreasing wages and the sudden (temporary) increase in wages when the effort demanded increases is likely to be due to both our design and the efficiency trade-off of the wage and effort instruments. Whereas wages could fall in units of 1 increment, effort could only be adjusted in units of 20 in the experiment. Thus, it seems likely that owners used wages to satisfy workers' individual-rationality constraints when shifting from a lower effort to a higher one. In addition, a jump in effort reduces the overall surplus of the firm. This may have made it difficult to satisfy the workers' individual-rationality constraints while making the high-wage/high-effort contract more desirable to the firm than the low-wage/low-effort contract that was being traded.[31]

Convergence rates in the Rat-Race treatment show some variation across sessions. Specifically, sessions 5 and 6 in the Rat-Race treatment display a slower rate of convergence. This can partially be explained by reinforcement learning. In sessions where firms ended with a high frequency in early periods, the jump from low-effort/low-wage contracts to a higher-effort/higher-wage contract was slow to occur. In Table 4, we present the survival rate of firms based on the number of times established firms were replaced by new firms by the 10th, 20th, 40th and last period. As can be seen in the first column, the probability that old firms continue operating is markedly low in the first 10 periods for sessions 2, 5, and 6 in the Rat-Race Treatment (56%, 56% and 67%, respectively, compared to 100%, 89% and 89% for the other sessions).

[31] The slow evolution of wages is consistent with adaptive learning models where an individual is learning about the expected value of becoming an owner over time. However, given the large expected surplus that workers were receiving early on in the experiment, we would have predicted that the wages path would fall rapidly if learning was only about the continuation value of the firm. It is also possible that individuals are adjusting their reservation utility — the amount of expected surplus they require in order to satisfy their individual rationality constraint — downward in response to contracts that they see others take. This adjustment of the reservation utility could occur if individuals reference contract is changing in response to observations about other players' behavior. As we have limited data to distinguish between different models, we remain cautious about how to interpret the wage path.

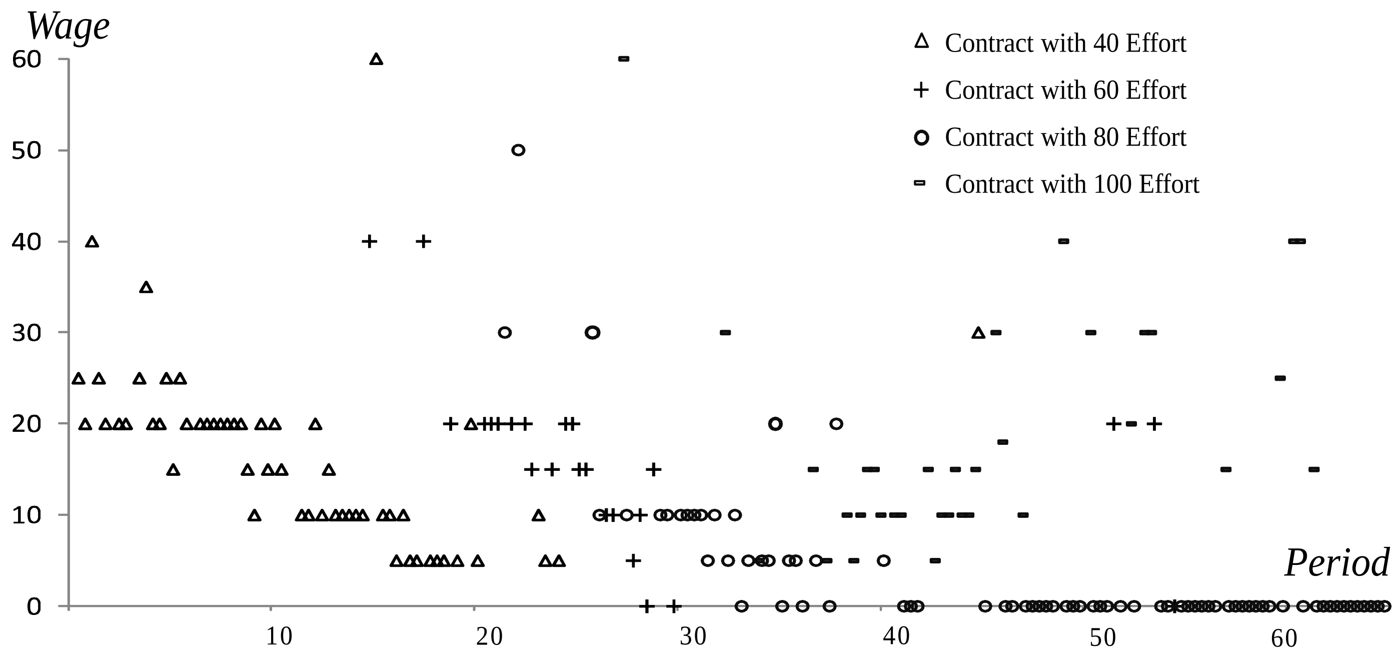

Figure 6: Wage and Effort of Contracts Traded in Session 2 of the Rat-Race Treatment.

