M. Ruth & B. Davidsdottir, eds. *Changing Stocks, Flows, and Behaviors in Industrial Ecosystems* (Edward Elgar 2008).

7 Changing a Firm's Environmental Performance from Within

Clinton J. Andrews

INTRODUCTION

Industrial ecologists are becoming interested in better understanding the influences on industrial structure and performance. It is clear that agency exists at several levels, with the fundamental sources of agency in industrial ecosystems being individuals acting as citizens, employees, investors, and consumers. Nations are the key actors on the global stage, but national policies emerge in part from interactions among citizens and organizations. Firms are key actors within sectors, supply chains, and symbioses. Corporate behavior, to some extent, emerges from interactions among employees, and market outcomes emerge from the myriad choices of individual consumers.

Individuals are not truly independent actors, however (Scott 2001). Institutions and organizations place formal, regulative constraints on individual choices: the chain of command, the order of work, and legal requirements. They also impose informal, normative constraints, such as a code of professional conduct, a work ethic, or an expectation of environmental stewardship. Some constraints are entirely unwritten and operate as cultural framing assumptions or cognitive biases: for example, humans should be fruitful and multiply, and humans should satisfy their immediate survival needs before addressing abstract concerns that are distant in time, space, or genetic similarity.

The conundrum of whether agency determines structure, or vice versa, disappears if one recognizes that they interact (Giddens 1984). Individuals have some leeway to change the organizations where they work and also their governing institutions through quitting their jobs, voting politicians out of office, and buying different products. Yet how much influence do agents have on structures at the time scales of interest to industrial ecologists? This chapter uses a simulation modeling framework to explore agency relationships in a particular corporate environmental management context.

COMPUTATIONAL STUDIES OF ORGANIZATIONS

Organizational behavior is a multidisciplinary field that describes, explains, and prescribes how agency and structure interrelate in firms. The origins of the field lie in the sociology of Weber's bureaucracies, the economics of Smith's specialized pin makers, and the engineering of Taylor's scientific managers. Modern studies of organizations draw on systems thinking and social network theory for inspiration.

Following World War II, theorists characterized organizations as controllable systems (Ashby 1956) with map-able structures and information feedbacks (Forrester 1961). Development of "open systems" theory represented a major breakthrough that emerged in stages. Economists noticed that organizational structures could be viewed as responses to varying external environments (Lawrence and Lorsch 1967). Social psychologists claimed that organizations were open systems that interfaced in multiple ways with the external environment (Katz and Kahn 1978). Systems dynamicists, in a nod to game theory, observed the explanatory power of "double interacts," involving linked decision makers (Weick 1979). Philosophers identified organizations as interacting assemblages of individuals, objects, and exogenous forces (Churchman 1979). Those attentive to the second law of thermodynamics noticed how "organizations are dissipative structures that can only be maintained when members are induced to contribute energy to them" (Anderson 1999; Barnard 1938; Prigogine and Stengers 1984).

Organizations are structured by their networks of internal and external relationships. Formal relationships between principals and agents, that is, owners and employees, only elicit desired behaviors if appropriate incentives are in place (Kerr 1975; Panayotou and Zinnes 1994; Gibbons 1998). More broadly, the economic characteristics of contracting internal and external relationships help researchers understand why there are firms and what their boundaries should be (Coase 1937; Williamson 1979; Holmström and Roberts 1998). Friendship networks and other informal group processes are important as well (Locke and Schweiger 1979).

There are several ways to study organizational behavior and none is fully satisfactory. Aggregate statistical analysis of firms' performance often yields trivial results and may fail to address the multi-level nature of the phenomenon. Case studies of individual firms can capture rich detail but in a static, retrospective snapshot. Experimental studies of social psychology in the workplace must abstract drastically from complex reality in their pursuit of controlled conditions. Longitudinal studies that track employees and firms over time are revealing but time consuming and expensive. Studies focusing on individual behavior in the context of external pressures are helpful for studying leadership but do not really advance understanding of the links between agency and structure.

Multi-agent simulation (MAS) modeling is a complementary way to conduct formal theorizing about organizational behavior, to explore interactions among the moving parts in a system, and to inform managerial decision making in a prospective, progressive manner. Although it too suffers from weaknesses, in a relatively few years it has become a useful and widely used method for organizational research. It complements causal modeling (Snijders 1998), case studies (Lin 2000), experiments (Burton 2003), and training activities (Fridsma and Thomsen 1998). Unlike its antecedents in game theory and systems dynamics, MAS supports speculation about links between agency and structure in organizations (Levitt 2004).

Most MAS researchers with training in economics do bottom-up modeling, meaning that they specify agent rules for interacting with one another and the external environment, then start the simulation and observe what happens. They hope to see social structure emerge from the interactions of individuals. They build in autonomy, decision making, and the possibility of adaptation and agent evolution. Researchers with training in sociology are more willing to specify structure in their models, acknowledging that in the real world, firms already exist, and employees have no choice but to work within existing structures. Some researchers take a Lamarckian view and allow firms to evolve in response to their employees and their external environment. See Conte, et al. (2001) for an extended discussion.

