Training and Turnover in the Evolution of Organizations

Natalie S. Glance   Tad Hogg   Bernardo A. Huberman

Organization Science **8**, 84–96 (1997)

Abstract

We elucidate the interplay between the free rider dilemmas faced by both management and employees of a firm. Both managers and employees face a free rider dilemma: organizations can hire workers trained by competitor organizations and workers that shirk can still receive a share of the profits from production. If all organizations and employees reason similarly, the overall level of production will be very low. We construct a dynamical model of the two-level social dilemma and find, by running computer experiments, that the dynamical behavior at the two levels is closely coupled and highly path-dependent. In some cases the double dilemma is resolved for the industry as a whole and productivity increases steadily over time. In others, the organizational level dilemma remains unresolved and workers contribute at fluctuating levels. In these cases the overall productivity stays low. Our computer experiments indicate that low turnover, large enterprise size and high productivity are positively correlated, a finding that is in line with the empirical literature on turnover and training in organizations.

1 Introduction

During periods of slow growth and a weak economy, corporations often cut programs in order to maintain profitability. Training programs in particular are often targeted because employee turnover is generally higher during times of economic uncertainty [19]. Even in the best of times, organizations must decide how much to invest in on-the-job training, balancing the benefits of increased productivity against the costs of training. Because trained workers can migrate easily between competing firms, another firm can potentially benefit from the increased productivity of workers trained by the former employer without paying the costs. For example, a survey of metalwork firms in Wisconsin indicated that managers are reluctant to train their workers because they fear competitor firms will lure their employees away before their investment costs are recouped [16]. Consequently, fear of losing trained employees to competitors can lessen a company's incentive to train and lead to less investment in skills than is economically desirable [3, 1].

Ironically, numerous studies have shown that untrained workers change jobs more often [19, 18]. The negative correlation between training and turnover has been documented in a number of companies, such as the Mariott Corporation, the Florida Power Corporation, IDS Financial Services Inc., and Target [13]. All of these firms experienced boosts in retention rates after investments in various training programs. Unfortunately, many firms are reluctant to train until some degree of stability is achieved within their workforce, and their hesitation may in turn be reinforced by observed high turnover rates [18]. Can the vicious circle be broken?

An organization's decision whether or not to train its workers also affects the economy, whether or not the firm factors this variable into its decision. If all firms within an industry fail to train their workers, the whole economy suffers. Hence, training workers is a type of public good [19], a category that encompasses social dilemmas such as the support of public radio, the so-called "tragedy of the commons" [11] and recycling programs. There is a long history of interest in such problems in political science, sociology and economics [22, 12]. Any resolution to the training dilemma will depend not only on the benefits and costs associated with a particular training program, but also on the firm's expectations concerning employee turnover and the policies of competing firms.

Employees face a similar dilemma in their choice of how much they contribute to the overall productivity of the organization. If employees receive a share of the profits regardless of their contribution, some may decide to free ride on the efforts of their fellow workers. If all employees decide to do so, the company will fail. Profit-sharing and employee-ownership can exacerbate the dilemma [7], and indeed the gains from profit-sharing plans are frequently lower than expected [5, 17]. In principle, the problem could be resolved by strict management, but in practice, worker monitoring is always imperfect and employee effort can vary from high to low within the range allowed [20].

The two dilemmas on the employee and organizational levels are tightly interlaced. On the one hand, the benefits of training accrue only to the extent that employees contribute to the organization. Thus, a firm should take into account how it expects a training program to affect employe effort as well as employee turnover. On the other hand, trained workers produce at higher rates, which in turn may affect how much they contribute and how often they migrate to other firms compared to untrained workers.

Because the two dilemmas are strongly coupled, they are hard to study in a natural organizational setting. Computer simulations provide an effective way of studying such problems and their evolution. If the assumptions are clearly stated, many scenarios can be explored without the disruptions they would cause if tried in the firm.

In this paper we explore the dynamics of training and turnover in firms facing both organizational and employee-level dilemmas. First we establish a simple model that captures these conflicts and incorporates imperfect information and both worker and organizational expectations. Organizations can be both created and dissolved, and employees can move between firms, start new ones, or leave the industry for good. Next we summarize the different ways the dilemmas can unfold over time, collated from several computer experiments. For example, under one set of conditions, the double dilemma can be resolved for the industry as a whole and productivity then increases steadily over time. Alternatively, the organizational level dilemma may remain unresolved and workers may contribute at fluctuating levels. In this case the overall productivity stays low. We find a positive correlation between high productivity, low turnover and enterprise size, a relation that has also been observed in the empirical literature on training, stability, and turnover in organizations [19, 21].

Our dynamical model of training and turnover in organizations both confirms the empirical observation that the two variables are tightly interlinked and reveals how the connections might be unraveled. In addition to supporting the empirical data on firms, it provides a way to understand how the interplay between different variables, such as turnover, training, enterprise size and productivity, comes about and evolves over time.

2 Modelling Organizational and Employee Strategies

In this section, we describe our model of organizational training, individual learning, and decision-making on both the worker and organizational levels. In our model, all organizations within an "industry" produce the same good, for which there is a completely elastic demand outside the industry. This assumption means that the industry can grow indefinitely as there is no ceiling for production. Employees, or "agents," can move between organizations, within the bounds allowed by the organizations' "managers." The managers must decide whether or not to train the agents in their own organization, and the agents must decide whether or not to contribute to production.

2.1 Interwoven social dilemmas

Our model of management training and employee production is a two-level social dilemma. At the level of the agent, each individual must decide whether or not to contribute to production (a binary approximation to the continuous range of effort they can deliver). For the case of profit-sharing assumed by the model, each agent receives an equal share of its organization's total production, independent of its contribution. Each agent is tempted to free ride on the industriousness of the other agents, but if all agents do so, nothing is produced and everyone loses.

On the higher level of management, organizations must decide whether or not to train their agents. If a manager decides to train, then members of its organization learn over time, and when its members do contribute to production, they do so at progressively higher levels as time passes. However, training agents comes with a cost to the total utility produced by the organization, which managers must take into account. An organization does not want to train its agents only to have them stolen by a competitor, but if all agents receive training the entire industry is better off, garnering higher utility over time.

2.2 Expectations

Recent work on the dynamics of single organizations suffering from the agent-level social dilemma has shown that high levels of production can be sustained when groups are small or hierarchically structured into smaller groups with fluid boundaries [8, 10, 9]. The ongoing nature of the social dilemma lessens its severity if the agents take into account the future when making decisions in the present. How an agent takes into account the future is wrapped into what we call its expectations. The barest notion of expectations comes from the economic concept of horizon length. An agent's horizon length is how far it looks into the future, or how long it expects to continue interacting with the other agents in its organization. The agent's horizon may be limited by its lifetime, by its projection of the organization's lifetime, by bank interest rates, and other factors.

Here our notion of expectations differs from the standard rational expectations treatment in economics [2], which assumes agents form expectations about the future using near-perfect knowledge of the underlying model. This notion is self-consistent, but circular: the agents predict the future exactly. In our model of expectations, agents believe their present actions will affect those of others in the future. The extent of the effect depends on the size of the organization and the present level of production. The larger the group, the less significance an agent accords its actions: the benefit produced by the agent is diluted by the size of the group when it is shared among all agents. An agent that free rides can expect the effect to be very noticeable in a small group, but less so in a larger group. The reasoning is similar to that a student uses when deciding whether or not to attend a lecture she would prefer to skip. Among an audience of 500, her absence would probably go unnoticed (and if all students in the class reason similarly...). On the other hand, in a small seminar of ten, she might fear the personal censure of her professor.

