Path Dependence and Adaptive Processes among Public Agencies

Andrew B. Whitford
Professor of Public Administration and Policy
University of Georgia
204 Baldwin Hall
Athens, GA 30602 United States
+1 706.542.2898
Fax: +1 706.583.0610
Email: aw@uga.edu

Abstract

This paper examines the twin roles path dependence and adaptation play in a model of bureaucratic performance. I illustrate the paths that bureaus take when moving along performance and technology search paths using evidence from a set of simulations.

Recent literature on path dependence in political and bureaucratic settings suggests "increasing returns" support a range of observed behavior in governments and polities. Political science and policy studies have used path dependence to help explanation that behavior partly because of the role of historical elements in shaping long-run government performance.

In contrast, an older but less-developed perspective on political and bureaucratic systems emphasizes the role of adaptation in decision-making. Bureaus are adaptive due to computational and epistemic uncertainty, among other factors.

I describe a model in which a bureau experiences performance, sets goals, and searches for new technologies in a fully adaptive manner. Theories of path dependency emphasize the importance of different institutional arrangements. I then show how the model allows the comparison of the impact of different institutional frameworks, including competition, the ability to imitate other bureaus, and the potential for "random punishment".

1 Introduction

Public agencies are rigid (Downs 1967). They are unyielding to change. They follow incremental decision processes (Lindblom 1959). The thrust of many modern analytic narratives about the state of public bureaucracy is that they are unable or unwilling to move from one path of "doing things" (or *not* "getting things done") to other pathways that make them more efficient, effective, or able to achieve higher-valued public goals. The impacts of these arguments have included a slew of attempts to reform bureaus (dating back at least two thousand years) and (more recently) attempts to create a "new public management" or to replace the state with an array of for- and not-for-profit entities.

The purpose of this paper is to make the claim that two broad theories of organizational change help explain these attributes of bureaus. First, the developing literature on "path dependent" processes in organizations (or societies, markets, firms, etc.) argues (broadly) that "history matters" or (narrowly) that processes are self-reinforcing (often due to the presence of increasing returns (Arthur 1989, North). Second, a parallel (but largely disconnected) literature on "adaptive processes" provides unique insights into the attributes of how the agents that compose those organizations make decisions, and how those decisions contribute to stasis or rigidity in organizational change (Levinthal and March 1981). The two literatures provide a more complete picture of the nature of bureaus and helps contribute to the empirical regularity of stasis. This paper investigates and illustrates an adaptive mechanism that can provide a formal backbone to the regularity. In a sense, the paper joins together the "history-shaping" focus of path dependency with the mechanics of a self-reinforcing process in an organization.

This approach also provides insight into several key aspects of path-dependent processes in public organizations. Recently Paul Pierson (2000) criticized the increasing use of "path dependence" in political science as an explanation for observed political or public behavior, in part because of the lack of rigor usually attained. Specifically, he sees the key to understanding path-dependent processes in political science as hinging on the relative presence of increasing returns. He also sees the presence of those processes as helping support key claims of the "historical institutionalism" movement in the social sciences, including the roles of timing and sequence, the consequences of contingent events, the rigidity of particular courses of action, and the possibility of "punctuated equilibria". While Pierson's views have helped support new initiatives on path dependency, they also lack mechanics. This lack of mechanics is particularly noticeable when considering application to bureaus, in part because political science has largely followed the mechanism design literature in increasing formalism regarding how and why bureaus make strategic decisions (e.g., Padget 1980, Bendor and Moe 1985).

For instance, many of the analyses of the 1993 Government Performance and Results Act (GPRA) in the United States have described concerns about the way bureaus have responded in terms of changing their levels of performance (National Academy of Public Administration 1999). Since GPRA, bureaus in the U.S. have regularly focused on organizational change, comparison, and redirection, with varying degrees of success. Agencies have adapted, partly because after GPRA they could be compared, partly because they compete with one another in terms of "results", and partly because they have learned from their own and others agencies' experiences. As always, the problem has been developing models of bureaucratic attention, investment, and adaptation that are manipulable for investigation of the relative contribution of their different parts.

Agencies adapt to their surroundings imprecisely, just as individuals do in Herbert Simon's bounded rationality and in other forms of incremental decision making(e.g., Bendor 1995; Cyert and March 1963; Lindblom 1959; Newell and Simon 1972; Simon 1947, 1957, 1976). The model here illustrates how this adaptive process, involving the bureau's willingness to innovate, make investments, and react to the performance it experiences, is a form of self-reinforcing process. It also shows how the outcomes of the process are different for different agencies based on concurrences of small events, and how bureaus experience different kinds of outcomes based on different "initial conditions".

Yet, these adaptive processes live inside (and are part of) larger path-dependent processes, such as whether and how bureaucratic performance is affected by comparison, competition, imitation, or strategic redirection. Of course, bureaus respond to external and internal circumstances, and (to a degree) can be redirected. Agencies may compete with other agencies over resources, constituencies, or policy areas, or an interested overseer may use comparisons to reward or to punish. Their counterparts may offer opportunities for imitation. Professional analysts (including the media, consultants, and academics) learn about bureaucratic performance and productivity by comparing.

This essay proceeds as follows. In the second section I review foundational arguments about adaptation and competition in bureaus, and place these in context by describing the overall shape of the literature on path dependency. I then illustrate the power of an adaptive mechanism as a baseline for understanding bureaus, and describe how these type of model can be used to "test bed" the specific structural features of comparison, imitation, and random punishment that form a path-dependent "backdrop" within which the adaptive process operates. Last, I conclude with thoughts on this modeling strategy, the results, and directions for future research.

2 Adaptation and Path Dependence in Bureaucracies

This essay combines two major initiatives in the study of public bureaus as organizations: the modeling of agencies as adaptive systems, and the power of path dependency as an explanation for the variety of behaviors we observe among bureaus in a rich array of case studies. In this section, I first describe how adaptation in public bureaucracies and technological search are key ways of describing and modeling the actions of agencies. Then I turn to comparison, competition, and information revelation in the context of the path-dependency paradigm.

We have long built complex adaptive systems for understanding organizations and societies, dating at least to cybernetics (Ashby 1957), largely because the technologies associated with this approach give us the tools to study organizational behavior through the lenses of process, benchmarking, and flexibility. The literature has moved a long way in the last two decades as computation and computer-aided simulation procedures have become widespread. The emphasis in this paper is on the adaptive mechanism rather than on simulation, agent-based modeling, or computational organizational theory per se (Burton and Obel 1995; Carley and Prietula 1994). The simple reason for this focus is that adaptation in public agencies has long been considered the basic way of describing public decision-making, in part due to the significant presence of computational and epistemic uncertainty in the public decision process (e.g., Grandori 2001).

In the public sector, we have described public bureaucracies as adaptive at least since Lindblom's work on incrementalism in policy (1959, Bendor 1995) and Simon's description individual decision making, bounded rationality, and incrementalism in public agencies (Newell and Simon 1972; Simon 1957, 1976). The growth in the number of dynamic models of adaptive behavior in the behavioral tradition is truly astounding (Axelrod 1976, 1997; Bendor and Moe 1985; Cohen 1981, 1984; Cohen and Axelrod 1984; Cohen, March, and Olsen 1972; Crecine 1969; Cyert and March 1963; Kollman, Miller, and Page 1992; March and Simon 1958; Padgett 1980; Simon 1947). The adaptive approach illustrated here is an agency that searches for technology in a couple of different ways in order to make probabilistic changes to its observed (either perceived or objective) performance. The bureau can use search to change the relationship between performance and action, but its ability to do so is error-prone.

Specifically, the bureau uses to "technology" to apply knowledge for practical purposes (policy-making, policy-setting, policy implementation, etc.); essentially, technology is a method. Of course, public bureaus constantly invest in the search for and implementation of new technologies, if only to develop new data processing and investigation technologies. Bureaus invest in “soft” technological search when they select and train personnel.