	Period 10	Period 20	Period 40	Last Period
Session 1	1.00	.84	.92	.88
Session 2	.56	.79	.79	.80
Session 3	.89	.95	.87	.86
Session 4	.89	.79	.79	.78
Session 5	.56	.74	.85	.83
Session 6	.67	.74	.79	.85

Table 4: Empirical Probability of Continuation in Rat-Race Treatments

5. Discussion

In many firms, workers' decisions are motivated by the prospect of being promoted to a senior role that entitles them to a share of the firm's future profits. Since the profits in these firms depend on the contracts that future workers will be willing to accept, the contract design is recursive. As we show in the theoretical framework presented in section 2, if the expected value of owning the firm is sufficiently high, then recursive contracts can give rise to a rat race with high effort levels, low wages and rent dissipation. Such rat races are likely to be more common in firms that are expected to continue operating for longer periods of time.

Findings from a laboratory experiment we designed to examine the relationship between recursive contracts, firm longevity and rat races are consistent with the theoretical predictions. Despite the computational complexity of the decision problem, the exposure to risk, the presence of payoff inequalities, and the well-documented difficulty of establishing common expectations between individuals, we find that increases in firm longevity lead to rat rates characterized by inefficiently high levels of effort and substantial inequality in earnings between owners and workers.

The finding that firm longevity can have adverse effects on efficiency is in contrast to findings in previous studies that have associated firm longevity with higher efficiency arising from the ability of individuals to enforce cooperation across generations of workers.[32] Unlike earlier models, where members of the young generation can punish those of the old generation, our setting gives full contracting power to the old generation. As a result, reducing the discount rate lowers efficiency in our environment.

The key feature of our design is that owners are residual claimants who have been internally promoted and who are given full contracting authority, although they cannot sell the firm to an outsider. This feature appears to be common in partnerships in which ownership of non-fungible assets (such as reputation or corporate skills) is eventually passed vertically within the firm from the older to the younger generation, and in which workers' effort is closely monitored. Partnerships with such features are common in many fields, including law, accounting, investment banking, management consulting, advertising, and medicine. Partners in these firms have full, if implicit, contracting authority vis-a-vis the promotion of associates. The reputation of the firm is embodied in its people and culture. Since selling the firm to outsiders would incur a high reputational

[32]See Cremer (1986), Kandori (1992), Salant (1991) and Smith (1992).

cost, the promise of future promotion is credible.[33]

The experimental environment and mechanism used in this experiment isolated the recursive contracting channel by mitigating or eliminating secondary forces that are also likely to influence real labor markets. By introducing an excess supply of workers, we likely reduced the effect that other-regarding, risk and time preferences may have on the overall market. While we believe that competition is an inherent characteristic of high reputation firms, its impact on the interaction between recursive contracting and preferences is an open question. Our design also eliminated differences in worker's productivity which would introduce adverse selection into the model. Future research can usefully investigate the importance of these factors in the emergence of rat races.

Given the inefficiency of the rat-race equilibrium, it would also be interesting to consider what institutions or policy interventions might prevent rat races from emerging. Restricting wages to a minimum will not prevent rat races as owners, in response, will simply demand higher levels of effort from workers. However, imposing a cap on the maximum number of hours could be effective. Reducing the rents available to owners and increasing the number of senior hires from outside the firm may also weaken the rat race. These remedies are a topic for future studies.

References

[1] Apesteguia, J., S. Huck, and J. Oechssler. 2007 "Imitation: Theory and experimental evidence," *Journal of Economic Theory*, 136, 217-235.

[2] Apesteguia, J., S. Huck, J. Oechssler, and S. Weidenholzer. 2010. "Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust," *Journal of Economic Theory*, 145(5), 1603-1617.

[3] Akerlof, G A. 1976. "The Economics of Caste and of the Rat Race and Other Woeful Tales," *Quarterly Journal of Economics*, 90, 599-617.

[4] Aliprantis, C.D. and C.R. Plott. 1992. "Competitive Equilibria in Overlapping Generations Experiments," *Economic Theory* 2, 389-426.