In our model, the agents expect their actions will be imitated by other agents and the extent of this mimicry depends on present levels of production. An agent expects that if it decides to free ride ("defect") in a group of contributors, or "cooperators," others will eventually choose to defect as well. The agent also believes that the rate at which the switchover occurs over time depends on the fraction of the group presently cooperating. The more agents already cooperating, the faster the transition to defection. Similarly, an agent expects that if it starts cooperating in a group of free riders, others will start cooperating over time. Once again the agent believes that the rate depends on the proportion of cooperators, which in this case is very low. Our key assumption is that agents believe their actions influence contributors, or "cooperators" more than sluggards, or "defectors." This difference in influence is taken to be proportional to the fraction agents already cooperating, and is used in deriving Eq. (6).

Consider the set of beliefs the agent expects of others in the context of recycling programs. Recyling has a strong public good component because its benefits are available to all regardless of participation. Not too long ago very few towns had such programs. Perhaps you would read in the paper that a small town in Oregon had started a recycling program. Big deal. But several years later, when you read that cities all over your state have jumped onto the recycling bandwagon, then suddenly the long-term benefits of recycling seem more visible: recycled products proliferate in the stores, companies turn green, etc. Alternatively, imagine some futuristic time when everyone recycles, in fact your town has been recycling for years, everything from cans to newspapers to plastic milk jugs. Then you hear that some places are cutting back their recycling efforts because of the expense and because they now believe that the programs don't do that much good after all. You think about all your wasted effort and imagine that the other towns still recycling are reaching the same conclusion. In view of this trend, your commitment to recycling may falter.

To some extent, this set of beliefs is arbitrary. We can imagine other scenarios for which another set of expectations would be more appropriate, but there is a class of expectations for which the general conclusions of our work hold. Specifically, our model can accommodate the class of expectations for which agents believe that the strength of their influence on the amount of cooperation extends into the future as far as their horizon, decreases with the size of the group and increases *roughly* with the current proportion contributing. Perhaps agents believe instead that their influence is greatest when a certain fraction cooperates, but declines at both extremes of full cooperation and full defection. Thus, they imagine their influence grows with the fraction cooperating only when that proportion is small. Alternatively, perhaps agents believe their influence is greatest at the extremes and declines when the group is a mix of cooperators and defectors. In this case, the agents imagine their influence grows with the fraction cooperating only when that proportion is large. Both of these cases fall within the range of expectations compatible with our model.

In our interlocking model of organizational training and agent cooperation, we also extend the formulation of expectations to the organizational level. Managers decide to train or not based on the number of organizations in the industry and on the number that presently train their agents. Folded into this decision and into their expectations is the behavior of the agents that comprise a manager's organization. A manager's horizon length depends on the tenure lengths of its agents: the longer its agents stay, the longer it expects them to stay in the future, and the more reason the manager has to train them. Likewise, a manager predicts greater future value from training when more of its agents are actively contributing instead of free riding.

2.3 Conditions for cooperation

In a profit-sharing organization in which individual agents receive equal shares of the utility produced by the organization, the utility to an agent is its share minus its costs. If an agent contributes, it incurs a cost which reduces its net gain; it suffers no cost if it does not contribute, but the total production to the organization declines. That is the utility U_i to agent i in organization m is its share minus its cost c for cooperation:

$$U_i = \frac{1}{n_m} \sum_{j=1}^{n_m} b_j^m k_j - c k_i \tag{1}$$

where k_i is 1 if the agent contributes and zero otherwise, n_m is the size of the organization and b_j^m is the benefit produced by agent j when it cooperates. The individual agent utility also depends indirectly on the managerial policies of its parent organization. If an organization trains, its agents will learn over time and produce at progressively higher levels. Otherwise, the benefit of cooperation for its agents stays fixed over time. Specifically we use a linear model of learning, which is given by the differential equation

$$\frac{db_i^m}{dt} = \gamma \kappa_m \tag{2}$$

where γ is the learning rate and κ_m is 1 if the organization trains and zero otherwise.

All agents in the industry start off at the same baseline benefit for cooperation, $b_{\min}$. When agents move between organizations within the industry, they retain only a fraction of the gain in their benefit for cooperation obtained over time through training, although the benefit is not allowed to fall below the baseline level. The loss in learning when agents migrate models the incomplete transfer of knowledge between organizations, i.e.

$$b_i^l = r(b_i^m - b_{\min}) + b_{\min} \tag{3}$$

so that r gives the fraction of learning that is transferred.

The organizational utility is the total utility produced by its constituent agents minus any training costs. For each agent that contributes, the organizational utility increases by that agent's contribution. If the agent is learning over time, the agent's contribution also increases over time, but is offset in part by the costs for training that agent. It is given by

$$U^m = \sum_{j=1}^{n_m} b_j^m k_j - n_m T \kappa_m \tag{4}$$

where T is the training cost per agent.

Agents and managers use their respective utility functions to guide their decisions on whether to contribute and train. They project future earnings in accordance with their expectations and their horizon lengths. For individual agents, the criteria for cooperation was derived in [9] for a simpler model and extends easily to the present case. Individuals cooperate if their observed share of production

$$\langle b \rangle^m = \frac{1}{n_m} \sum_{j=1}^{n_m} b_j^m k_j \tag{5}$$

exceeds the critical amount

$$b_{\text{crit}}^m \equiv \frac{b_{\min}}{H\alpha} \left(\frac{n_m c - b_i^m}{b_i^m + \gamma \kappa_m H - c} \right) \tag{6}$$

where H is the evaluation horizon and α is the rate at which agents reevaluate their choices. This critical amount was derived by computing the net benefit that an individual would accrue if it decided to cooperate, based on the fraction of individuals perceived as cooperating at that time and how long the game is expected to last, as given by the horizon H. If the individual cooperates only when the benefit is positive, defects when is negative and chooses at random when the benefit is zero, the condition for cooperation can be expressed in terms of a critical size. According to this criterion, beyond a critical group size, no agent will cooperate, and below a second critical group size, all agents will cooperate. Notice that the longer the horizon, the smaller the critical group for agent cooperation. Conversely, the larger the group the larger the critical size, and the harder it is to secure voluntary cooperation. Between the two limits are two equilibrium points, one of mostly cooperation and the other of mostly defection. The group dynamics tends towards the equilibrium closest to its initial starting point. Generally, one of the equilibria is metastable, while the other is the long-term equilibrium. By metastable we mean an equilibrium that is stable against small perturbations but unstable against large ones. (An example would be a ball in a trough situated on top of a hill). If a group falls into a metastable state, it may remain there for very long times (exponential in the size of the group). Because of uncertainty the group will eventually switch over to the global equilibrium very suddenly (in time logarithmic in the size of the group), as shown in [8].