In the 1990s, MAS was limited to modeling highly stylized versions of phenomena such as cognition (Cooper, et al. 1996), interactions among members of generic teams (Carley and Prietula 1998), and social network dynamics (Zeggelink, et al. 1996). Recently, the field has been able to tackle more applied topics such as absenteeism (Sanders and Hoekstra 1998), employee turnover (Harrison and Carroll 2002), employee promotion (Phelan and Lin 2001), affinity and animosity (Costa and de Matos 2002), and the imposition of quality standards (Torenvlied and Velner 1998). The current paper presents an applied MAS model of organizational behavior affecting environmental management, tests its sensitivity to key assumptions, and shows the effects of changing policy variables.

A STYLIZED FIRM

This paper looks at selected internal dynamics within a stylized firm. It formalizes certain abstractions of organization theory in a multi-agent simulation model of a firm that includes a factory, its employees, and the external environment. The factory provides structure by specifying the physical characteristics of the technology and the regulative constraints of the production process. The employees have structured, hierarchal relationships among the plant manager, shift supervisors, and machine operators. But employees also have informal social network ties that influence their behavior. The external environment provides inputs and exogenous, driving factors to which the employees (and thus, the firm) must respond, and it also accepts outputs, including products and pollution, from the firm.

The time tick of the simulation model is hourly in order to capture the short time scale at which human behavior is the dominant dynamic. The model can run, however, for tens of thousands of time ticks, entering the multi-year time scale over which technological change takes place. The model thus allows exploration of the relationship between human resource management and technology management within firms.

To provide an empirical anchor for the research, this modeling effort focuses on a specific industry and context, chosen for its relative simplicity. It models a small polymer processing firm that manufactures injection-molded products such as cafeteria trays for fast food restaurants, molded housings for consumer electronic products, and plastic coffee mugs. Prior to developing the model, the research team visited and conducted detailed case studies of several actual firms in New Jersey, US, and Suzhou, China, summarized in Andrews (2006).

The production technology for such a firm is fairly simple, involving a set of injection molding machines, a costly stainless steel mold for each product type, and a generic factory building with a warehouse for raw materials, a factory floor for the many parallel production lines, and a shipping area. The production process is also simple: plastic pellets are heated, the melted plastic is injected into a mold, once cooled, the product is released from the mold, the product is cleaned and finished, and it is shipped. The human resource relationships in such a firm are also not too complex: a plant manager hires, fires, and makes capital investments; a marketing manager brings in new business; shift supervisors assist machine operators and also report on their performance to the plant manager; machine operators operate the injection molding equipment; a materials mixer delivers plastic pellets to the production lines; a shipping clerk sends completed products out the factory door; and a janitor regularly cleans the factory floor.

DESCRIPTION OF MODEL

At this idealized firm, maximum capacity is fixed at ten production lines. This avoids the complexity of expansion planning and allows a focus on managing existing assets. Inputs to the production process include labor, electricity, and plastic pellets. Outputs include products (here called “widgets”), scrap plastic, and air pollution. Scrap plastic can be recycled, and air pollution is an increasing function of process temperature. Widget quality is a non-monotonic function of process temperature, that is, there is a quality-maximizing setpoint, so there is a tradeoff between pollution reduction and product quality objectives. Machine operators set the process temperature at what they view as the optimal point.

It is possible to automate parts of the production process, thereby eliminating much labor, reducing human error, and improving product quality. The plant manager will invest in automation when it appears to be economically prudent given the relative costs of inputs, especially labor, and the prices obtained for the widgets that are the factory's outputs. Automation is the sole technological choice available to the plant manager.

The plant manager makes several types of human resource decisions. Based on shift supervisors' reports of machine operators' attendances and absences, and the frequency with which they need assistance in operating their molding machines, the plant manager can fire underperforming machine operators. The plant manager can also hire replacement machine operators from a pre-established pool using criteria including aptitude, attitude, experience, and cultural background. Finally, the plant manager can promote a well-performing machine operator to become a shift supervisor when an opening arises.

Individual employees also have the autonomy to make several decisions. They choose whether to come to work each morning and whether to alter their attitude toward work when their supervisors admonish them for misconduct. Machine operators each choose what they consider to be the optimal temperature set point for their injection molding machine, weighing the firm's desire for high quality widgets against their own possible interest in reducing air pollution.

Individuals are quite heterogeneous in terms of innate aptitude and preferences, socially-influenced attitudes, and acquired experience. These factors influence their decisions and the plant manager's responses, leading in some cases to changes in firm-wide performance.

IMPLEMENTING THE MODEL

This model was implemented in Java using the Ascape multi-agent simulation framework. Object-oriented programming languages such as Java and C++ provide a straightforward pathway to agent-based modeling, because agents can be implemented as software objects that interact with one another. Ascape was developed at Brookings Institution, serving as one of the first libraries of Java routines to facilitate research on social science topics (Parker 2001). Well known competing frameworks include Swarm and RePast (Swarm 2006, SourceForge 2006).