[5] Andersson, F. 2002. "Career Concerns, Contracts, and Effort Distortions," *Journal of Labor Economics*, 20(1), 42-58.

[6] Ando, M., and H. Kobayashi. 2008. "Intergenerational Conflicts of Interest and Seniority Systems in Organizations," *Journal of Economic Behavior & Organization*, 65, 757-767.

[7] Baker, S., C.J. Choi, and M. Gulati. 2006. "The Rat Race as an Information-Forcing Device," *Indiana Law Journal*, 81, 53-82.

[8] Baker, G., M. Gibbs, and B. Holmstrom. 1994a. "The Wage Policy of a Firm," *The Quarterly Journal of Economics*, 109(4), 921-55.

[9] Baker, G., M. Gibbs, and B. Holmstrom. 1994b. "The Internal Economics of the Firm: Evidence from Personnel Data," *The Quarterly Journal of Economics*, 109(4), 881-919.

[10] Bardsley, P. and K. Sherstyuk. 2006. "Career Paths in Partnerships: An OLG Agency Model," *Topics in Theoretical Economics*, 6(1), 1-37.

[11] Bar-Isaac, H. 2007. "Something to Prove: Reputation in Teams," *Rand Journal of Economics*, 38, 495-511.

[12] Becker, G.S. 1962. "Investment in Human Capital: A Theoretical Analysis," *Journal of Political Economy*, 70, 9-49.

[13] Binmore, K. 2007. "Game Theory: A Very Short Introduction," Oxford University Press, Oxford.

[14] Bulow, J. and K. Rogoff. 1989. "A Constant Recontracting Model of Sovereign Debt," *Journal of Political Economy*, 97(1), 155-78.

[33] Tadelis (1999) and Mailath and Samuelson (2001) discuss models in which reputation can be traded in the external market, not just within the firm.

[15] Chaudhuri, A., A. Schotter, B. Sopher. 2009. "Talking Ourselves to Efficiency: Coordination in Inter-Generational Minimum Effort Games with Private, Almost Common and Common Knowledge of Advice," *Economic Journal*, 119(534), 91-122.

[16] Chermak, J. M. and K. Krause. 2002. "Individual Response, Information, and Intergenerational Common Pool Problems," *Journal of Environmental Economics and Management*, 43(1), 47-70.

[17] Cremer, J. 1986. "Cooperation in Ongoing Organizations," *The Quarterly Journal of Economics*, 101, 33-49.

[18] Duffy, J. 2008. "Macroeconomics: A Survey of Laboratory Research," mimeo.

[19] Ferrall, C. 1996. "Promotions and Incentives in Partnerships: Evidence from Major U.S. Law Firms," *The Canadian Journal of Economics*, 29(4), 811-827.

[20] Fischbacher, U., 2007. "Z-Tree: Zurich Toolbox for Readymade Economic Experiments," *Experimental Economics*, 10, 171-178.

[21] Gibbons, R., Murphy, K.J. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns," *Journal of Political Economy*, 100(3), 468–505.

[22] Greiner, B., 2004. "The Online Recruitment System ORSEE 2.0 — A Guide for the Organization of Experiments in Economics," University of Cologne, Working paper series in economics 10.

[23] Hammond, P., 1975. Charity: Altruism or Cooperative Egoism? In: Phelps, E. (Ed.), Altruism, Morality, and Economic Theory. Russel Sage Foundation, New York, pp. 115-131.

[24] Hussam, R. N., D. Porter and V. L. Smith. 2008. "Thar She Blows: Can Bubbles Be Rekindled with Experienced Subjects?" *American Economic Review*, 98(3), 924–937.

[25] Holmstrom, B. 1999. "Managerial Incentive Problems: A Dynamic Perspective," *Review of Economic Studies*, 66, 169–182.

[26] Holt, C. A. and S. K. Laury. 2002. "Risk Aversion and Incentive Effects," *American Economic Review*, 92(5), 1644–1655.

[27] Johnson, Eric J., Camerer, Colin F., Sen, Sankar and Rymon, Talia. 2002. "Detecting Failures of Backward Induction: Monitoring Information Search in Sequential Bargaining," *Journal of Economic Theory*, 104(1), 16–47.

[28] Kandori, M. 1992. "Repeated Games Played by Overlapping Generations of Players," *Review of Economic Studies*, 59, 81-92.

[29] Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263-291.

[30] Landers, R.M., J. B. Rebitzer, and L. J. Taylor. 1996. "Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms," *The American Economic Review*, 86(3), 329-348.

[31] Lei, V., C.N. Noussair and C.R. Plott. 2001. "Non-Speculative Bubbles in Experimental Asset Markets: Lack of Common Knowledge of Rationality vs. Actual Irrationality," *Econometrica*, 69(4), 831-859.