The training criterion for organizations follows by analogy. A manager trains when the observed fraction of organizations training exceeds the critical amount

$$f^m_{\text{crit}} \equiv \frac{1}{H_m \alpha_m} \left(\frac{NT - \gamma f^m_c}{\gamma f^m_c - T} \right) \tag{7}$$

where N is the number of organizations, f^m_c is the estimated fraction cooperating in the organization, and H_m and α_m are the horizon and reevaluation rate for the managers, respectively. This criterion has the following properties. Managers are more likely to train when their horizon lengths are long, training costs are low compared to the agents' learning rate, the number of organizations is small, and they estimate a large proportion of their agents to be cooperating. A manager can estimate the fraction cooperating from the production level observed by inverting the organizational utility given by Eq. (4). This estimate will differ from the actual fraction cooperating since an organization's agents may have received different amounts of training and will consequently have different benefits for cooperation. However, for the sake of simplicity, we will model the manager's estimate of the fraction cooperating using Eq. (5) as

$$f^m_c = \frac{\langle b \rangle^m}{b_{\min}} \tag{8}$$

Although this estimate somewhat overstates the amount of cooperation and worsens as the agents learn over time it captures the essential feature that the manager's perception of the workers is based on their overall production.

We intend the two conditions for action to be taken as heuristic guidelines rather than precise formulas. While the agent-level condition for cooperation was derived from the expectations sketched out earlier, its qualitative features are what interest us. We expect the heuristic form of the criteria to hold for a wide range of expectations. For some sets of expectations they may not hold, in which case a different model would then be appropriate. While these heuristics may differ from those used by real organizations, we believe that they are indicative of the qualitative behavior that one expects to see in the real world.

2.4 Fluidity

We also model the changing structural nature of industries over time. We use the term fluidity to describe the ease with which structure can change. The parameters which govern the amount of fluidity in an industry are listed in Table 1. For the purposes of our model, we consider them as given exogenously; they could also be thought of us under the control of some metalevel agent (say, some regulatory mechanism) that adjusts the fluidity parameters to optimize the overall utility of the industry or perhaps even under individual agent control.

Fluidity describes the ease with which agents can move within an organization from subgroup to subgroup, how promptly they leave the organization for another one or leave the industry completely seeking higher personal utility, and how readily they start an organization of their own. Organizations restrict structural fluidity to the extent they make it difficult for agents to join and difficult for them to leave or move within their organizations.

μ	moving threshold
η	break away threshold
Ω	entrepreneurial rate
ρ	joining threshold

Table 1: Fluidity parameters.

In our model of structural fluidity, managers control the rate at which constituent agents choose to move between organizations and the rate at which agents from a pool of agents exterior to the industry can join, but do not restrict agents from leaving. Specifically, agents move between organizations or join an organization only when invited by a manager. Agents accept or decline the invitation according to moving and joining strategies that optimize utility and take into account moving and joining costs (set at the metalevel). Say that agent i in organization m is invited to join organization l. Agent i compares its organization's production level with that of organization l. Agent i will move only if

$$\langle b\rangle^{l} - \langle b\rangle^{m} > \mu b_{\min} \tag{9}$$

where $\langle b\rangle^{m}$ is defined in Eq. (5) and $\mu < 1$. Similarly, if agent j is invited to join organization m from the outside pool of agents, the agent will join organization m only if the organization's production level exceeds the agent's costs:

$$\langle b\rangle^{m} > \rho c \tag{10}$$

with $\rho > 1$ generally.

Agents can also decide to "break away" or leave the industry for good. In our model, an agent will break away when its organization's production level falls below a lower threshold parametrized by the break away variable η:

$$\langle b\rangle^{m} < \eta c \tag{11}$$

Some (small) fraction of the time, parametrized by the entrepreneurial rate, Ω, the agent will start a new organization within the industry instead of leaving. In this fashion, the number of organizations in the industry can grow over time. The number of organizations decreases whenever all agents from one organization have left.

In previous work, we described how structural fluidity within a single organization enables agent-level cooperation [9]. In this paper, we assume that the time-scale of structural change on the organizational level is much shorter than on the industry level so that we can ignore intra-organizational fluidity and better pinpoint the effects of training and inter-organizational fluidity.

2.5 Computer Experiments

As described in detail in the appendix, the simulation of our model runs on two levels: the agent level and the organizational level. Agents wake up asynchronously according to a Poisson process described in the appendix. When they wake up, they either (1) reevaluate their decision to cooperate or not according to the condition for cooperation; or (2) reevaluate their choice to stay in their organization, or start a new organization, or break away from the industry completely.

Each manager also wakes up asynchronously, but according to a Poisson process whose mean time increases linearly with the size of its organization. This reflects both the more ponderous decision-making of larger organizations and the longer time-scales over which organizations reevaluate their decisions compared to agents. When a manager wakes up, it either (1) reevaluates its decision whether or not to train its agents; or (2) invites an agent from a competitor organization to join. In the second case, if the invited agent refuses to join, the manager invites an agent from the outside pool to join. Organizations prefer to steal away agents from competitors since they most likely produce at higher levels, as a result of training, but agents will switch only if they perceive a gain in personal utility.

This is only one of many ways to simulate such a model. Our experience running similar types of simulations indicates that one of the most important features is that the agent and managerial states be updated asynchronously [14], not synchronously, to accurately model continuous time.

3 Results

The dynamics on the organizational level mirrors the agent-level description given earlier: when the number of organizations in the industry exceeds a critical number, none train, and when it falls below another critical number, all train. Again, between these two critical sizes is a middle region with two equilibria: one in which all managers train, and one in which none train. The transition from the metastable state to the global equilibrium may not happen for a time exponential in the number of organizations and is very sudden when it finally occurs. The critical numbers depend on the learning rate of the agents and the training cost for the organizations.

However, for fluid industries in which agents can move in and out of various organizations, the critical regions for cooperation and defection shift for both agents and organizations. For agents, the critical regions shift because the size of their parent organizations changes over time. A small cooperating organization will tend to grow over time because outside agents see its high productivity. If the organization becomes too large and its agents do not receive training, eventually a transition to overall defection will take place. Once all the agents in the organization are defecting, the group's size will shrink because many (or all) will break away from the industry or move to another organization. At some point, the group will again be small enough to support cooperation. This cycle of cooperation-growth to defection-attrition and back again repeats over and over for each organization when managers do not train. The amount of cooperation within different organizations and their sizes are coupled because of the agents moving between organizations.

The critical regions also shift in time for each organization, depending on how many of its agents cooperate and how long the agents stay in the same organization. Over time, what was originally an unresolvable dilemma for the managers (so none train) becomes resolvable, and eventually the dilemma can disappear completely. The behavioral regions shift (1) as the agents' tenure lengths change; and (2) as the agents' production levels increase. The agents' tenure in a particular organization increases when agents remain loyal to the their parent organization. Generally, agents are loyal when their colleagues cooperate. Tenure lengths are short when few cooperate within an organization since agents will move often or break away. Agents' production levels increase when their parent organizations train them and when the agents cooperate among themselves.

The detailed parameter values used in the figures presented below are included in the appendix.

3.1 Dynamics of industry growth

The dynamics of agent and organizational behavior are closely coupled. Cooperation at one level encourages cooperation on the other level and similarly with defection. At both levels, metastable states can trap the industry in lower-performing states (or higher performing states). For certain parameters the industry is in the two-equilibria region on both the agent level and the organizational level. We concentrate primarily on the behavior of the industry for this regime.