Witness, for example, the recent debate over the “proper” way to regulate financial markets in the U.S. The search is on for new ways to form and implement regulation, ranging from the use of information technology (new surveillance), to soft technologies like the involvement of different accounting forensics specialists. Perhaps the most innovative search for new technologies is the search for new organizational forms, such as the possible agency for consumer protection (in current debates) or the independent regulatory commission (at the turn of the Twentieth Century). In the adaptive mechanisms illustrated in this paper the latter form of technology is the least-investigated; the main focus is on hard or soft technologies like those regularly invested in by public health organizations, regulatory bodies, or the military.

The focus here on public organizations is a lower-level application of an approach regularly followed in the study of firms and other economic organizations. Indeed, in economic history studies, adaptation is a feature of and a solution to problems in large, modern organizations. Technology and business institutions coevolve, with agents choosing new forms if they solve organizational problems (Chandler 1990). Over the last several decades, economists have sought to better understand how technology and efficiency entwine as mechanisms of adaptation and selection. Arthur (1989) and David (1985) show that technology at a point in time is truly cumulative since today’s advances build from and improve upon status quo technologies.

A full description of the role of technology in economic organizations is beyond the scope of this paper, but technology plays a two fundamental roles in this paper.[1] The

[1]For at least the past half-century, economists have recognized that knowing how countries attain different technologies helps explain economic growth (Schumpeter 1934; Solow 1956, 1957; Fagerberg 1994; Sen 1999). For instance, endogenous growth theory argues that improved technology caused the improved living standards experienced after the Industrial Revolution (Grossman and Helpman 1994). In 1957, Solow showed

first is that technological search is at the core of the adaptive process, and the since the search process is fraught with error, the organization cannot guarantee outcomes. This adaptive mechanism is fundamentally dynamic, with reinforcing processes, and likely sensitive to initial conditions. Each of those aspects is consistent with the stories embedded in the incrementalist and bounded-rationality perspectives of public organizations described above.

The second role comes in the connection to path dependency. Specifically, path dependency can be conceived as a situation where "preceding steps in a particular direction induce further movement in the same direction" (Pierson 2000, 252). Yet, this is really a description of increasing returns, which is particularly important when focusing on political and public organizations and systems. Systems with increasing returns are unpredictable, can be inflexible, have feed back, and are potentially path-inefficient (Arthur 1994). Sequences are critical in these systems. And the situation of increasing returns is particularly important in economies, especially so due to their relative rarity (e.g., see the debate in endogenous growth theory). For Arthur, tech-

that in the U.S. economy, technical change (including both hard technologies and improved education) explained seven-eighths of the growth in output per hour worked; in contrast, changes in labor and capital (long-held to be the main causes of growth) together explain little (Solow 1957; see also Maddison 1987). One view is that firms learn by doing, and that those differences in knowledge alter how labor and capital combine across countries (Arrow 1962; Romer 1986; Lucas 1988). A parallel, incentives-based view is that firms invest in technology (not just knowledge and human capital) based on their expected private returns (Grossman and Helpman 1991); yet, the returns to society may actually exceed firms private returns depending on how investment in technology expands the overall knowledge base. In turn, policy helps explain whether firms (or other groups such as universities) make those investments (Romer 1994). Two ways technology affects the economy include the specific benefits firms gain from investment, and investments overall effect on the knowledge base (also called knowledge capital). Some innovations are especially likely to have broad effects. General purpose technologies (GPTs) affect multiple economic sectors and complement other, more specific innovations (Bresnahan and Trajtenberg 1995); for example, this happened with the introduction of the telephone. GPTs help expand the knowledge base and also help firms innovate, but GPTs may benefit some sectors more that others, just as some countries may be more likely to adopt them (Helpman 1998, 2). Empirical studies show that technology helps firms compete at both the national and international levels (Fagerberg 1994). If technology varies with geography, so will the degree to which firms and organizations embed or embody technological capabilities (Solow 1960). In some markets firms simply do not have incentives to invest in technologies that help spur the growth cycle (Stiglitz 1989, see also Pietrobelli 1994). In practice, the focus in economics both theoretical and empirical on technology has pushed people to study how technology varies across countries (e.g., Archibugi and Coco 2005), mainly by measuring national differences in the access people have to technology like GPTs or advanced education.

nology and its social context create the conditions for increasing returns when there are setup costs, learning effects, coordination effects (i.e., network externalities), or adaptive expectations. Economists recognize the possible presence of such returns in economic geography, where clusters form and persist, and even in the case of "packages" of institutions that give some economies distinct advantages (North 1990). As Pierson notes, the problem for economists is to describe how likely are increasing returns, and the answer probably depends on how well other features (such as competitive markets) "weed" them out.

But just as public organizations can search for and implement technologies, so do those technologies affect the likelihood of increasing returns in governance. Pierson's view is that political life is relatively more prone to experience increasing returns, with the result that we are more likely to observe certain features in the areas of government where they operate. He sees four features as especially important; we are more likely to observe multiple equilibria, the contingent effects of small events on long-term consequences, the importance of timing and sequences, and observable inertia (Pierson 2000). While his real interest is broader political life, it is remarkable how well these features work as observable empirical regularities for bureaus. Some work well, while some do not (multiple equilibria). Small events do seem to matter (otherwise it is difficult to explain some agencies' peculiar jurisdictions). Timing and sequencing are especially important in policy formulation processes, where it sometimes seems that people decide early on the solution and then go looking for a problem to solve with it. And, as noted at the beginning of the paper, many commentators think the words "inertia" and "bureau" are interchangeable.

The position in this paper is that the attributes of increasing returns that enhance path-dependent processes act as shaping forces that structure the way adaptive processes inside bureaus encounter the world (and each other). Like all agencies, an adaptive agency resides in a larger social system that is concerned about social efficiency, the output of government, and political responsiveness. Political overseers shape public agencies and what they do. The degree of network externalities, high fixed costs, or learning effects can enhance increasing returns, and thus lead to the path dependent processes that describe why agencies end up where they are.

The next step is to illustrate how an adaptive process can itself provide insight into key observations about public organizations, such as how where the agency moves depends on its initial or starting conditions, or how self-reinforcing mechanisms inside the agency can lock it into a behavior pattern even when there is no possibility for the kinds of externalities that help explain path dependent processes. Given this context, the model then allows us to concentrate on the relative contribution of different degrees or types of interaction among adaptive agents. I illustrate three types below: comparison, punishment, and imitation. In the baseline model the bureau has measurable attributes of performance, goals, and technology, but can only search for technology (a new means of implementing policies, by procedures, information services, human knowledge, etc.). The bureau sets goals/performance targets, observes performance, adopts a technology, and sets new goals.

The specific structural features that can be compared to this are comparison and punishment of agencies, through simple competition, differential punishment of agencies, imitation ("learning by others' doing"). I also discuss an alternative to these deterministic systems in the form of nave stochastic/random punishment.

3 Simple Illustrations of Self-Reinforcing Processes, with Applications to Theories of Path Dependency

3.1 Basic Agency Adaptation

I start with Levinthal and March's (1981) model of adaptive organizational search. Organizations change through the adaptive search for new technology; an organization's behavior and performance reflects the consequences of simple adaptation in an ambiguous environment. This is a generalized search model in an organizational learning framework. This model serves as a baseline (Levinthal and March 1981, 327). I assume organizations change their performance and goal-attainment behavior by searching across possible implementation technologies. Agencies search by innovating (Radner 1975; Knight 1967; Nelson and Winter 1978). An agency discovers exogenous technology, but the process of searching depends on past performance, goals, and search investments.

Agencies perform, have goals (that may be attained), and make expenditures to determine the best way to implement policies. To start, the agency sets a performance goal based on its past experience. The agency has only one way to attain that goal: to search for a new implementation technology. When it begins the search process, the agency examines its past performance and expenditures to assess the effectiveness of different types of technology search. Doing so helps determine the agency's propensity to search, its search efficiency, and the amount of resources it allocates to search.

The key is how many times the bureau will undertake either a refinement search or an innovation search. Refinement involves only local search; innovation is far-ranging (but not global). The agency compares the results of the search processes and its status quo technology. Once implemented, the agency experiences a new performance level based on its technology, expenditures, and environmental variation. The agency then adapts its goals based on this performance experience and starts the process again.