[32] Mailath, G. J., and L. Samuelson. 2001. "Who Wants a Good Reputation?" *Review of Economic Studies*, 68, 415-441.

[33] Miyazaki, H. 1975. "The Rat Race and Internal Labor Markets," *Rand Journal of Economics*, 8(2), 394–418.

[34] Morrison, A. D. and W. J. Wilhelm. 2004. "Partnership Firms, Reputation, and Human Capital," *American Economic Review*, 94(5), 1682-1692.

[35] Offerman, T., J. Potters, H.A.A. Verbon. 2001. "Cooperation in an Overlapping Generations Experiment," *Games and Economic Behavior*, 36(2), 264-275.

[36] Rebitzer, J. B. and L. J. Taylor. 2007. "When Knowledge is an Asset: Explaining the Organizational Structure of Large Law Firms," *Journal of Labor Economics*, 25(2)., 201-229.

[37] Roth, A. E., V. Prasnikar, M. Okuno-Fujiwara, and S. Zamir. 1991. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review*, 81, 1068-1095.

[38] Salant, D. J. 1991. "A Repeated Game with Finitely Lived Overlapping Generations of Players," *Games and Economic Behavior*, 3, 244-259.

[39] Smith, L. 1992. "Folk Theorems in Overlapping Generations Games," *Games and Economic Behavior*, 4, 426-49.

[40] Smith, V.L., G. L. Suchanek and A. Williams. 1988. "Bubbles, Crashes and Endogenous Expectations in Experimental Spot Asset Markets," *Econometrica*, 56, 1119-1151.

[41] Sobel, J. 2005. "Interdependent Preferences and Reciprocity," *Journal of Economic Literature*, 93, 392-436.

[42] Stiglitz, J. E. 1975. "Incentives, Risk and Information: Notes towards a Theory of Hierarchy," *Bell Journal of Economics*, 6(2), 552–579.

[43] Tadelis, S. 1999. "What's in a Name? Reputation as a Tradable Asset," *American Economic Review*, 89, 548-563.
[44] Wallis, P. "Apprentice and Training in Premodern England," *The Journal of Economic History*, 68, 832-861.

Appendix: Uniqueness of the Subgame-Perfect Equilibrium

We consider the version of the model where one individual is born at each date, and the owner makes a take it or leave it offer to the worker. To reduce notation we set $\kappa = 1$.

The contract space is $\mathcal{S} = \{(e, \pi) : 0 \le e \le \bar{e},\ \pi \le e\}$, where we write $\pi = e - w$ for the current period profit to the owner. A history is a finite sequence $h = (c_0, c_1, c_2, \ldots)$ where $c_t = (e_t, \pi_t) \in \mathcal{S}$ is a contract offered at date t. At history h the worker chooses a set $\mathcal{A} \subset \mathcal{S}$ of contracts that will be accepted, and the owner chooses a contract $c \in \mathcal{S}$.

We consider subgame perfect equilibria. In such an equilibrium the owner and the worker solve the interlinked Bellman equations

$$V_h = \max\left[0, \sup_{(e,\pi)\in\mathcal{S}} \chi_{\mathcal{A}}(e, \pi)\, \pi\right],$$
$$U_h = \max\left[0, \sup_{\mathcal{A}\subset\mathcal{S}} \chi_{\mathcal{A}_+}(e, \pi)\left(e - \pi - c(e) + \alpha V_{h|c}\right)\right].$$

The supremum over an empty set is, by convention, $-\infty$. The set $\mathcal{A}_+ = \{(e, \pi) \in \mathcal{A} : \pi \ge 0\}$ is the set of contracts in $\mathcal{A}$ that are not dominated by the owner's outside option. The indicator functions are defined by $\chi_{\mathcal{A}}(e, \pi) = 1$ if $(e, \pi) \in \mathcal{A}$ and $\chi_{\mathcal{A}}(e, \pi) = 0$ if $(e, \pi) \notin \mathcal{A}$; and similarly for $\chi_{\mathcal{A}_+}$. The history $h|c$ is the history h with the contract $c = (e, \pi)$ appended at the end. Since we are at an equilibrium, the suprema are actually attained, and we could write max rather than sup.