The dynamics of the industry is highly path-dependent, a phenomena observed in several economic systems, particularly those influenced by technological innovation [23]. For the same initial conditions and parameter choices, the industry can evolve to a number of different states. Fig. 1 and 2 show a series of snapshots taken from the time evolution of two industries starting from the same initial conditions. Initially, both industries consist of four organizations, with eight agents each. The total number of agents in the industry is printed at the top of the schematic tree. The agents cooperate initially, as indicated by the filled lower-level circles. None of the managers are training, as indicated by the open upper-level circles (the same code, filled circles for cooperation/training and open circles for defection/no training are used for both agents and managers). Both industries grow in size at first since their agents cooperate and new agents from outside the industry are attracted by the high levels of production (increasing the size of the industry as a whole). Once an organization grows too large, its agents switch to defection and move to another organization or break away completely (decreasing the size of the industry).

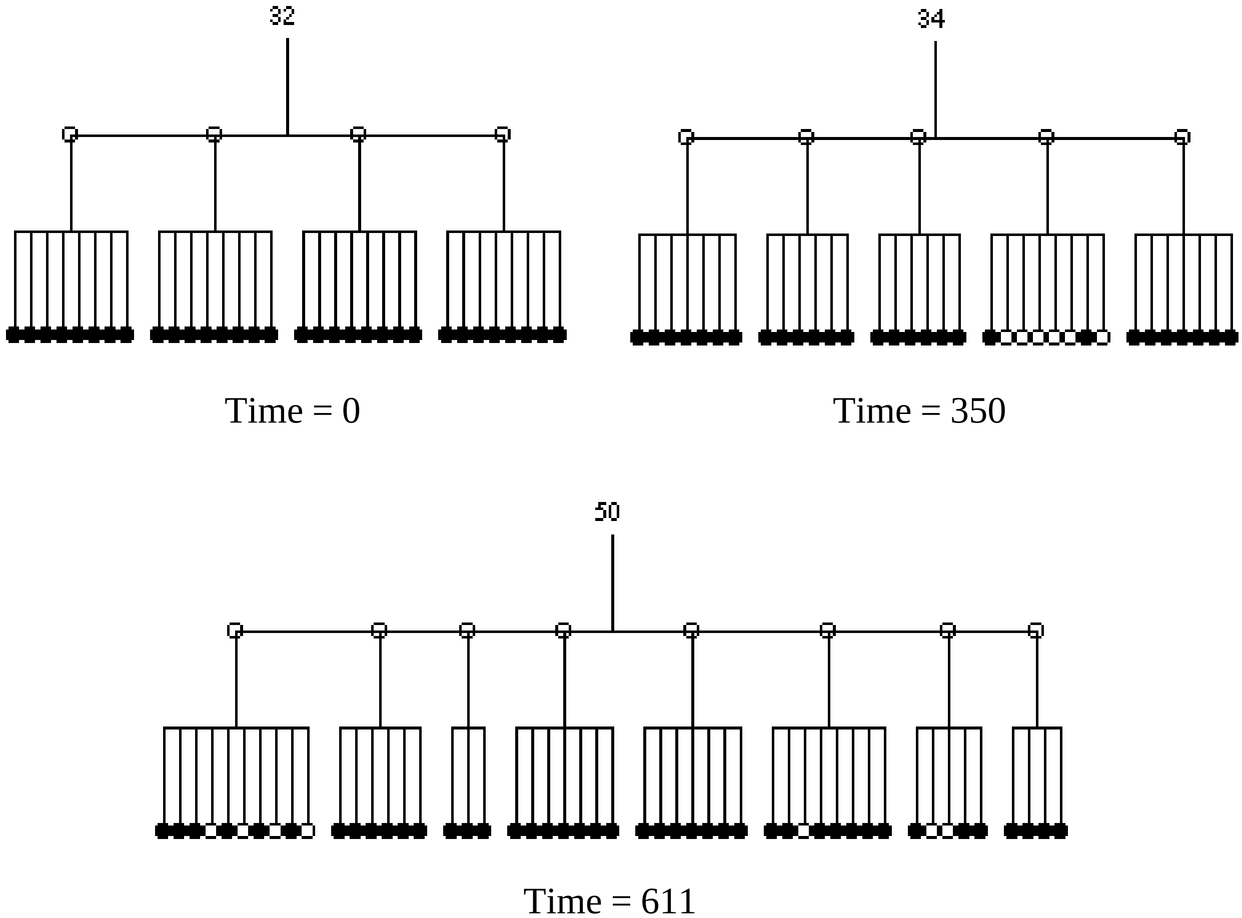

Figure 1: Snapshots of the time evolution of an industry faced with social dilemmas at both the individual agent and organizational levels. Agents must decide whether or not to cooperate knowing that they receive a share of their organization's production regardless. Organizations must decide whether or not to train knowing that the costs of training will be lost if their agents switch to another organization. The dynamics of the industry is highly path-dependent. For the same initial conditions and parameters, the industry can evolve to a number of different states. The snapshots above are taken from a simulation for which the number of organizations proliferates over time and the dilemma on the organizational level becomes untenable—there is no training of the agents. Since there is no training, the industry's utility can increase only because agents join.

The number of organizations varies stochastically: organizations die whenever all of their constituent members leave, and new organizations form because entrepreneurs strike out on their own. The balance between these two trends depends on the average rates of the various events and on chance. When the number of organizations happens to grow over time, the dilemma on the organizational level becomes untenable—the switchover to overall training never happens. Instead, the number of organizations proliferates over time and the industry tends towards a state of many organizations, each with a small number of members who cycle between states of cooperation and defection. This is the process indicated by Fig. 1. On the other hand, if the number of organizations happens to stay constant or shrink, all managers eventually decide to train their agents. In this case, the industry tends towards a state with a small number of very large, highly productive organizations. Fig. 2 depicts such an industry.

The overall utility to the industry over time depends strongly on the path the industry follows. Fig. 3 shows the abrupt deviation in overall utility between the two industries of Fig. 1 and 2. Once the organizations in the second industry switch over to the training equilibrium, the industry's utility rises steadily as the industry attracts more agents who learn and produce more over time.