Specifically, performance in a given time period, P_t, is a function of the technology it implements (T_t), search costs, and an exogenous and varying environmental variable (a_t). The agency makes expenditures for two types of search: refinement (R_t) and (I_t):

$$P_t = (1 + a_t)T_t - R_t - I_t \tag{1}$$

At the beginning of the period, the agency sets a performance target or goal (G_t) before experiencing performance. It modifies that target given performance as:

$$G_t = b_1 P_{t-1} + (1 - b_1)G_{t-1}. \quad (2)$$

P_{t-1} is organizational performance in time $t - 1$. $b_1 > 1$ makes the target an exponentially-weighted moving average of past performance; I assume $b_1 = 1$. Note that a moving target that depends on past performance is an observable consequence of path dependency; here targets are strictly fluid, rather than my forcing them to weight past performance heavily. The observable outcomes and targets are due to the behavior of the agent, rather than a "history-forcing" approach. Alternative simulations can easily account for that possibility, though.

Performance, goals, and technology are simple indices (or complex combinations that comprise complex phenomena) (Jones 2001). Organizations may pursue multiple goals simultaneously; this can affect search behavior in an organization, but may even enhance organizational performance (Cohen 1984).

Technology largely determines the agency's performance and the search process is what the agency controls. The agency never forgets its past technology, which forms a baseline against which all new technologies are compared. Technological search proceeds by refinement and innovation. The agency chooses the best of three technologies: that obtained from refinement (T_r) and innovation (T_i) search, or its technology from the past period:

$$T_t = max(T_r, T_i, T_{t-1}). \quad (3)$$

A general model would allow technology to change over time (to decay or improve). I assume constant technology to isolate the effects of organizational search on agency performance.

The key decision is for the agency to decide the resources it will commit to search for a new technology. The agency compares its performance and the search resources it committed in the past period. The agency determines whether it met its goal: if it is met, it deems "successful" the type of search in which it engaged (innovation or refinement). Meeting a goal is the message the agency receives about the appropriateness of its past behavior. It is interested only in whether it meets its goal, not in how much it is exceeded.

The agency assesses the value of engaging in any search at all. This general propensity for search ($S_{s,t}$) depends on whether search resources were made available and the agency met its performance target or goal ($Q_{s,t-1}$):

$$S_{s,t} = Q_{s,t-1}b_2 + S_{s,t-1}(1 - b_2) \tag{4}$$

$Q_{s,t-1}$ is a test that takes the value "1" in two cases: if search resources were expended and the performance goal was met; or, if search resources were not expended and the performance goal was not met. b_2 determines the path of this propensity; it can limit the agency's ability to make sudden changes - failure would not easily translate into lower search investments, and success would not necessarily cause more.

Second, the agency's propensity to undergo a particular kind of search ($S_{i,t}$ or $S_{r,t}$) depends on a test relating the size of the type of expenditure and past performance: innovation ($Q_{r,t-1}$) or refinement ($Q_{i,t-1}$).

$$\begin{aligned} S_{r,t} &= Q_{r,t-1}b_3 + S_{r,t-1}(1 - b_3) && (5) \\ S_{i,t} &= Q_{i,t-1}b_4 + S_{i,t-1}(1 - b_4) && (6) \end{aligned}$$

$S_{r,t}$ and $S_{i,t}$ do not necessarily sum to one because an agency may not search at all; their sum may exceed one given the differential allocation of resources between the two types.

While this model allows testing the differential effects of learning rates, I first assume that $b_2 = b_3 = b_4 = 1$ so that the search propensities are functions of the past period's performance only. These specific learning rates make the change for a search propensity incremental in goal attainment; this is a restrictive assumption, but it is consistent with past work on learning, organizations, and policy formation (Lindblom 1959). After that, I show the impact of allowing for making search propensities incremental in past search expenditures.

These search propensities drive both the amount of sampling and the search space through which agencies learn about alternate technologies. The agency's search resources in a time period are limited, and the amounts allocated for refinement and innovation search can differ. First, the agency determines the total amount of resources it can dedicate to search ($U_{s,t}$, given $S_{s,t}$):

$$U_{s,t} = S_{s,t}(P_{t-1} + R_{t-1} + I_{t-1}) \tag{7}$$

The restrictive assumption here is that the agency is not limited in search, performance, or goal attainment by a lack of resources.

Second, this total is allocated to the competing types. For public agencies, it is likely that the resources available for refinement are greater in lean times and that innovation resources are greater when the organization attains goals. If the performance goal was achieved, the refinement resources are greater than those for innovation:

$$U_{i,t} = U_{s,t} \tag{8}$$

$$U_{r,t} = (U_{s,t})^{\frac{1}{h_r}} \tag{9}$$

h_r, the refinement slack coefficient, controls the shift from one type of search to another given performance achievement. If the performance goal was not achieved,

$$U_{i,t} = (U_{s,t})^{\frac{1}{h_i}} \tag{10}$$

$$U_{r,t} = U_{s,t} \tag{11}$$

where h_i is the innovation slack investment coefficient.

Not all available resources must be expended in this model. The resources expended on a search type depend on the propensity to undertake a kind of search and the amount of search resources available for that type. On one hand, agencies with limited resources may expend them all. On the other, agencies meeting performance goals may not expend any available resources when they reach a balanced state of performance and goal achievement.

The search propensities and the resources available determine the innovation (I_t) and refinement (R_t) search resources expended:

$$I_t = S_{i,t}U_{i,t} \tag{12}$$

$$R_t = S_{r,t}U_{r,t} \tag{13}$$

These search resources help determine the agency's sampling efforts for refinement ($K_{r,t}$) and innovation ($K_{i,t}$):

$$K_{r,t} = k_r R_t E_{r,t} \tag{14}$$

$$K_{i,t} = k_i I_t E_{i,t} \tag{15}$$

$E_{r,t}$ and $E_{i,t}$ are search efficiencies that reflect that efficiency increases with increasing search, but at a decreasing rate (Levinthal and March 1981, 313).

$$E_{r,t} = (E_{r,t-1} + K_{r,t-1})^{\frac{1}{w_r}} \tag{16}$$

$$E_{i,t} = (E_{i,t-1} + K_{i,t-1})^{\frac{1}{w_i}} \tag{17}$$

All together, the bureau has four interlocking self-reinforcing processes on its ability to generate changes in the technology it implements. It is limited by the existence of resources to commit to search. Its propensity to search, regardless of resources, is limited by its past performance and expenditure experience. The agency's search efficiency is determined by its past search experience. Last, it is constrained by the weight applied to control the shift from one search type to the other. Most importantly, these are constraints on how the bureau searches for technology - not how it discovers.

The technology an agency implements depends on the number of technologies it samples, the underlying distribution of possible (or undiscovered) technology, and its current technology. The returns to the search agencies undertake depend on their expenditures, the efficiency of search, and the current opportunities for development. The likelihood of finding an appropriate change depends on the number of opportunities sampled.

Given $K_{r,t}$ and $K_{i,t}$, the search process is a set of draws from different distributions based on the technology implemented in the past time period. For refinement, in each time period, let $n = K_{r,t}$. A drawn procedure is a change from T_{t-1} where

$$T_{r,n} \sim N(0, V_{r,t,n}) \tag{18}$$

For $n = 1$:

$$V_{r,t,n} = c_1 T_{t-1} \tag{19}$$

For $n > 1$:

$$V_{r,t,n} = c_2 V_{r,t,n-1}. \tag{20}$$

In a time period, if search resources are great enough, the agency may make a number of draws from this distribution. As it does, the variance of the distribution falls over multiple draws. c_1 means that the variance in refinement depends on the past technology; c_2 means that that variance falls in the time period when an agency makes multiple refinement searches. The space of possible refinements becomes smaller over multiple searches.

For innovation, a drawn procedure is a change from T_{t-1} where

$$T_i \sim LN(0, V_i) \tag{21}$$

$$V_i = c_3 T_{t-1}. \tag{22}$$

c_3 plays the same role as c_1 above.