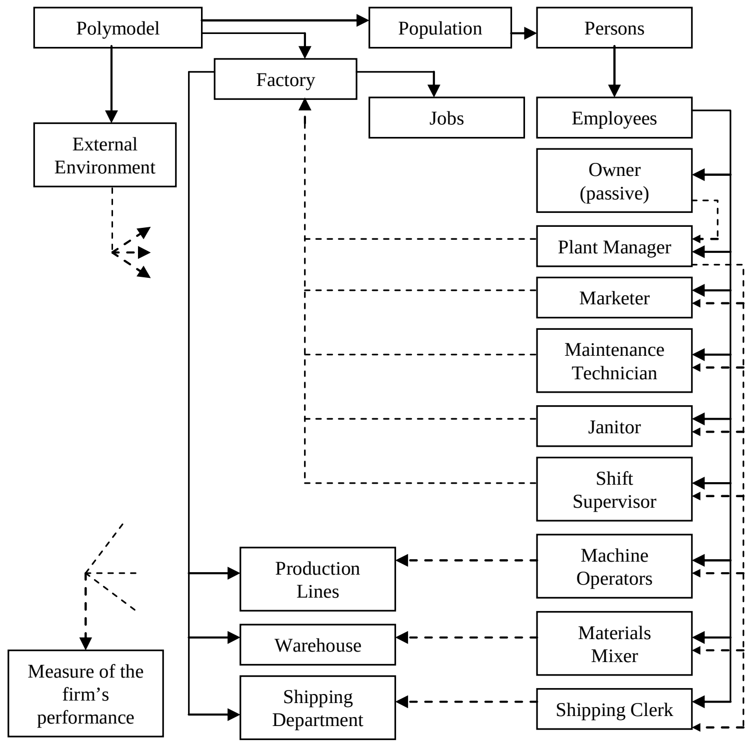

Figure 7.1 Structure of PolyModel (Note that solid lines represent the modeling hierarchy, specifying which objects inherit traits of other objects, whereas dashed lines represent key modeled interactions among objects)

The model includes 22 classes representing the categories of objects that interact (see Figure 7.1). These include PolyModel (the overall framework for the model), the external environment, the factory (containing production lines), the people (some of whom acquire jobs), and miscellaneous supporting objects. Both the raw code and executable versions of the model are available online (Andrews 2006).

Key elements of the model's logic are the bounded rationality and biases of the employees, the way employees relate in social networks, the factors that drive employees' decisions, the hiring and firing processes used by the plant manager, the production process, and the plant manager's automation decision process.

Employee Decisions

It is common when modeling principal-agent problems to assert that agents know their utility functions and act to optimize them. A familiar formulation posits a labor–leisure tradeoff (Block and Heineke 1975). An alternative view is that employees have bounded rationality and do not fully understand their own utility functions or the firm's cost function (Simon 1947). Adopting this latter view brings behavioral realism that approximates utility maximization under some circumstances but diverges in interesting ways under others, as employees employ heuristics and decide without complete knowledge (Cattaneo and Robinson 2000). The following are key heuristics in PolyModel:

- An employee decides to go to work if he feels healthy enough (random draw), acts responsibly enough (random draw), and a large enough fraction of his friends are also going to work (based on a poll of his social network).
- An employee decides to quit his job if the net present value of quitting exceeds that of staying, given the difference between his wage and the prevailing wage, the expected time to re-employment based on the current unemployment rate relative to the historical rate, and his discount rate.
- A machine operator decides what temperature to set his injection molding machine at based on his environmental views. If he is an environmentalist, he will attempt to set the temperature lower to minimize air pollution. If not, he will attempt to set the temperature higher to minimize product defects. The degree to which he hits these target temperatures is a function of worker error.

Employee Characteristics

The standard rational actor model, *homo economicus*, adopts a Hobbesian view of individuals as atomistic, pre-social agents. Yet a commitment to methodological individualism does not imply adherence to such a naïve model, and there is every reason to include a richer social dimension (Heath 2001). As Elster (1989:13) observes, "the elementary unit of social life is the individual human action." Likewise, whereas the simple agent models typical of game theory typically employ the unrealistic assumption that agents are homogenous, there is much benefit to relaxing that assumption and allowing heterogeneity (Lempert 2002). It is also useful to characterize agents as "computational" in the sense that they have bounded information processing abilities (Carley 2002) that give rise to mistakes. Thus PolyModel characterizes employees and their firm further as follows (see also Figure 7.2):

- Worker error affects most stages of the production process. Many production parameters are set by drawing a number from a random distribution, and higher worker error increases the dispersion around the mean, "target" value of the production parameter. Worker error is a weighted, decreasing function of the employee's aptitude, experience, and happiness, and an increasing function of tiredness.
- Employee aptitude is assigned randomly at the beginning of the simulation; their work experience grows over time from an initial random endowment; and tiredness increases with number of consecutive hours worked without a break. Happiness increases with income and friendships on a normalized, weighted basis, and is adjusted slightly for cleanliness of the workplace. Happiness weights are assigned exogenously so that the model user can adjust the fraction of "Type A" persons (0.9 money, 0.1 friends) and "Type B" persons (0.1 money, 0.9 friends) in the population.