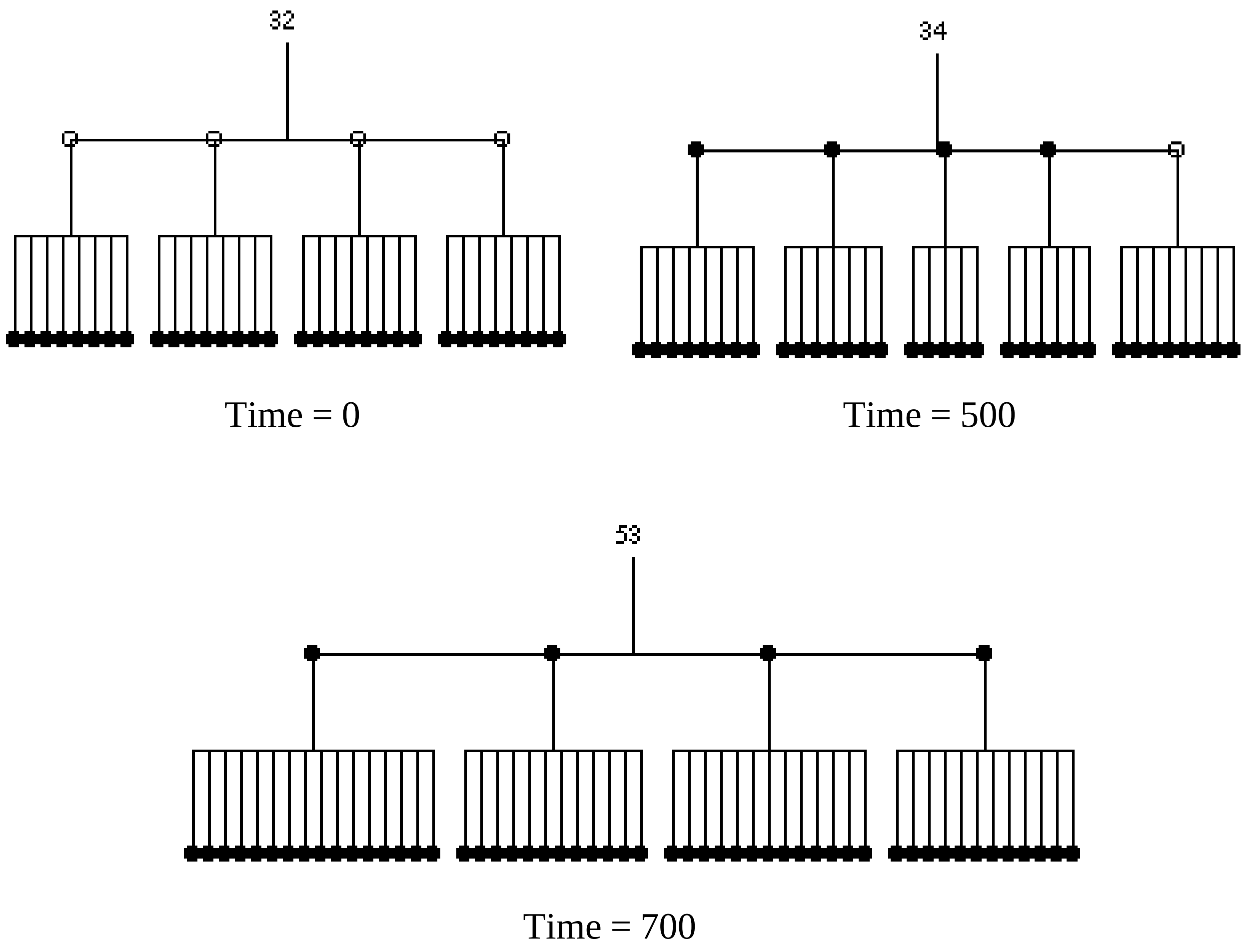

Figure 2: Snapshots of the time evolution of an industry starting from the same initial conditions and for the same choice of parameters as in the previous figure. The dynamical path followed in this case is very different. The number of organizations remains small for long enough that the organizations switch over to the equilibrium in which all organizations train. Once settled in the training equilibrium, the agents produce at higher and higher levels attracting more agents from outside the organization to join, further increasing the total utility produced by the industry as a whole.

3.2 Maximizing industry-wide productivity

Is there a relation between the utility produced by the industry as a whole and the average tenure lengths of its members? This question is very relevant in today's world of downsizing and rapid turnover. We ran a hundred simulations of the model using the same parameters and initial conditions, given in the appendix, to address this question. Fig. 4(a) indicates the correlation found between short tenure lengths and lower overall utility for the industry.

We also studied how sensitive the performance of the industries is to the values of various parameters in the model. We found two parameters were most significant, given the constraint that the model be kept in the regime of the two-level social dilemma. These were the entrepreneurial rate (the rate at which agents that break away start a new company) and the ratio of the learning rate to the training costs. When the entrepreneurial rate is high, the number of organizations proliferate rapidly and the likelihood that the organizations spontaneously decide to train drops. On the other hand, if the entrepreneurial rate is low, the number of organizations remains small, and the transition to overall training becomes much more likely. Low entrepreneurial rates also limit the overall size of the industry.

The effect of varying the learning rate is more interesting since companies or industries may have some

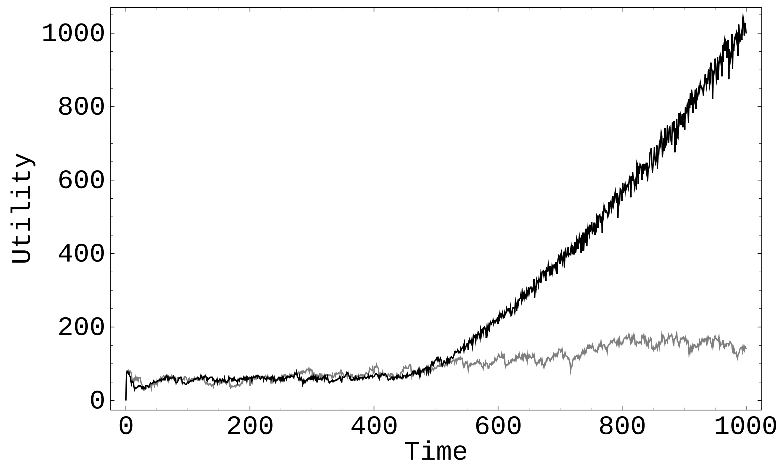

Figure 3: Utility as a function of time for the two industries described in Fig. 1 and 2, in gray and black respectively. The utility at time step 1000 for the industry of Fig. 1 is more than seven times greater than that of Fig. 2 (1000 compared with 140).

control over this variable through their policies on the level of training. In order to determine the average effect of increasing the learning rate while keeping training costs fixed, we ran the simulation many times for the same choice of parameters and initial conditions. Fig. 4(b) shows the average utility over 30 runs for each data point. The average utility increases exponentially with increasing learning rates. Increasing the learning rate by less than 50% results in a factor of six explosion in average utility for this set of simulations. The large increase in utility is the expected value; the actual change in utility for a given industry can vary widely because of the path-dependency described earlier. Such behavior has been observed in other organizational models with different assumptions [6].

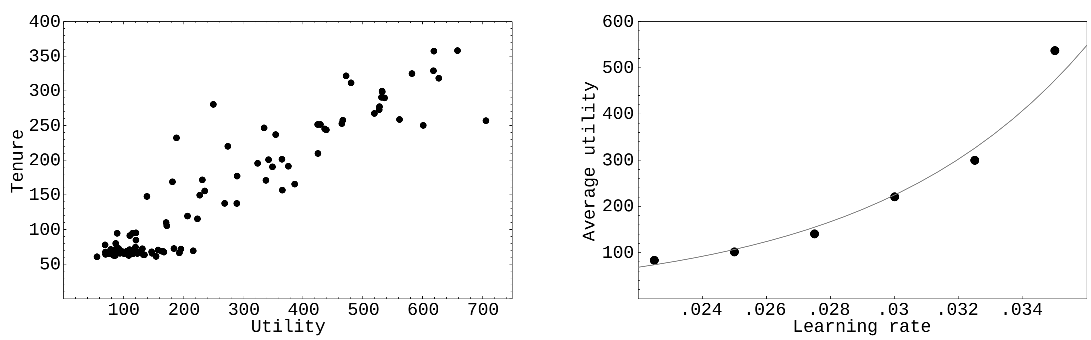

Figure 4: (*Left*) Scatter plot of average utility versus average agent tenure lengths for 100 simulations of an industry starting from the same initial conditions and identical parameter choices. (*Right*) The average utility produced by an industry over time typically increases exponentially with increasing agent learning rates. The data points were obtained by averaging over 30 runs for each value of the learning rate. The gray curve is an exponential fit to the data.