It is possible to learn to not search due to estimation errors:

$$EE_{r,t} \sim N(0, V_{e,r,t}) \tag{23}$$

$$EE_{i,t} \sim N(0, V_{e,i,t}) \tag{24}$$

$V_{e,i,t}$ is initially greater than $V_{e,r,t}$, but each declines with the implementation of a given type of change in procedure. The value of either a refinement or an innovation is known with certainty once it is implemented. It is known only with error when it is not implemented. Once drawn, T_r, T_i, and T_{t-1} are compared in each time period and the largest value is selected and implemented. $V_{e,i,t}$ shrinks if innovation is the final technology; $V_{e,r,t}$ shrinks if refinement is successful.

In this model, the agency chooses a technology, experiences performance, compares that level with its goal, and starts the search process again. I simulate this baseline model of simple adaptation for two agencies given the initial conditions and parameters in Table 1.

[Insert Table 1 about here.]

One way to present simulations is to generate a population of representative agencies, allow the population to evolve, and summarize the results for the population by either summary statistics or by plotting the trajectory for the average member (for an example, see Cohen, Riolo, and Axelrod 1999). The figures for the simple adaptive agencies are representative of those obtained from multiple simulations. The figures for the extensions (except where indicated) are also representative. These simulations were written in Gauss.

Figure 1 shows the performance and goal paths for two agencies where each agency evolves over one hundred time periods. Both agencies quickly settle down to equilibrium levels of performance. Under these initial conditions, agency performance and goals are in steady state. By giving the agency a "shock" in the first time period (forcing the agency to make three technological searches), changes in technology spur changes in performance and goals. Both agencies explore new technologies and alter their performance and goal attainment; for each, the result is a higher performance level.

[Insert Figure 1 about here.]

Note that the agencies, though they face the same initial conditions, do not end up in the same place. Small differences in the stochastic processes of innovation and refinement (of search) lead to lasting differences in the performance levels of the two organizations. Note, too, that differences in performance experienced based on those stochastically-driven different processes of search have change the behavior of the organization over time, by shaping future search processes. In Figure 1, $b_2 = 1.0$, but in

Figure 2 $b_2 = 0.5$. The pattern shows that changing the initial condition about how the bureaus "mix" the impact of two experiences (successful search and past search expenditures), leads over time to compression of the two bureau's experiences. We do not observe the level differences in their performance and goal attainment behavior. Last, note that in Figure 2 that goals slightly lag performance, while in Figure 1 goals closely track observed performance.

[Insert Figure 2 about here.]

Figure 3 shows the time paths of one agency's propensities to undertake search for technology (general, innovation, and refinement); only one agency's results are shown because of their similar experiences. The general and innovation propensities track one another (due to the averaging process) and settle down over time. As this clearly shows, the equilibrium performance level occurs because the agency becomes unlikely to search after a period of time. As expected, refinement dominates innovation because this type of search is most closely related to an adaptive theory of organizational behavior. In fact, the adaptive model specified here should produce this result.

[Insert Figure 3 about here.]

Figure 4 shows the outcome of search, or the likelihood of choosing an innovation, a refinement, or retaining the status quo. On this scale, the agency is choosing innovations when the score is close to two, refinements when the score is close to one, and the status quo when the score is close to zero. The agency chooses innovation in the first period, and then quickly rejects search outcomes in favor of the status quo. This also shows that the likelihood of success (performance exceeding or matching goals) goes to one as the likelihood of retaining the status quo goes to zero. Figure 5 shows how the number of searches approaches zero rapidly.

[Insert Figures 4 and 5 about here.]

This baseline model demonstrates three basic propositions about self-reinforcing processes in agency learning and change. First, technology search, as defined here, produces changes in performance and goals, subject to constraints on search. It is the central mechanism by which the organization changes its operating behavior. Second, this system reaches a steady state of performance, goal attainment, and acceptance of the status quo technology. Without disturbances, the agency is unlikely to alter its technology, its goals, or its search for new technologies. Third, multiple indicators of agency activity are useful for examining how agencies use the search mechanism to change their performance. Here, indicators go beyond simple performance attainment to include search propensities, success, and outcomes.

3.2 A Competition Illustration

Over the past four decades, competition has been offered to fuse the public goods provision with the efficiency of markets (Tullock 1965, Downs 1967, Niskanen 1971). Essentially, the competition movement is an exercise in the politics of bureaucratic structure, shaping the incentives of agencies to redirect their behavior to new ends. Niskanen argues that competition allows the comparison of relative prices of competing bureaus and so shifts power from bureaucrats to politicians. As Boyne (1998) notes, these models' core hypotheses are that under competition total spending on services will fall and technical efficiency will increase; they make no specific prediction about allocative efficiency under competition. Direct descendents are "contracting out" and compulsory competitive tendering (Niskanen 1968; Boyne 1998; Kettl 1993).[2]

[2] Refinements of Niskanen argue that the model inaccurately characterizes the interaction between bureaucrats and legislators (e.g., Blais and Dion 1991). Migue and Belanger (1974) note that bureaucrats may maximize goals other than the supply of public services. Conybeare (1984) reveals an implicit assumption of perfect price discrimination (see also Bendor, et al 1985; Breton and Wintrobe 1975). Miller and Moe (1983) show that modeling the legislature alters predictions about the power of privatization and competition. Empirically, there are mixed findings at best on competition and agency costs and output (Boyne 1998; Conybeare 1984; Higgins, Shugart, and Tollison 1987). Conybeare (1984) notes that even if multiple, competitive bureaus are competing for funding, rather than producing equivalent goods (as in McGuire, Coiner, and Spancake 1979), there may be negative side effects such as high monitoring costs.

The value of competition, as Miller and Moe (1983) show, is that it reveals information about actual supply costs, and thus places monitors in better decisional positions and enhances their power. The advantage of competition is how it reveals information by allowing comparison. Another reason for multiple and comparable agencies is the value of redundancy or parallelism to reduce system-level performance errors (Bendor 1985, Landau 1969, Heimann 1993). Agencies would respond to comparison, competition, and information revelation because of the real-world implications of failure. While bankruptcy is not possible in the public sector, there is agency deletion (e.g., the Interstate Commerce Commission), reductions in budgets (e.g., the Reagan-era Environmental Protection Agency), media coverage (e.g., the Bureau of Alcohol, Tobacco, and Firearms), external options for agency heads (e.g., the Department of the Treasury), the simple incentive just to do well in the organization (Edwards, Nalbandian, and Wedel 1981), and other nonmonetary incentives (Crewson 1995).

In the second illustration, two adaptive agencies compete on the basis of performance. The losing agency is punished, and agencies adapt their goals given their realized performance. A fundamental reason for competition is that comparison makes reveals information about what is possible for agency performance. The baseline model provides a point of comparison to this second, institutional layer. This is competition neither over finite resources nor over a pool of customers through price competition or rivalry. Perhaps the best metaphor is athletes competing in time trials. Specifically, the agencies have independent search processes, but their performance is compared at the end of each period. The lower-performing agency is punished and is given only partial credit. Its observed performance is recalculated, and the losing agency's new goal now depends on this partial credit. In this application, two agencies are compared with one another. "Thin" competition allows me to isolate the effects of competition, and may be realistic for public agencies.

The key is to consider the value of comparison in a world of increasing returns. Comparison is information - information that does not need to be discovered through investment and search. As such, comparison is a form of externality. But comparison is also a limiter of increasing returns, since it can enforce remediable increasing returns and force a type of "market competition".

In Figure 6, the losing agency's performance measure is halved and punishment is symmetric (agencies are not distinguished between). Figure 6 shows the time path for the original ($b_1 = 1.0$) initial conditions, where the model now includes a competition step. The paths in this figure are substantially different from those in Figure 1. Both two display greater variation in their performance over the iterative process; the goal path of each agency tracks its performance path. The strong oscillation in the agencies' performance and goal paths is linked. Essentially, the "meeting-separating" pattern is a core result of this competitive process as over time the aggressive movement of one bureau to increase its performance leads to lagged increases in the other's performance. Eventually, both agencies settle into separate performance equilibria.