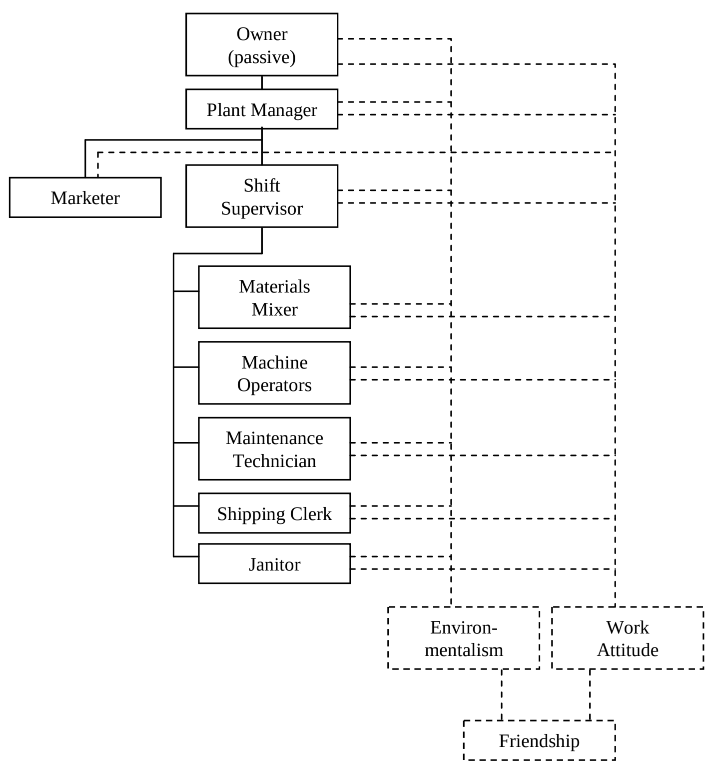

Figure 7.2 Formal reporting hierarchy versus social network

- The normalized number of friends, which is calculated every Saturday night, itself involves weighting and screening calculations, such that acquaintances with more similar cultural backgrounds (e.g. "environmentalist," "non-environmentalist") count more heavily than those with different backgrounds, people with different attitudes (e.g. "good," "bad") cannot become friends, and people become better friends with people they have previously and repeatedly met at work.
- Employees can change both their attitude and culture. If an employee with a "bad" attitude is threatened with firing by the plant manager, he chooses (with some probability) to adopt a "good" attitude. If a majority of other employees are "environmentalists," then an employee chooses (with some probability) also to be one, and vice versa.

Hiring and firing

Each day, the plant manager evaluates all of the employees. He asks the shift supervisor (who transmits information with some probability of error) about the employee's technical performance (based on amount of worker error), absenteeism, and number of times the shift supervisor has had to help the employee correct their mistakes. He fires some fraction of the employees whose performance falls below acceptable thresholds on these three dimensions.

When a job opening arises because an employee quit or was fired, the plant manager hires someone. He draws a random sample of potential employees from the population, ranks the potential employees according to an exogenously set hiring bias (favoring some weighted combination of aptitude, attitude, experience, and culture), and hires the highest ranked person.

Automation

Every year the plant manager decides if any production line should be automated. If there are enough funds and a line is not automated, then he estimates whether the cost of automation is lower than the annual cost of running the existing line. The running cost of a line consists of the machine operator's salary, the cost of supplying plastic pellets and electricity, the cost of disposing waste, and the cost of recycling. Automated lines do not suffer from worker error.

MODEL CALIBRATION AND VALIDATION

Model calibration and validation was a multi-step process. The first step was to derive initial conditions of the models from the case study evidence, secondly, drive the models using exogenous conditions identified in the case studies, and finally, validate the models against outcomes from the case studies. The case studies are available online (Andrews 2006).

Criteria for determining the acceptability of the multi-agent simulation model included ease of implementation and use, validity, and reliability. Ease of use was tested by asking several groups of non-programmer graduate students to conduct a classroom exercise using the model. By following good programming practices, reliability was ensured.

Model validation proved to be the least straightforward task. Since this was a highly stylized model of an extremely complex reality, it was unreasonable to expect to replicate in detail the experiences documented in the case studies. Instead, "we were mostly concerned with validating the insights we gained . . . from the simulation" (Shannon 1975: 29). The goal was to achieve enough face validity to approximate the major features of the case studies and enough robustness to perform plausible "what-if?" simulations that help users understand both the model's relative sensitivity to different assumptions and how the model behaves after altering policy variables. This chapter typically reports relative results that compare a simulation to a base case that approximates the case study experience.

RESULTS

This section shares illustrative modeling results that provide insight into the effective sources of agency in firms. The first set of results shows the base case, in which the model responds to historical data. The second set of results shows four policy scenarios applied to a steady state version of the base case in which the exogenous drivers are fixed.

Base Case

The base case mimics performance of the industry as observed in the case studies, and includes parameter values calibrated to be "reasonable." The key assumptions in the base case are that (1) there is no hiring or firing bias, (2) the plant manager fires only one-fifth of those employees whose performance is inadequate in any given month, (3) when a shift supervisor helps a machine operator this is the equivalent of increasing the operator's experience by 5 per cent, (4) one-fifth of the population call themselves environmentalists, and (5) one-third of the population has utility functions weighted towards preferring friendships over money.