3.3 Changing environments and exogeneous shocks

Since the number of firms, their sizes, the extent of cooperation and training all change over time, we can say that the environment of the industry changes endogenously. We have seen that the qualitative aspects of these changes is case-dependent. However, the environment of the industry could also change exogenously.

We choose to model the changing environment as affecting the increased benefits of cooperation due to training. Alternatively, a changing environment could affect the baseline benefits and costs of cooperation or the costs of training. However, we are most interested in the effects of the environment on learning. For example, the introduction of a new technology may render past training more or less useful. If the introduction happens gradually, the industry can adapt to it smoothly. On the other hand, if the introduction is sudden, then the change may be very disruptive.

When the environment changes, agents that have been trained may partially lose their advantage over agents that have not been trained. Alternatively, the advantages of training may be heightened. Agents that have not been trained are assumed to be unaffected. If the external environment changes smoothly over time, the dilemmas on the organizational and employee level will gradually become either harder or easier to resolve for firms that train, depending on the direction of change. However, industries that train will still on average perform better than ones that don't. If the changing environment acts to lessen the benefits of training, then the likelihood that the firms will train decreases, but as long as the rate of environmental loss is not too high, industries that train will on average accrue higher utilities then ones that don't.

If the environment changes abruptly, the effect on trained agents can be sudden and large. For example, if the agents were trained to exploit one technology, they may not have the set of skills necessary to deploy a radically new one. We model this extreme case by imagining that an exogenous shock decimates the accumulated learning of trained agents. Before a shock trained agents are much more productive than untrained agents. After a shock trained and untrained agents produce at equal levels.

Consider as a concrete example that an exogenous shock occurs at time step 800 for the industry pictured in Fig. 2. By time step 800, the industry is made up of four large firms. All of the firms are training, and all of the 62 total employees are cooperating. An exogenous shock will render all of the employees' learning useless, taking their benefits of cooperation back down to the baseline benefit, i.e. the agents are now basically untrained. The sudden downward change in the benefit of cooperation makes it impossible for the firms to sustain employee cooperation because of their large size. In Fig. 5, we see in the first snapshot the sudden outbursts of defection already by time step 801. Four time steps later, most of the employees have fled the industry. By time step 820, the number of firms has decreased to two and then to one by time step 850. However, the one firm remaining still trains, and that firm is able to slowly recover as its agents adapt and learn. By time step 1000, the recovery is well underway. However, since, in this example, only one firm has survived, the rate of growth of the industry will not be as high as for the four-firm industry before the shock.

Fig. 6 shows the effect of the exogenous shock on the overall utility produced by the industry. At time step 800, when the shock occurs, there is a sudden and rapid decrease of the total utility. As the industry recovers, the utility once again starts to increase, but at a slower rate than previously. In other cases, the effect of the shock might be somewhat different since the dynamics of the industry is itself also highly path-dependent. In this case, the number of firms decreased after the catastrophic shock. In others, the number of firms might not decrease or might even increase. Whether or not the industry continues to train after the shock depends on what happens to the number of firms.

Note that in the example given, the average utility (over 1000 time steps) for the industry of Fig. 5 is still higher than that of the industry in Fig. 1, which never trains. We find that the average utility of an industry which is training before a shock is almost always greater than one which is isn't. Thus, the overall increased utility to industries that train generally makes up for the disastrous effect of exogenous shocks over short time scales. In addition, if we run many simulations with an exogenous time shock introduced, then we once again obtain a similar tenure-utility profile as the one shown in Fig. 4(a), except with the axes rescaled.

4 Discussion

To understand the interplay of social dilemmas at both the organizational and agent level, we constructed a simple model that encompasses cost-benefit analyses and expectations at both levels. At the organizational level, managers decide whether or not to train based on both the costs of training compared to the benefits and on their expectations and observations of the number of other firms that also train. Managers take into account the sum of their employees' contributions and the average tenure length within their organization. At the agent level, employees decide whether or not to contribute to company production based on their expectations as to how other employees will act. When trained, agents learn over time and fold their increased

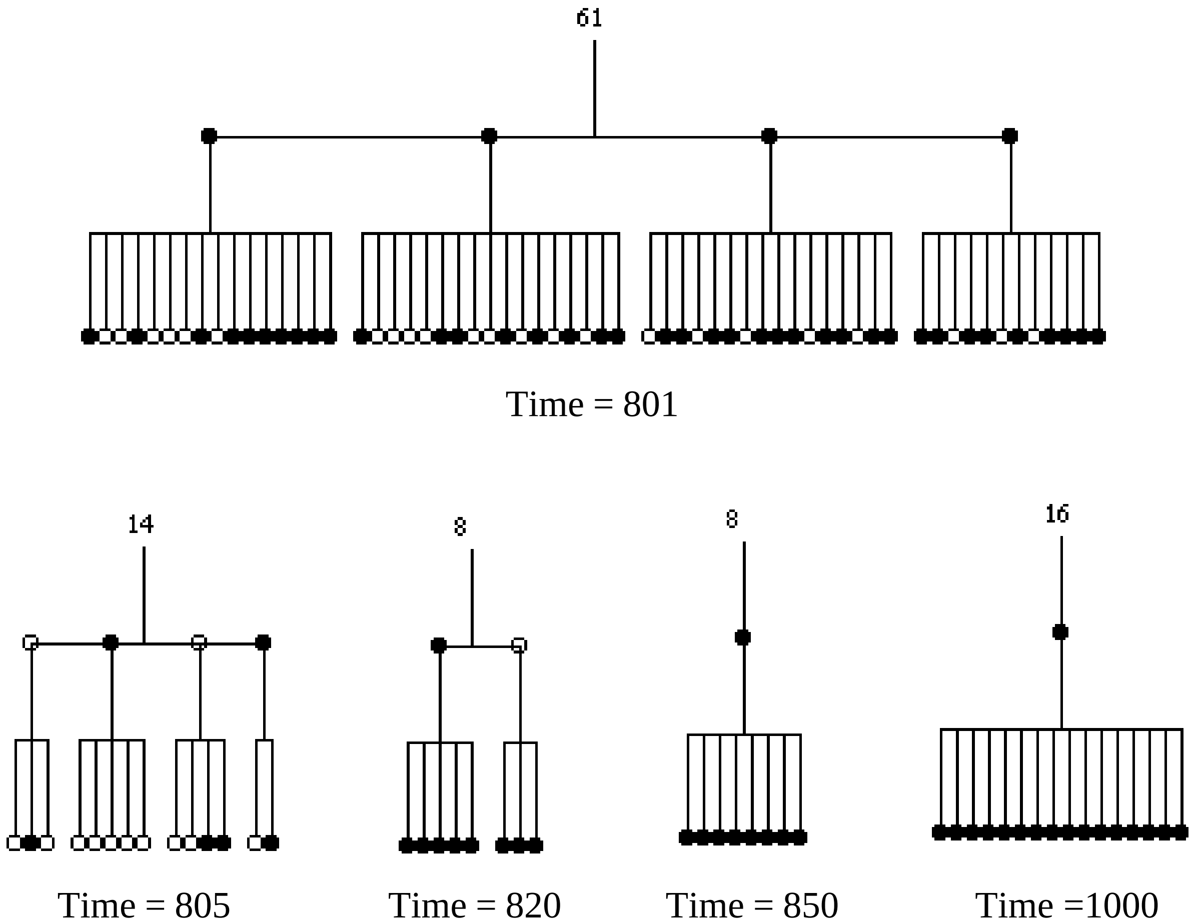

Figure 5: The industry of Fig. 2 undergoes an exogeneous shock at time step 800 which brings the benefit of cooperation of the trained employees back down to the baseline benefit of cooperation of untrained employees. One time step after the shock, many of the employees have switched over to defection and by time step 805 many have flown the industry. Over succeeding time steps the industry contracts further, until only one firm remains. Since its manager is still training, the firm slowly recovers from the exogenous shock and gradually grows over time once again.

productivity into their decision whether or not to contribute.