By subgame perfection, the worker will always accept any individually rational contract. Thus

$$\mathcal{A}(h) = \left\{(e, \pi) \in \mathcal{S} : e - \pi - c(e) + \alpha V_{h|c} \ge 0\right\}.$$

Consider the owner's problem. The set $\mathcal{A}$ is bounded above by the functions $\pi = e$ and $\pi = e - \pi - c(e) + \alpha V_{h|c}$. Since $c(e) \to \infty$ as $e \to \bar{e}$, $\mathcal{A}$ has compact upper sections and there exists a point in $\mathcal{A}$ that maximizes π. It is clear that this maximum must occur at a value of e in the interval where the constraint $e - \pi - c(e) + \alpha V_{h|c} \ge 0$ binds, since otherwise a small increase in e would increase π. But $e - \pi - c(e) + \alpha V_{h|c}$ is strictly concave on this interval, so the maximum must be unique. It may, however, be below the owner's reservation value of 0. But the owner's optimal choice is unique in this case as well: she chooses the outside option with payoff 0.

We thus conclude that the strategies of both the owner and the worker at history h, and in particular the value V_h, are entirely determined by the number $\alpha V_{h|c}$. Thus the equilibrium history is completely determined by the sequence of values $V_t = V_{h_t}$, where h_t is the history truncated at date t, and these values satisfy the recursion

$$V_t = \max\left[0, \max_e \min\left[e, \alpha V_{t+1} + e - c(e)\right]\right].$$

Since $V_{t+1} \geq 0$, and the owner can always choose the efficient contract yielding a strictly positive surplus, it follows that $V_t > 0$ and we can write the simpler recursion

$$V_t = \max_e \min\left[e, \alpha V_{t+1} + e - c(e)\right].$$

It is convenient to write $Z_t = \alpha V_t$ for the expected valuation prior to knowing whether the firm continues. To simplify notation, we write Z for Z_t and Z_+ for Z_{t+1}. We will consider how V is determined by Z_+.[34] We consider separately the cases where the non-negativity constraint $\pi \leq e$ does and does not bind.

If the non-negativity constraint does not bind then $V = v_0(Z_+)$, where the function $v_0(z)$ is determined by

$$\begin{aligned} c'(e) &= 1, \\ v_0(z) &= z + e - c(e) \end{aligned}$$

so

$$v_0(z) = e^{FB} - c\left(e^{FB}\right) + z.$$

Note that $v_0(0) = e^{FB} - c\left(e^{FB}\right) > 0$, and $v_0'(z) = 1$.

If the non-negativity constraint does bind then $V = v_1(Z_+)$, where the function $v_1(z)$ is determined by

$$\begin{aligned} e &= z + e - c(e), \\ v_1(z) &= e \end{aligned}$$

so

$$v_1(z) = c^{-1}(z)$$

is the inverse to the cost of effort function.

It follows that the valuation is determined by $V = v(Z_+)$, where

$$v(z) = \begin{cases} v_0(z) & \text{if } z \leq z^{FB} \\ v_1(z) & \text{if } z \geq z^{FB} \end{cases}$$

where $z^{FB} = c\left(e^{FB}\right)$. To see this, notice that $v_0(z) \geq v_1(z)$, $v_0\left(z^{FB}\right) = v_1\left(z^{FB}\right)$, and $v_0'\left(z^{FB}\right) = v_1'\left(z^{FB}\right)$. Thus $v_0(z)$ is the tangent line that supports $v_1(z)$ at z^{FB}. For $z \geq z^{FB}$, $v(z)$ is the inverse cost function, while for $z \leq z^{FB}$ it is the tangent supporting the inverse cost function at z^{FB}. Thus it is continuously differentiable, monotonic increasing with $0 \leq v'(z) \leq 1$, and concave. It is strictly non-negative and bounded, with $v(0) = e^{FB} - c\left(e^{FB}\right) > 0$ and $\lim_{z \to \infty} v(z) = \bar{e}$.

Contract dynamics are determined by the relations

$$\begin{aligned} V &= v(Z_+), \\ V &= \frac{Z}{\alpha} \end{aligned}$$

[34] We use upper case letters V, Z for equilibrium values, and lower case variables π, z, e for other variables.

and can be determined in the usual way from the loci of these two relations in (Z, V) space. Since the loci cross once, it is clear that for $0 < \alpha \leq 1$ there exists a unique stationary equilibrium. But this is the only equilibrium. Any other equilibrium can only be supported by a value trajectory that either increases without bound, violating the assumption that $e \leq \bar{e}$, or decreases without bound, violating the constraint that $e \geq 0$.

We note, in conclusion, that the steady state condition $V = v(\alpha V)$ can be written equivalently as

$$V = \max_{e} \min \left[e, \frac{e - c(e)}{1 - \alpha} \right].$$

This can be verified by considering separately the cases where the non-negativity constraint $\pi \leq e$ does or does not bind.