Time trends for the model's exogenous drivers, for example, current costs of inputs and prices of products shipped, are shown in Figure 7.3 for the years

2001 to 2005. Prevailing wages of labor have increased steadily since the 1950s, electricity costs increased rapidly during the 1970s before recently leveling off, plastic pellets have risen in cost sporadically and have shown remarkable acceleration in recent years, and the price obtained for the product outputs ("widgets") has increased only modestly over time, with a recent significant flattening in prices. Since 2001, the competitive position of the US polymer processing industry has deteriorated dramatically. New low-cost manufacturers have sprung up in Asia, and a significant amount of US manufacturing capacity has shut down or moved off shore. More than ever, US manufacturers are price takers, and the average price per widget has dropped in real terms. At the same time, the cost of inputs has increased. Labor costs have slowly but steadily increased over decades, whereas plastic costs have skyrocketed. This squeeze has threatened the viability of the US injection molding industry.

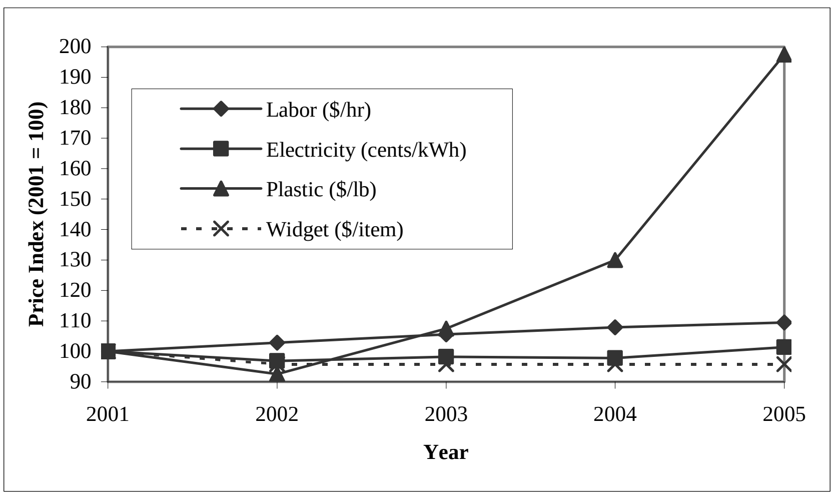

Figure 7.3 Cost and price trends in the polymer processing industry

Figure 7.4 shows how the model responds to the exogenous drivers. Profits drop dramatically over time, with the firm on the verge of losing money. Most other variables (not shown) achieve steady performance, although they exhibit substantial amounts of random noise caused by worker error, absenteeism, equipment malfunctions, and similar factors. These results are plausible and indicate that the model has face validity.

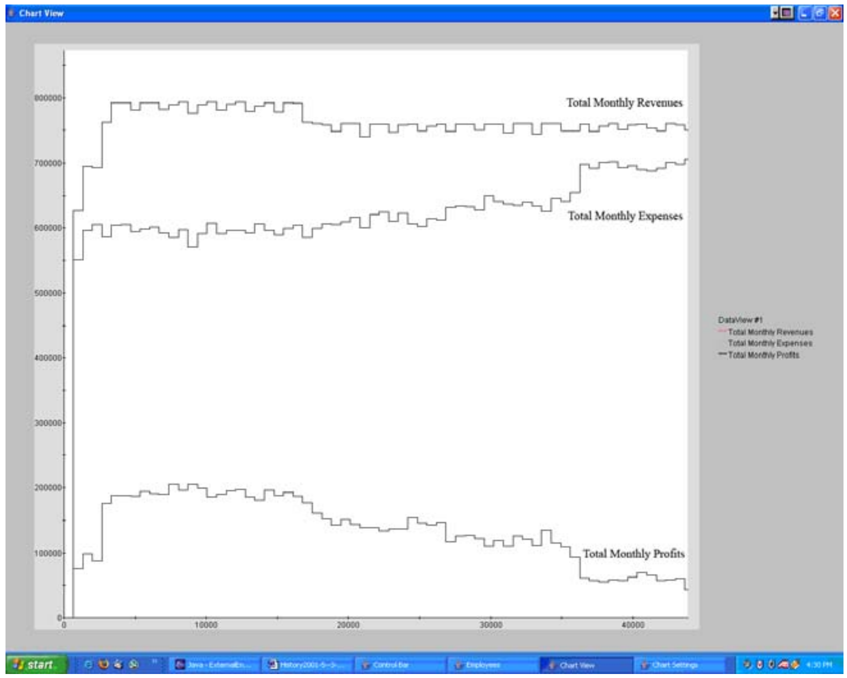

Figure 7.4 Historical simulation of revenues, expenses, and profits

Steady State Policy Scenarios

Extending from the base case by (1) freezing the values of the exogenous drivers at 2005 levels, and (2) changing selected parameters of the model, this section compares four possible ways to reduce the environmental impacts of economic activity. Specifically, these four scenarios portray different approaches to reducing the amount of air pollution produced by the plastic injection-molding firm. Table 7.1 shows average per cent differences between the base case and each policy scenario. Results are based on 20 repeated simulations of 10 000 hours for each scenario, and they are evaluated for significance using a paired, two-tailed difference-of-means test. The values of the table entries should be treated as indicative rather than precise measures of differences among scenarios.