We also modeled how easily employees can move between firms, a property we call "structural fluidity." In addition, agents can leave the industry for good, and new ones can join. Our modelling turnover as a social dilemma differs from other approaches [6]. New firms may be created when an agent leaves its parent organization to start a new one. We described how fluidity relieves the dilemma at the agent level by allowing a large, low-productivity organization to break up into smaller pieces. In extreme cases, the organization may dissolve completely. However, when firms break apart in this way, the total number of organizations in the industry proliferates, exacerbating the dilemma on the organizational level.

The dynamical behavior at the two levels is closely coupled because of these interlinked effects. As a result, the dynamical unfolding of the dilemmas on the employee and organizational levels is path-dependent. The evolution of the industry over time depends not only on the characteristics of training programs, learning curves, and cost-benefit analyses, but also on the vagaries of chance. Starting from the same conditions, an industry can evolve to one of many states. In some cases, it evolves to a stable collection of firms that train their agents and become more productive over time. In other cases, the number of firms increases over time, and each firm experiences high worker turnover and low productivity because of the lower contributions of untrained and, at times, unmotivated, workers. These results are in line with the widespread empirical

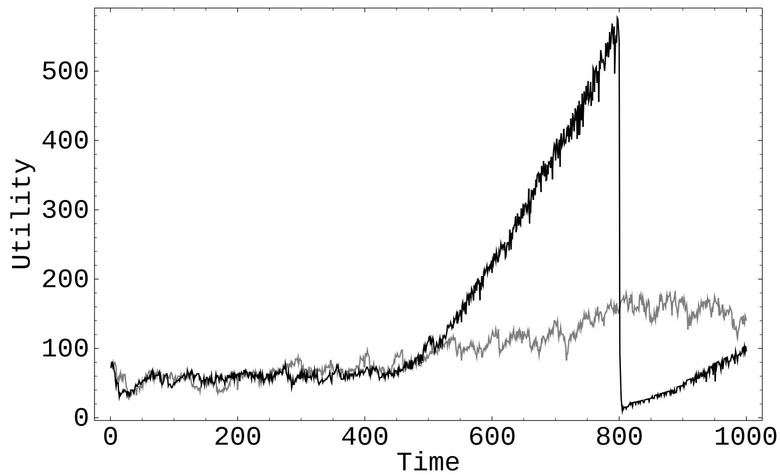

Figure 6: Utility as a function of time for the industry described in Fig. 5, which undergoes an exogenous shock at time step 800. The utility to the industry falls abruptly but starts to climb again once the industry recovers, albeit at a slower rate than before the shock.

observation that enterprise tenure is longer in larger firms and that the extent of training may differ between small and large firms [19]. Our computer experiments also show a correlation between high turnover and low overall utility to the industry, a correlation that has also been observed in several sociological studies that define performance as work-group productivity [21].

These results were obtained for both fixed and changing environments. In the more general case, the environment changes over time, perhaps setting the employees back in their training programs, or bankrupting firms. An environment that changes continuously may effectively offset some of the benefits of training, but the dynamics of the industry will be qualitatively similar. In such a case, organizations that train still have an advantage over those that do not. The effect of an environment that changes intermittently and abruptly is more dramatic. For industries that are training when the shock occurs, the change is catastrophic—employees stop contributing and flee the industry until the industry and its constituent firms become small enough to once again support cooperation. At this point new employees enter the firm, not necessarily those that previously left. However, we found that the effect of an exogenous shock is not disastrous enough to offset the gains of training to the industry over time. Even in an environment that changes abruptly, industries that train generally do better than ones that don't. Note that we did not include in our model of a exogenously changing environment any possible effects on employee and managerial expectations of the future. We expect that any such effects would probably be further destabilizing, perhaps in some manner decreasing agent and managerial horizon lengths.

In summary, our results indicate that organizational training can foster spontaneous cooperation in large firms, to some extent obviating the need for more complex management policies of employee-monitoring. Training can continue indefinitely if managers are able to constantly exploit improvements in technology, leading to a continuous rise in the organization's productivity. The ever-evolving nature of the dynamics of industries that we observe in our model contradicts the existence of a static economic equilibrium typically assumed when studying the economics of firms.

How well our results apply to human organizations will depend on the match between a particular industry and the characteristics of our model [4, 15]. Our assumption of common good problems on both the managerial level and the employee level will be an approximate description for a variety of industries and a poor description for others. Even for those industries which face the situation developed in the model, the dilemmas may remain dormant either because of the firms' small sizes or because of low costs or because of different managerial and employee decision-making criteria. However, in a variety of cases for which the free rider component of the problem of organizational training versus worker training is important, we believe that our results will provide insights into its dynamics.

In particular, our approach is useful because it addresses the *dynamics* of organizations and indicates how the interplay between organizational variables, such as the extent of training, the rate of turnover, enterprise size, and work-group productivity, is manifested as an industry evolves over time. Our method elucidates both the static relationships between organizational variables and the dynamics of path-dependent states, although several simplifying assumptions are necessary. If the model retains enough descriptive power to indicate cause-effect relations, this simplification is acceptable.

Our study suggests this kind of computer simulation can be used to design more efficient organizations. As we show, there is an advantage in being able to explore the unfolding of many possible scenarios to choose policies conducive to generating desired behavior.

A Computer Experiments

In this appendix we provide the details of the simulations that we used. It is characterized by a number of parameters that describe agent and organizational attributes. These parameters and their definitions for agents and organizations are given in the tables.

$b_{\min}$	Baseline benefit (per unit time) of cooperation
b_i^m	Benefit (per unit time) of cooperation for agent i belonging to organization m
c	Cost (per unit time) of cooperation
H	Horizon length
k_i	Binary variable: $k_i = 1$ if agent i contributes, 0 otherwise
γ	Learning rate
r	Fraction of learning transferred across organizations
t_i^m	Tenure length of agent i in organization m
α	Reevaluation rate
p	Measure of uncertainty

Table 2: Agent attributes.

N	Total number of organizations in the industry
n_m	Number of agents in organization m
κ_m	Binary variable: $\kappa_m = 1$ if organization m trains, 0 otherwise
T	Training cost per agent per unit time
H_m	Horizon length for manager m
α_m	Reevaluation rate for manager m
q	Measure of uncertainty
f_c^m	Estimated fraction cooperating in organization m

Table 3: Organization attributes

A.1 Algorithm

As described in the text, the simulation of our model uses two Poisson processes: one at the agent level with mean $1/\alpha$, and the other for the managers whose mean, n_m/α, depends on the size, n_m, of its organization. The conditions for the agents to move and join organizations are given in Eq. (9) and (10).

Our model has no prescribed limit to the number of agents in the industry: there is an infinite pool outside the industry which supplies the organizations and to which workers can return. The model also has no limit to the number of organizations in the industry. Each time an agent breaks away to form a new organization the total number of organizations increases. The number of organizations decreases whenever all of the agents in one particular organization break away from the industry completely to return to the external pool of agents.