[Insert Figure 6 about here.]

As above, the propensities help gauge how an agency's performance and goal find equilibrium. Figure 7 shows the time path for one agency for the three propensities. In contrast to Figure 3, the search propensities do not settle down quickly. Moreover, innovation spikes follow punishment phases, and refinement spikes follow innovation spikes: punishment induces innovation and local search follows innovation. Figures 8 and 9 visually confirms punishment's role in driving the agency's internal search process. The amount of search switches for the first periods of the simulation, but the number of searches falls over time.

[Insert Figures 7, 8 and 9 about here.]

When an agency is docked for underperformance, its key indicator is whether its credited performance does not meet its internal goal. This disparity generates any internal changes that the agency makes in response to competition and punishment.

Specifically, these signals knock the agency off the performance-goal achievement steady state it naturally obtains in the underlying adaptive model. In this model, the central effect of competition and comparison is that agencies can find new levels of performance because performance inequities motivate their enhanced search for new technologies. Essentially, then, this model visually illustrates the point of Pierson (2000) that external mechanisms may (in some cases) reduce the impact of increasing returns and force bureaus off comfortable and inertial paths. However, the competition shown in this illustration does not force a common equilibrium. Note in Figure 6 that the agencies again separate after initial periods of wrangling over position. While early gains are wiped out when other bureaus invest in search, eventually the difference in performance (and the impact of a continuing path of punishment) keep the devalued agency from investing, and so they stay in a position of lower performance.

3.3 An Illustration with Differential Punishment

An alternate case is where competing agencies are punished differentially. In this illustration, the simulation involves a "90/10" punishment scenario. In this exercise, one underperforming agency receives 90 percent credit for its performance; the second receives 10 percent credit if it has lower performance. The logic of this experiment is to uncover the potential differential impact of punishment on agency performance. Are lower punishments reflected in smaller agency increases in search, expenditures, and performance achievement? Figure 10 shows the paths for the performance and goals of the two agencies in one simulation.

[Insert Figure 10 about here.]

This simulation highlights the differential responsiveness of an agency given the two punishment scenarios. The shallow valleys shown for Agency 2 reflect the 90 percent credit rule; the deep valleys for Agency 1 reflect the 10 percent rule. Note that Agency 1 - even given the strict 10 percent rule - finds a way to recover to globally higher performance levels. At a point in the model's evolution, Agency 1 and 2 diverge to different equilibrium performance levels. Figure 11 shows pairs of equilibrium performance values for 100 simulations. For two agencies, 60 percent of the "10 percent" agency's equilibrium values are higher than that for the "90 percent" agency. However, the situation separates: for the other 40 percent, the "10 percent" agency's values are substantially lower; the "90 percent" agency's mean value is significantly higher. This difference is visually confirmed in Figure 12, and also is confirmed statistically [unpaired t test with assumption of unequal variances; $t = 4.7902$; Satterthwaite's degrees of freedom, df = 108.345; difference significant at better than $p = 0.0001$].

[Insert Figure 11 about here.]

[Insert Figure 12 about here.]

The reason for this separation is how the agencies search for technology. Punishment means a lower credited performance and necessitates the enhanced search for innovation, but this can be sustained only for a short period of time. If the payoff to search is slow coming, search stops and an bureau falls into a position with low performance and no innovation. The core test within the agency for determining propensities remains whether performance fails to meet expectations (given punishments). Within this model, systematic significant punishments can spur greater global performance than light punishments (the punishment magnitude affects only $U_{s,t}$). An agency switches to refinement search when its performance exceeds its goals. If it is refinement mode and a significant punishment occurs, the agency's propensity for refinement falls. While the punishment may reduce $U_{s,t}$, it encourages innovation by discouraging refinement. Performance increases come from small levels of available funds because of the disjunction between credited performance and internal goals. On average, agencies that receive relatively heavy punishments (less credit for their performance) produce higher performance levels in the long run.

Perhaps more importantly, not in Figures 11 and 12 that multiple equilibria result from a model where the punishment rule differently affects the self-reinforcing processes in each bureau. The relative positions are roughly the same; what matters is which agency ends up where first. That part of the outcome is separate from the fact there are two equilibria (in a nutshell, equilibrium selection is stochastic, while the structure of the choice of the interaction produces the two equilibria).

3.4 Illustrating Imitation

Altering an agency's credited performance is just one way to alter its goals and search behavior. Instead, competing, adaptive agencies may imitate one another. Specifically, a losing agency may imperfectly imitate a dominant agency's technology. After it searches for a technology, experiences a level of performance, and competes with a counterpart, the losing agency imperfectly imitates a dominant agency's technology. Its technology is now a simple average of the two technologies in the system at that point in time. This is a fairly pure illustration of large externalities in a (admittedly small) network. Here, the relative position at each point in time (measured in terms of performance) dictates which way the technology flows in the network.

Figure 13 shows the time path for the same initial conditions as in the adaptive system, except that the agencies now compete (with symmetric punishments set at 50 percent credit), and losing agencies imitate. The simulation initially proceeds similar to that under simple competition: initial oscillation is followed by a period of stability.

[Insert Figure 13 about here.]

However, imitation produces a third dynamic. Here, a long period of stability is interrupted as one agency's performance jumps once its technology evolves to a point where real performance changes occur. At that point, the formerly dominant agency is punished and responds with a short period of innovation. Notably, in this case, the agencies actually switch positions - the dominant agency is replaced by the imitating one. As Figure 14 shows, the dominant agency's search propensities signal the subordinate agency's intrusion in the last innovation spike. Agency 2 shows a similar propensity to search, generating increased search, and causing a reversal of position.

[Insert Figure 14 about here.]

In this model, imitation plays a peculiar role: it provides a basis for displacing dominant agencies with agencies that have been punished by competition. Competition creates the opportunity for the first agency to dominate. Yet, imitation actually produces a globally-higher level of performance in the system, when it reaches equilibrium. For bureaus, the benefits of competition may be short-term benefits for the first agency. Under competition, imitation may produce a second set of benefits when the second agency matches and then replaces the first. Even so, imitation may have a downside when the dominating agency is beaten, and then imitates the second agency's past, underperforming technology. In this model, this agency's response is to adopt the underperforming technology, resulting in a lower equilbrium performance path.

The basic premise of this paper is revealed again: the process by which technology spreads (and thus the degree to which there are increasing returns) shapes the observed behavior of the two agencies. The bureaus settle into two different equilibria, each marked by inertia. After a long period of time, they switch positions due to changes in their relative willingness to search, and thus they also switch positions in terms of their relative performance. This is a classic punctuated equilibrium.

3.5 An Illustration with Naïve Random Punishment

The last illustration replaces competition and comparison with a system of random punishment. Random punishment alters the performance the agency is credited with; this credited performance feeds back into the agency's adaptive process of goal-setting and technology search. Neither competition nor imitation is allowed: the only shock to the system is random punishment. Specifically, I draw a uniform random number and impose random punishment to each agency one-third of the time. If an agency is punished, it receives a 10 percent reduction in its performance credit. Two-thirds of the time, an agency will not be punished. In this exercise, agencies are never punished simultaneously. Figure 15 shows the performance and goals for two agencies in one simulation.

[Insert Figure 15 about here.]

Across many simulations, the step-level increases in agency performance appear general. Again, the principal mechanism is self-reinforcing search: once either agency has arrived at an equilibrium performance level, the agency no longer searches for new implementation technology. As Figure 16 shows, the propensities to search for an agency are markedly different in this system. Search ebbs and wanes in direct relationship to the incidence of random punishment.

[Insert Figure 16 about here.]

We can debate whether the relevant indicator of agency performance is raw performance level or search propensity. The baseline agency lacking external shocks usually fails to search for new technologies in order to enhance its performance once in equilibrium. Different external shocks produce more complex behavior by this "model" agency; each generates internally-consistent agency search for new technology. Naïve random punishment produces the greatest amount of search and the most predictable pattern of increasing performance, but at the expense of predictability for the agency. And under even random punishment, there is significant nonlinear variation in the performance of agencies over time.