Simulation #1: promote environmentalism as a social movement

A bottom-up approach to changing the environmental performance of the firm is to promote environmentalism in the population. If the proportion of environmentalists increases, then they will eventually infiltrate the firm through normal hiring practices. Bansal (2003) argues that individual concerns can translate into an organizational response provided that employees enjoy adequate discretion and the organization enjoys excess resource slack. Thus the simulation explores what happens when those hired as machine operators have discretion to adjust the temperature set points of their injection molding machines to reduce air pollution.

Table 7.1 Summary of results (per cent change from base case after reaching steady state)

Scenario Metric	1. Promote Environmentalism	2. Hire Environmentalists	3. Impose quality control	4.Automate Production
Air Pollution (monthly emissions total)	-8***	-7***	0	0
Profits (monthly net)	-2***	-1**	+4***	+14***
Employee Happiness (average per employee)	-3	-3	-2	+16
Employee Experience (average per employee)	+2	0	+9***	+22***
Environmentalists Present (among employees)	+240***	+217***	-4	-63***
Worker Error (average per employee)	-1	-1	-26***	-18***
Pollution/Widget (average per widget shipped)	-8***	-5***	-2***	-4***

* Difference of means test is significant at $p = 0.05$
** Difference of means test is significant at $p = 0.01$
*** Difference of means test is significant at $p = 0.001$

In this simulation, the probability of encountering an environmentalist in the population is increased exogenously from 20 per cent to 100 per cent. Through normal firing and hiring practices, environmentalists soon infiltrate the firm and operate the injection molding machines in a way that reduces overall air pollution and pollution per widget produced by 8 per cent. Profits decrease an average of 2 per cent. The spread of environmentalism has insignificant effects on employee happiness, experience, and worker error rates.

Simulation #2: selectively hire environmentalists

Champions can powerfully influence organizational environmental performance (Andersson and Bateman 2000), especially if the champions have senior management positions (Howard-Grenville 2005). In this scenario, it is the plant manager who joins the environmental movement. Instead of waiting for society as a whole to become more environmentally conscious, she preferentially hires them, accelerating the infiltration of environmentalists into the firm. This top-down initiative ensures that more machine operators will optimize their injection molding machines to reduce air pollution.

The environmentalist hiring bias is implemented in the model by exogenously increasing the weight given to the cultural attribute (environmentalism) relative to other attributes such as aptitude, attitude, and experience. When hiring, the plant manager screens job applicants on that basis. Some employees are susceptible to peer pressure and will change their environmental views if they are outnumbered. Peer pressure thus amplifies the preferential hiring strategy once it has advanced far enough.

This hiring strategy succeeds in reducing total air pollution by 7 per cent and pollution per unit of production by 5 per cent. The strategy fails, however, to capture the indirect benefits of cultural homogeneity (increased happiness, reduced error, increased profits). It also hinders profitability (by 1 per cent) by displacing high aptitude new hires with environmentalists who may or may not have other good qualities.

Simulation #3: impose a strict quality control regime

Eco-efficiency arguments from industrial ecology assert that much pollution is unintended, and that tighter quality control and reduced worker error rates can simultaneously reduce pollution and increase profits (Manahan 1999). Similar assertions appear in the quality management literature (Evans and Lindsay 2005). This simulation explores a top-down initiative that focuses solely on reducing the worker error rate.

This strategy is implemented in the model by exogenously setting (1) the employee hiring bias to favor candidates with high aptitude, and (2) increasing the amount of helpful supervision that shift supervisors provide to machine operators.

The environmental results are visible only at the margin: total air pollution does not change at all relative to the base case, and pollution per unit of production drops by 2 per cent. On the other hand, profits increase by 4 per cent. Employee experience also increases because better supervision leads to fewer firings. These factors together lead to a large 26 per cent drop in the worker error rate.

Simulation #4: automate the production process

Rather than rely on human resource management tools, the plant manager could instead choose a technical fix calculated to promote eco-efficiency. Automating portions of the production process would displace error-prone and expensive employees, thereby reducing pollution and increasing profits, assert systems analysts such as Reimann and Sarkis (1996).

The model has a switch that allows the plant manager to choose automation when it appears to be cost-effective. Once this feature is exogenously switched on, the plant manager evaluates current input costs (labor, electricity, plastic pellets), product prices, and capital costs, proceeding to automate a production line when the expected net present value of the new technology investment exceeds the status quo option.

The results show that total pollution does not drop but pollution per widget produced does decrease by 4 per cent. Monthly profits shoot up by 14 per cent and worker error drops by 18 per cent relative to the base case.

DISCUSSION

The four scenarios described above represent pure management strategies, whereas real managers would combine elements in their pursuit of eco-efficiency, profits, and environmental improvement. Nevertheless, the results reveal interesting subtleties and tradeoffs that hint at the complexity of the managerial challenge.