The actual algorithm we used is as follows:

Initialize

- Structure of industry: number and size of organizations.
- Worker actions over all organizations: contribute or shirk?
- Managerial actions for each organizations: train or not train?
- Worker and organizational attributes.
- Wake-up (reevaluation) times Δt for all the workers and agents.
  - For manager of organization m, $\Delta t = -ln(\text{random number})/\alpha n_m$.
  - For a worker, $\Delta t = -ln(\text{random number})/\alpha n$, where $n = \sum n_m$.

While current_time < final_time:

- Wake-up earliest reevaluator and advance current_time.
- Move each worker being trained up the learning curve in proportion to its tenure length within its organization: $b_i^m = b_{\min} + t_i^m * \gamma\kappa_m$
- If earliest reevaluator is a worker, pick a worker at random. Worker reevaluates either (1) its decision to contribute or shirk; or (2) its position in the industry.
  - Worker reevaluates decision to contribute or shirk:
    - Evalute worker's observed share of production $\langle b \rangle^m$ from Eq. (5). Workers intending to contribute do so with probability p ($k_j = 1$ with probability p); workers intending to shirk also do so with probability p.
    - Evaluate critical threshold for cooperation b_{crit}^m from Eq. (6)
    - If $\langle b \rangle^m > b_{\text{crit}}^m$ worker contributes fully, otherwise worker shirks.
  - Worker reevaluates position in industry:
    - Evalute worker's observed share of production $\langle b \rangle^m$ from Eq. (5).
    - If $\langle b \rangle^m < \eta c$, ($\eta > 1$), then worker leaves organization.
      - If random number $< \Omega$, ($\Omega \ll 1$) then worker starts a new organization
      - Otherwise worker leaves industry entirely
- Otherwise earliest reevaluator is manager m, who either (1) reevaluates decision whether or not to train memebers of its organization; or (2) invites worker from another organization or from the outside pool of workers to join:
  - Manager decides whether or not to train:
    - Estimate fraction of its workers cooperating f_c^m from Eq. (8). Managers intending to train do so successfully with probability q ($\kappa_j = 1$ with probability q); managers intending to not train also do so with probability q.
    - Evaluate critical threshold for training f_{crit}^m from Eq. (7) using $H_m = \frac{1}{n_m} \sum t_i^m$
    - If $f_c^m > f_{\text{crit}}^m$, manager m trains, otherwise not
  - Manager invites outside worker to join:
    - Manager picks worker from other organization at random
      - Outside worker evaluates its current share of production in its current organization l, $\langle b \rangle^l$ from Eq. (5)
      - Outside worker compares $\langle b \rangle^l$ with the share $\langle b \rangle^m$ available to a worker in organization m.
      - If $\langle b \rangle^m > \langle b \rangle^l + \mu b_{\min}$, ($\mu < 1$), then worker accepts invitation to join organization m. Worker only retains part of the benefit of any training it received, specified in Eq. (3), with $r < 1$.
    - If worker declines invitation to join, then manager recruits from outside pool.
      - If $\langle b \rangle^m > \rho c$, ($\rho > 1$), then recruit joins organization m.
- Update wake-up time for worker or manager that just reevaluated its strategy.

A.2 Parameters

For the results we report, we used the following parameter values. For the agent attributes: $b_{\min} = 2.5$, $c = 1$, $H = 5$, $\gamma = 0.03$ (except in Fig. 4b where the learning rate was varied), $r = 0.9$, $\alpha = 1$ and $p = 0.95$.

For the fluidity parameters: $\mu = 0.1$, $\eta = 1.5$, $\Omega = 0.05$ and $\rho = 2$. The organizational attributes were $T = 0.02$ and $q = 0.95$.

Initially, we had four organizations ($N = 4$) each with eight agents ($n_m = 8$) all of whom were cooperating, but none of the managers were training.

References

[1] J. Bishop. Employer training and skill shortages: A review of the state of knowledge. Technical Report 91-32, Cornell Univ., Center for Advanced Human Resource Studies, 1991.

[2] Olivier Jean Blanchard and Stanley Fischer. *Lectures on Macroeconomics*. MIT Press, Cambridge, Massachusetts, 1989.

[3] A. S. Blinder and A. B. Krueger. International differences in labor turnover: A comparative study with emphasis on the U.S. and Japan. Council on Competitiveness and Harvard Business School project on U.S. corporate investment and the time horizons of American industry, 1991.

[4] R. Bronson and C. Jacobsen. Simulation and social theory. *SIMULATION*, 47(2):58–62, 1986.

[5] R. J. Bullock and Edward E. Lawler III. Gainsharing: A few questions and fewer answers. *Human Resource Management*, 23:23–40, 1984.

[6] Kathleen Carley. Organizational learning and personnel turnover. *Organization Science*, 3(1):20–46, 1992.

[7] Christine L. Cooper, Bruno Dyck, and Norman Frohlich. Improving the effectiveness of gainsharing: The role of fairness and participation. *Administrative Science Quarterly*, 37:471–490, 1992.

[8] Natalie S. Glance and Bernardo A. Huberman. The outbreak of cooperation. *Journal of Mathematical Sociology*, 17(4):281–302, 1993.

[9] Natalie S. Glance and Bernardo A. Huberman. Dynamics of social dilemmas. *Scientific American*, 270(3):76–81, March 1994.

[10] Natalie S. Glance and Bernardo A. Huberman. Social dilemmas and fluid organizations. In K. Carley and M. Prietula, editors, *Computational Organization Theory*, pages 217–239. Lawrence Elbraum, New Jersey, 1994.

[11] G. Hardin. The tragedy of the commons. *Science*, 162:1243, 1968.

[12] Russell Hardin. *Collective Action*. Johns Hopkins University Press, Baltimore, 1982.

[13] Marc Hequet. Can training stop turnover? *Training*, 30(10):82–87, October 1993.

[14] Bernardo A. Huberman and Natalie S. Glance. Evolutionary games and computer simulations. *Proceedings of the National Academy of Science USA*, 90:7716–7718, August 1993.

[15] C. Jacobsen, R. Bronson, and D. Vekstein. A strategy for testing the empirical adequacy of macro-sociological theories. *J. of Mathematical Sociology*, 1989.

[16] Jobs for the Future. *New Training Strategies for a High Performance Metalworking Industry*. Cambridge, MA, December 1991.

[17] Rosabeth Moss Kanter. Pay: From status to contributions: Some organizational implications. *Personnel*, 64:12–37, 1987.

[18] L. Lynch. The impact of private sector training on race and gender wage differentials and the career patterns of young workers. Final Report submitted to the U.S. Department of Labor, Bureau of Labor Statistics, 1991.

[19] OECD. *Employment Outlook*. Paris, July 1993.

[20] Paul Osterman. Turnover, employment security, and the performance of the firm. In Morris M. Kleiner and Steven G. Allen, editors, *Human Resources and the Performance of the Firm*, pages 275–318. Industrial Relations Research Association, Madison, WI, 1987.

[21] Joseph Price. *The Study of Turnover*. Iowa State University Press, Ames, 1977.

[22] Thomas C. Schelling. *Micromotives and Macrobehavior*. W. W. Norton and Company, Inc., United States, 1978.

[23] Joseph Schumpeter. *The Theory of Economic Development*. Oxford University Press, New York, 1961.