4 Discussion

Path dependence and adaptation play two roles in a model of bureaucratic performance. While recent literature on path dependence in political and bureaucratic settings centers on increasing returns support a range of observed behavior in governments and polities, most of the focus in those types of studies has been on how historical elements shape long-run government performance. A second, and probably older, perspective on political and bureaucratic systems emphasizes the role of adaptation in decision-making where computational and epistemic uncertainty (and other factors) cause agencies to move slowly along paths that are self-reinforcing.

This paper illustrates the paths that bureaus take when moving along performance and technology search paths using evidence from a set of simulations. In this model a bureau experiences performance, sets goals, and searches for new technologies in a fully adaptive manner. Because theories of path dependency emphasize the importance of different institutional arrangements, I illustrate the impact of different institutional frameworks, including competition, the ability to imitate other bureaus, and the potential for "random punishment".

Taken together, the simulations common elements drawn from Arthur and Pierson (among others) and shows that an adaptive mechanism can be used as a "test bed" for the assessment of increasing returns in organizations. Self-reinforcing process and path dependency are intrinsically interesting approaches for explaining observed empirical regularities in organizational and societal performance. They resonate largely because the stories told by economic historians and students of technology are sharp and clear. By and large, the study of self-reinforcing processes and path dependency have moved forward without strong mechanics. The Levinthal-March model is an appropriate platform for future development because of the strong focus on technology and its attainment.

If that represents an area of under-development, then the application to public organizations is even less-developed. There are many reasons for this lack, but none of these has to do with the lack of applicability. In fact, the inertia described above should be reason alone to study path dependency. Public management and administration scholars understand well the need to focus on technology in the organization. Policy historians see the impact of small events. These all contribute to a need, so the purpose of this paper is to start to help fill that gap. There are no "principles of administration" here, though, for the lessons are really about contingency and fitting, not global solutions.

5 Bibliography

Arthur, Brian. 1989. "Competing Technologies, Increasing Returns, and Lock-in by Historical Events." *Economic Journal.* 99(394): 116-131.

Ashby, W. Ross. 1957. *An Introduction to Cybernetics.* London: Chapman and Hall.

Axelrod, Robert M. 1976. *Structure of Decision: The Cognitive Maps of Political Elites.* Princeton, NJ: Princeton University Press.

Axelrod, Robert M. 1997. *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration.* Princeton, NJ: Princeton University Press.

Bendor, Jonathon. 1985. *Parallel Systems: Redundancy in Government.* Berkeley, CA: University of California Press.

Bendor, Jonathan. 1995. "A Model of Muddling Through." *American Political Science Review.* 89(4):819-840.

Bendor, Jonathan and Terry M. Moe. 1985. "An Adaptive Model of Bureaucratic Politics." *American Political Science Review.* 79(3):755-774.

Bendor, Jonathon, Serge Taylor, and Roland Van Gaalen. 1985. "Bureaucratic Expertise versus Legislative Authority: A Model of Deception and Monitoring in Budgeting." *American Political Science Review.* 79(4):1041-1060.

Blais, Andr and Stphane Dion. Editors. 1991. *The Budget-maximizing Bureaucrat: Appraisals and Evidence.* Pittsburgh, PA: University of Pittsburgh Press.

Boyne, George A. 1998. "Competitive Tendering in Local Government: A Review of Theory and Evidence." *Public Administration.* 76(Winter):695-712.

Breton, Albert and Ronald Wintrobe. 1975. "The Equilibrium Size of a Budget-maximizing Bureau: A Note on Niskanen's Theory of Bureaucracy. *Journal of Political Economy.* 83(1):195-208.

Burton, Richard M. and Borge Obel. 1995. *Design Models for Hierarchical Organizations: Computation, Information, and Decentralization.* Boston: Kluwer.

Carley, Kathleen M. and Michel J. Prietula, Editors. 1994. *Computational Organization Theory.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Chandler, Alfred. 1990. *Scale and Scope: The Dynamics of Industrial Capitalism.* Cambridge, MA: Harvard University Press.

Cohen, Michael D. 1981. "The Power of Parallel Thinking." *Journal of Economic Behavior and Organization.* 2:285-306.

Cohen, Michael D. 1984. "Conflict and Complexity: Goal Diversity and Organizational Search Effectiveness." *American Political Science Review.* 78(2):435-451.

Cohen, Michael D., James G. March, and Johan P. Olsen. 1972. "A Garbage Can Model of Organizational Choice." *Administrative Science Quarterly.* 17(1):1-25.

Cohen, Michael D. and Robert M. Axelrod. 1984. "Coping With Complexity: The Adaptive Value of Changing Utilities." *American Economic Review.* 74:30-42.

Cohen, Michael D., Rick L. Riolo, and Robert Axelrod. (1999) "The Emergence of Social Organization in the Prisoner's Dilemma: How Context-Preservation and other Factors Promote Cooperation." Santa Fe Institute Working Paper 99-01-002.

Conybeare, John A.C. 1984. "Bureaucracy, Monopoly, and Competition: A Critical Analysis of the Budget- Maximizing Model of Bureaucracy." *American Journal of Political Science.* 28(3):479-502.

Crewson, Philip E. 1995. "A Comparative Analysis of Public and Private Sector Entrant Quality." *American Journal of Political Science.* 39(3): 628-639.

Crecine, J.P. 1969. *Governmental Problem-solving.* Chicago: Rand-McNally.

Cyert, Richard and James G. March. 1963. *A Behavioral Theory of the Firm.* Englewood Cliffs, NJ: Prentice-Hall.

David, Paul. 1985. "Clio and the Economics of QWERTY." *American Economic Review.* 75(2): 332-337.

Downs, Anthony. 1967. *Inside Bureaucracy.* Boston: Little Brown.

Edwards, Terry J., John Nalbandian, and Kenneth R. Wedel. 1981. "Individual Values and Professional Education." *Administration and Society.* 13(2):123-143.

Grandori, Anna. 2001. *Organization and Economic Behavior.* London: Routledge.

Higgins, Richard S., William F. Shughart II, and Robert D. Tollison. 1987. "Dual Enforcement of the Antitrust Laws." In Robert J. Mackay, James C. Miller III, and Bruce Yandle, Editors. *Public Choice and Regulation.* Stanford, CA: Hoover Institution Press.

Jones, Bryan D. 2001. *Politics and the Architecture of Choice.* Chicago: University of Chicago Press.

Kettl, Donald F. 1993. *Sharing Power: Public Governance and Private Markets.* Washington, DC: The Brookings Institution.

Knight, Jack. 1992. *Institutions and Social Conflict.* Cambridge: Cambridge University Press.

Knight, Kenneth. 1967. "A Descriptive Model of Intra-firm Innovative Process." *The Journal of Business*. 40:478-496.

Kollman, Ken, John H. Miller, and Scott E. Page. 1992. "Adaptive Parties in Spatial Elections." *American Political Science Review*. 86(4):929-937.

Landau, Martin. 1969. "Redundancy, Rationality, and the Problem of Duplication and Overlap." *Public Administration Review*. 29(4): 346-358.

Levinthal, Daniel and James G. March. 1981. "A Model of Adaptive Organizational Search." *Journal of Economic Behavior and Organization*. 2:307-333.

Lindblom, Charles E. 1959. "The Science of Muddling Through." *Public Administration Review*. 19:79-99.

March, James G. and Herbert A. Simon. 1958. *Organizations*. New York: Wiley.

McGuire, Thomas G., Michael Coiner, and Larry Spancake. 1979. "Budget Maximizing Agencies and Efficiency in Government. *Public Choice*. 34(3 and 4):333-357.

Migue, J. and G. Belanger. 1974. "Toward a General Theory of Managerial Discretion." *Public Choice*. 17: 27-43.

Miller, Gary J. and Terry M. Moe. 1983. "Bureaucrats, Legislators, and the Size of Government." *American Political Science Review*. 77(June):297-322.