Long ago, Miles and Rosenberg (1982: 26) noted that "organizations that fail to redesign internal roles and relationships to tap the underutilized capabilities of human resources not only experience built-in inefficiencies, but also strain our social fabric." Indeed, the results shown here suggest that human resource tools of hiring, training, supervision, and firing are most effective when coordinated to achieve clear objectives. The underlying labor pool bounds the firm's potential. A rogue manager who preferentially hires environmentalists (or cronies, or people from a particular ethnic group) can affect the overall performance of the firm.

The most profitable strategies emphasize quality control or automation, and they provide minor environmental benefits given the characteristics of this injection molding technology. For other technologies, quality management may pay higher environmental dividends.

It is noteworthy that the behavioral/human resource fixes provided benefits that were of similar magnitude to the technical fix (automation). This result supports previous claims that organizational practices deserve attention from industrial ecologists (e.g. Andrews and Swain 2000).

CONCLUSIONS

Locating the sources of agency in the industrial ecosystem is a high priority task, because environmental improvements depend wholly on their actions. This paper has examined the relationships among actors within the firm, and it has found that social movements, champions, organizational rules and procedures, and technological change can each affect the firm's overall performance. More generally, as Carley (2002: 7262) observes:

> In many cases, simply building the model brings value in and of itself to the policy maker or manager as it lays bare hidden assumptions and potential limitations in the current system. As we move toward more realistic agents and tasks, the value of these models can only grow.

Such models are just beginning to be taxonomized and standardized (Chang and Harrington 2005) as this research community expands. The multi-agent simulation model described here has demonstrated its value as a tool for formal theorizing about complex, hard to study organizational dynamics. Given its relatively user friendly interface, PolyModel can also be used as a management training tool. In the latter role, models like it may help future managers become more effective sources of agency within the industrial ecosystem.

ACKNOWLEDGEMENTS

The author gratefully thanks Preetham Mysore and Shawn Patton for their programming efforts, Ana Baptista for performing case study research, Robert Axtell for his collegiality, and the Science To Achieve Results program of the US Environmental Protection Agency for financial support.

REFERENCES

Anderson, P. (1999), "Complexity theory and organization science," *Organization Science*, **10**(3): 216-232.

Andersson, L.M. and T.S. Bateman (2000), "Individual environmental initiative: Championing natural environmental issues in US business organizations," *Academy of Management Journal*, **43**(4): 548-570.

Andrews, C.J. (2006), *PolyModel Home Page*, http://policy.rutgers.edu/andrews/projects/abm. Last accessed 2 September 2007.

Andrews, C.J. and M. Swain (2000), "Institutional factors affecting life-cycle impacts of microcomputers," *Resources, Conservation and Recycling*, **31**: 171-188.

Ashby, R. (1956), *An Introduction to Cybernetics*, London, UK: Chapman and Hall.

Bansal, P. (2003), "From issues to actions: The importance of individual concerns and organizational values in responding to natural environmental issues," *Organization Science*, **14**(5): 510-527.

Barnard, C.I. (1938), *The Functions of the Executive*, Cambridge, MA, USA: Harvard University Press.

Block, M.K. and J.M. Heineke (1975), "Factor allocations under uncertainty: An extension," *Southern Economic Journal*, **41**(3): 526-530.

Burton, R.M. (2003), "Computational laboratories for organization science: Questions, validity, and docking," *Computational & Mathematical Organization Theory*, **9**: 91-108.

Carley, K.M. (2002), "Computational organization science: A new frontier," *Proceedings of the National Academy of Sciences*, **99**(3): 7257-7262.

Carley, K. and M. Prietula (eds) (1998), *Simulating Organizations: Computational Models of Institutions and Groups*, Cambridge, MA, USA MIT Press.

Cattaneo, A. and S. Robinson (2000), "Empirical models, rules, and optimization: Turning positive economics on its head," *TMD Discussion Paper No. 57*, Washington, DC: International Food Policy Research Institute. Last accessed 2 September 2007 http://www.ifpri.org/divs/tmd/dp/papers/tmdp53.

Chang, M.H. and J.E. Harrington, Jr. (2005), "Agent-Based Models of Organizations," *Economics Working Paper 515*, Johns Hopkins University, http://www.econ.jhu.edu/papers.html. Last accessed 2 September 2007.

Churchman, C. (1979), *The Systems Approach and Its Enemies*, New York, NY: Basic Books.

Coase, R. (1937) "The nature of the firm," *Economica*, **4**: 386-405.

Conte, R., B. Edmonds, S. Moss, and R.K. Sawyer (2001), "Sociology and social theory in agent-based social simulation: A symposium," *Computational & Mathematical Organization Theory*, 7: 183–205.

Cooper, R., J. Fox, J. Farringdon, and T. Shallice (1996), "A systematic methodology for cognitive modeling," *Artificial Intelligence*, **85**: 3-44.

Costa, L.A. and J.A. de Matos (2002) "Toward an organizational model of attitude change," *Computational & Mathematical Organization Theory*, **8**: 315-335.

Elster, J. (1989), *Nuts and Bolts for the Social Sciences*, Cambridge, UK: Cambridge University Press.

Evans, J.R. and W.M. Lindsay (2005), *Management and Control of Quality*, Mason, OH: Thomson Southwestern, pp. 28.

Forrester, J. (1961) *Industrial Dynamics*, Cambridge, MA, USA: MIT Press.

Fridsma, D.B. and J. Thomsen (1998) "Representing medical protocols for organizational simulation: An information processing approach," *Computational & Mathematical Organization Theory*, **4**(1): 71–95.

Gibbons, R. (1998), "Incentives in organizations," *Journal of Economic Perspectives*, **12**(4): 115-132.

Giddens, A. (1984), *The Constitution of Society: Outline of the Theory of Structuration*. Cambridge, UK: Polity Press.

Harrison, J.R. and G.R. Carroll (2002), "The dynamics of cultural influence networks," *Computational & Mathematical Organization Theory*, **8**: 5-30.

Heath, J. (2001), *Communicative Action and Rational Choice*, Cambridge, MA, USA: MIT Press.

Holmström, B. and J. Roberts (1998), "The boundaries of the firm revisited," *Journal of Economic Perspectives*, **12**(4): 73-94.

Howard-Grenville, J.A. (2005), "Review Essay: Explaining shades of green: Why do companies act differently on similar environmental issues?", *Law and Social Inquiry*, **30**(3): 551-581.

Katz, D. and R.L. Kahn (1978) *The Social Psychology of Organizations*, 2nd edition, New York, NY: Wiley.

Kerr, S. (1975) "On the folly of rewarding A, while hoping for B," *Academy of Management Review*, **18**: 769-783.

Lawrence, P. and J. Lorsch (1967), *Organization and Environment*, Cambridge, MA, USA: Harvard University Press.

Lempert, R. (2002), "Agent-based modeling as organizational and public policy simulators," *Proceedings of the National Academy of Sciences*, **99**(3):7195-6.

Levitt, R.E. (2004), "Computational modeling of organizations comes of age," *Computational & Mathematical Organization Theory*, **10**: 127-145.

Lin, Z. (2000), "Organizational performance under critical situations: Exploring the role of computer modeling in crisis case analysis," *Computational & Mathematical Organization Theory*, **6**(3): 277-310.

Locke, E. and D. Schweiger (1979), "Participation in decision making: One more look," *Research in Organizational Behavior*, **1**: 265-339.

Manahan, S.E. (1999), *Industrial Ecology: Environmental Chemistry and Hazardous Wastes*, Boca Raton, FL: Lewis/CRC Press, pp. 53.

Miles, R.E. and H.R. Rosenberg (1982), "The human resources approach to management: Second-generation issues," *Organizational Dynamics*, **Winter**: 26-41.

Panayotou, T. and C. Zinnes (1994), "Free-lunch economics for industrial ecologists," in R.H. Socolow, C.J. Andrews, F. Berkhout and V.M. Thomas (eds), *Industrial Ecology and Global Change*, Cambridge, UK: Cambridge University Press, pp. 383-397

Parker, M. (2001), "What is Ascape and why should you care?", *Journal of Artificial Societies and Social Simulation*, **4**(1): 5-20.

Phelan, S.E. and Z. Lin (2001), "Promotion systems and organizational performance: A contingency approach," *Computational & Mathematical Organization Theory*, **7**: 207–232.

Prigogine, I. and I. Stengers (1984), *Order Out of Chaos: Man's New Dialog with Nature*, New York, NY: Bantam Books.

Reimann, M.D. and J. Sarkis (1996), "An infrastructure for agile product development systems," in L.F. McGuinness et al. (eds), *Flexible Automation and Intelligent Manufacturing*, New York, NY: Begell House, pp. 57-66.

Sanders, K. and S.K. Hoekstra (1998), "Informal networks and absenteeism within an organization," *Computational & Mathematical Organization Theory*, **4**(2): 149-163.

Scott, W.R. (2001), *Institutions and Organizations*, 2nd edition, Thousand Oaks, CA: Sage Publications.

Shannon, R.E. (1975), *Systems Simulation: The Art and Science*, Englewood Cliffs, NJ: Prentice-Hall.

Simon, H.A. (1947), *Administrative Behavior*, New York, NY: Macmillan.

Snijders, T.A.B. (1998), "Methodological issues in studying effects of networks in organizations," *Computational & Mathematical Organization Theory*, **4**: 205-215.

SourceForge (2006) *Repast Home Page*, https://sourceforge.net/projects/repast. Last accessed 2 September 2007.

Swarm Development Group (SDG) (2006), *Swarm Development Group Wiki*, http://www.swarm.org. Last accessed 2 September 2007.

Torenvlied, R. and G. Velner (1998), "Informal networks and resistance to organizational change: The introduction of quality standards in a transport company," *Computational & Mathematical Organization Theory*, **4**(2): 165-188.

Weick, K.E. (1979), *The Social Psychology of Organizing*, Reading, MA: Addison-Wesley.

Williamson, O.E. (1979), "Transaction-cost economics: The governance of contractual relationships," *Journal of Law and Economics*, **22**: 233-271.

Zeggelink, E.P.H., R. van Oosten, and F.N. Stokman (1996), "Object oriented modeling of social networks," *Computational & Mathematical Organization Theory*, **2**(2): 115-138.