National Academy of Public Administration. 1999. *Helpful Practices in Improving Government Performance: A Compilation of Papers*. Washington, DC: National Academy of Public Administration.

Nelson, Richard R. and Sidney G. Winter. 1978. "Forces Generating and Limiting Concentration under Schumpeterian Competition." *Bell Journal of Economics*. 9:524-548.

Newell, Allen and Herbert A. Simon. 1972. *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

Niskanen, William A., Jr. 1968. "The Peculiar Economics of Bureaucracy." *American Economic Review*. 58:293-305.

Niskanen, William A., Jr. 1971. *Bureaucracy and Representative Governmen*t. Chicago: Aldine Atherton.

North, Douglass C. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge: Cambridge University Press.

Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge: Cambridge University Press.

Padgett, John. 1980. "Bounded Rationality in Budgetary Research." *American Political Science Review.* 74:354-372.

Pierson, Paul. 2000. "Increasing Returns, Path Dependence, and the Study of Politics." *American Political Science Review.* June.

Radner, Roy. 1975. "A Behavioral Model of Cost Reduction." *Bell Journal of Economics.* 6:196-215.

Simon, Herbert A. 1947. *Administrative Behavior.* New York: Macmillan.

Simon, Herbert A. 1957. *Models of Man: Social and Rational; Mathematical Essays on Rational Human Behavior in Society Setting.* New York: Wiley.

Simon, Herbert A. 1976. *Administrative Behavior: A Study of Decision-making Processes in Administrative Organization.* Third edition. New York: Free Press.

Tullock, Gordon. 1965. *The Politics of Bureaucracy.* Washington, DC: Public Affairs Press.

6 Tables and Figures

Table 1: Variable Definitions and Initial Conditions

Variable Name	Initial Condition
Environmental variable (a)	0.05
Weight given performance in changing goals (b_1)	1
Control of rate of change of refinement search propensity (b_2)	1
Control of rate of change of innovation search propensity (b_3)	1
Control of rate of change of general investment propensity (b_4)	1
Relation of variance in refinement to new technology (c_1)	0.01
Reduction of variance in refinement draws with experience (c_2)	0.98
Relation of variance in innovation to new technology (c_3)	0.01
Efficiency of refinement ($E_{r,t}$)	0
Efficiency of innovation ($E_{i,t}$)	0
Goal in current time period (G)	40
Refinement slack investment exponent (h_r)	1.1
Innovation non-slack investment exponent (h_i)	1.1
Innovation investment in current time period (I)	5
Number of refinement draws this period ($K_{r,t}$)	3
Number of innovation draws this period ($K_{i,t}$)	3
Constant for converting to refinement draws (k_r)	0.2
Constant for converting to innovation draws (k_i)	0.2
Performance in the current time period (P)	40
Refinement investment in the current time period (R)	5
Refinement search propensity in the current time period ($S_{r,t}$)	0.5
Innovation search propensity in the current time period ($S_{i,t}$)	0.5
Propensity to invest this period ($S_{s,t}$)	0.5
Standard deviation of refinement estimation error ($V_{e,r,t}$)	0.5
Standard deviation of innovation estimation error ($V_{e,i,t}$)	0.75
Standard deviation of distribution of refinement opportunities ($V_{r,t}$)	$c_1 T_t$
Standard deviation of distribution of innovation opportunities ($V_{i,t}$)	$c_3 T_t$
Refinement efficiency exponent (w_r)	3.5
Innovation efficiency exponent (w_i)	3.5

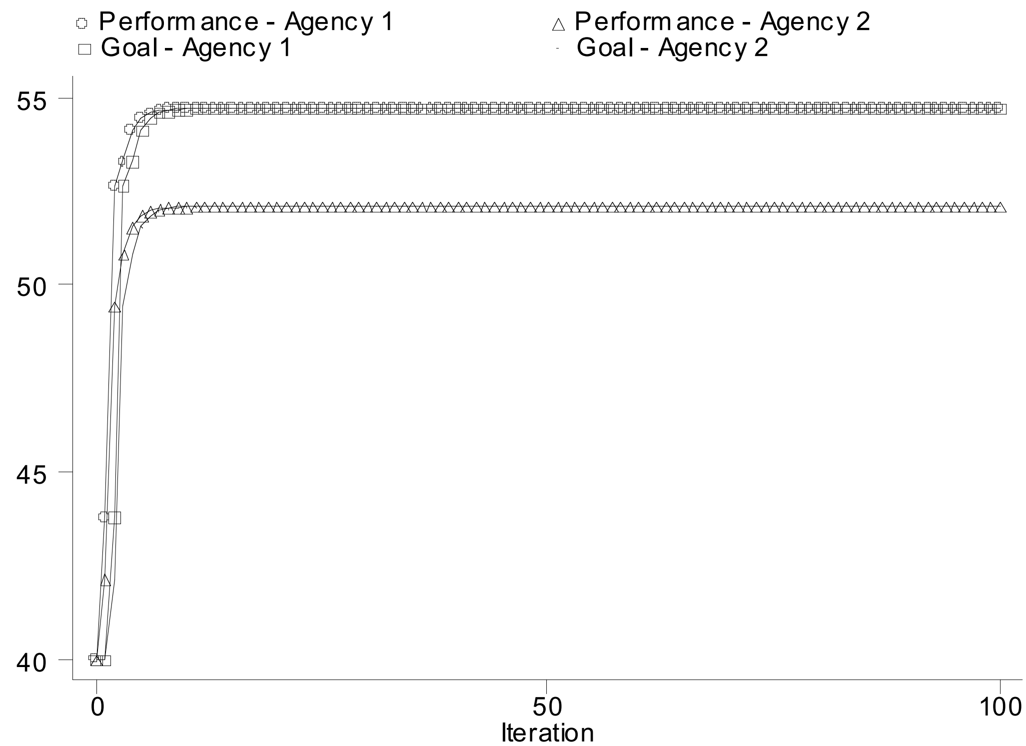

Figure 1: Adaptation: Performance and Goals: $b_2 = 1.0$

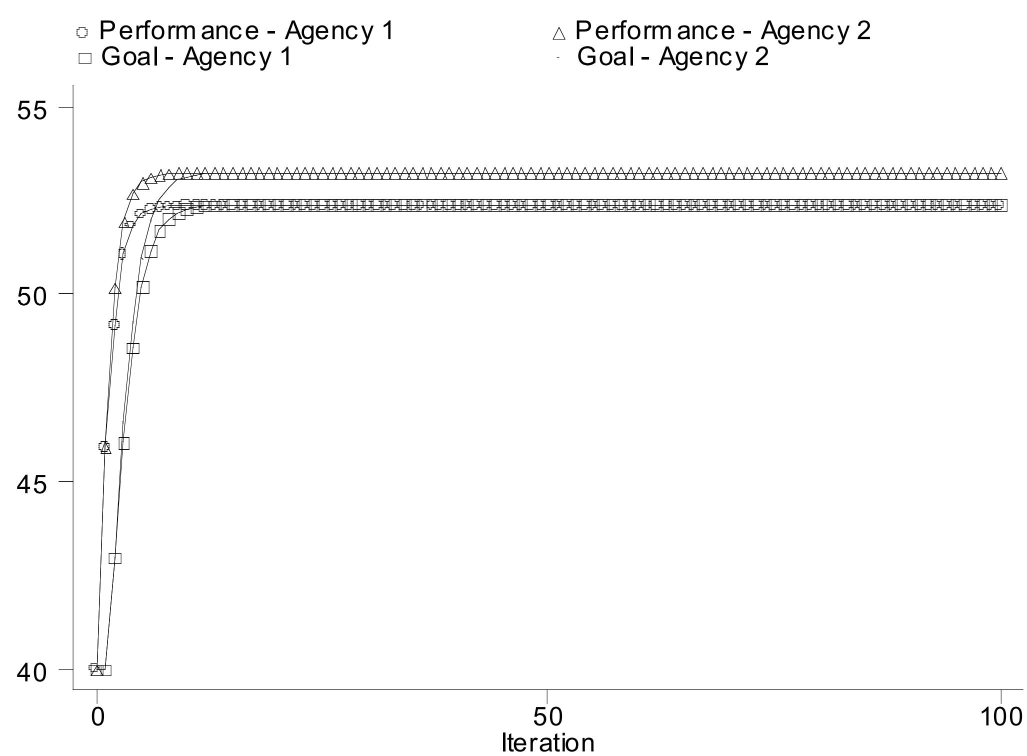

Figure 2: Adaptation: Performance and Goals: $b_2 = 0.5$

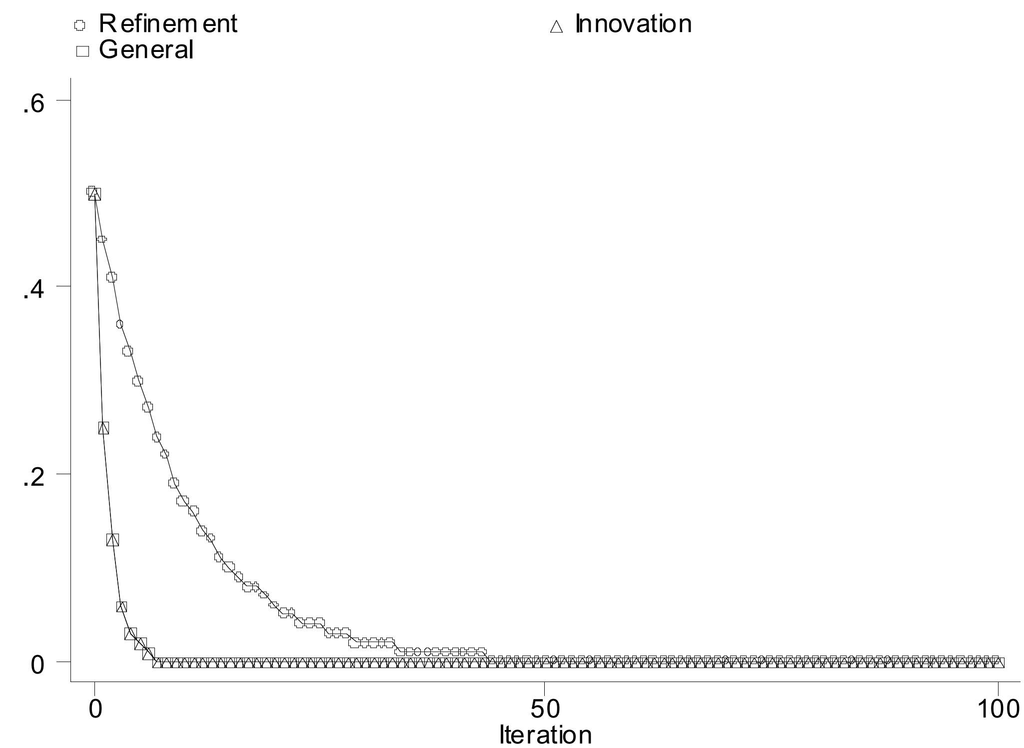

Figure 3: Adaptation: Propensities

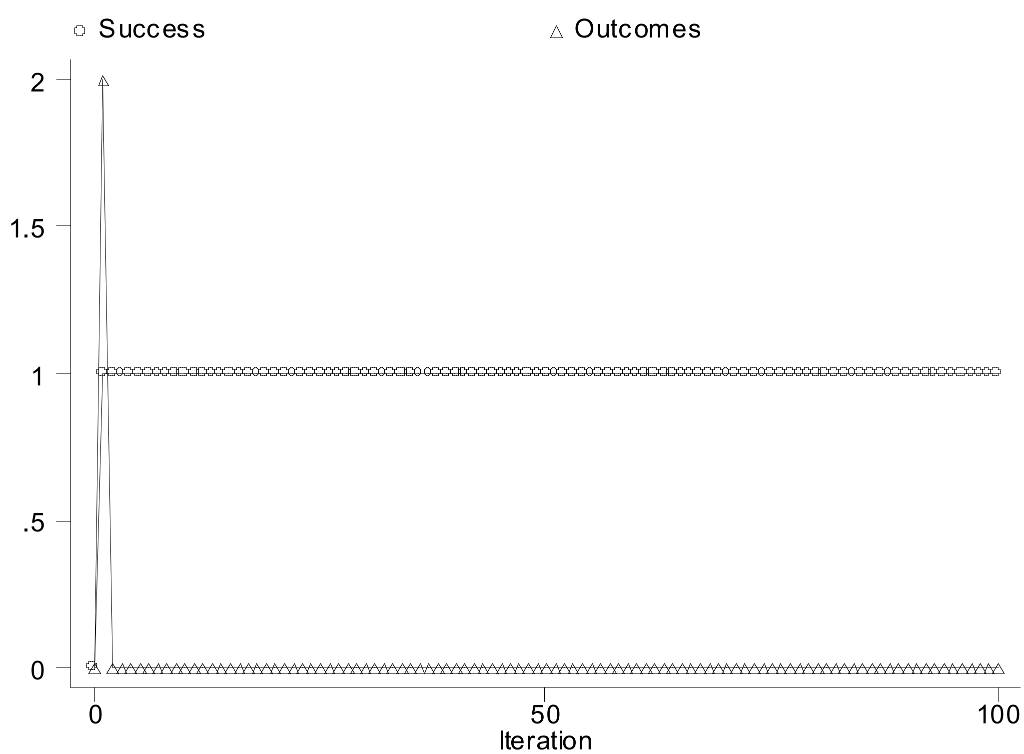

Figure 4: Adaptation: Success and Outcomes

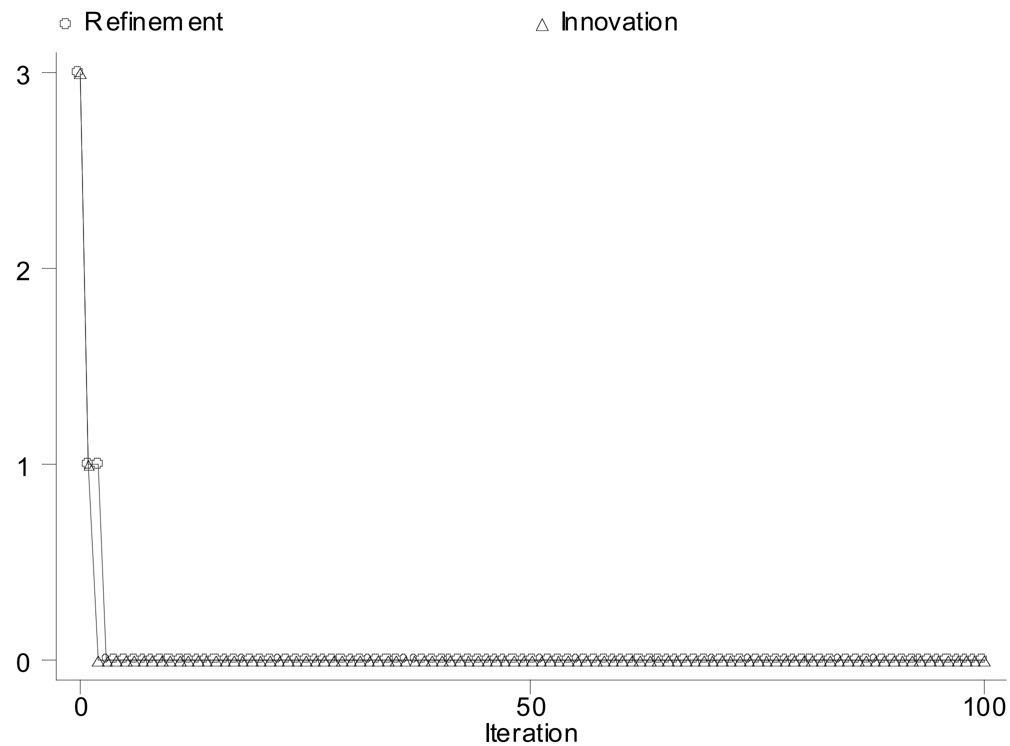

Figure 5: Adaptation: Number of Searches

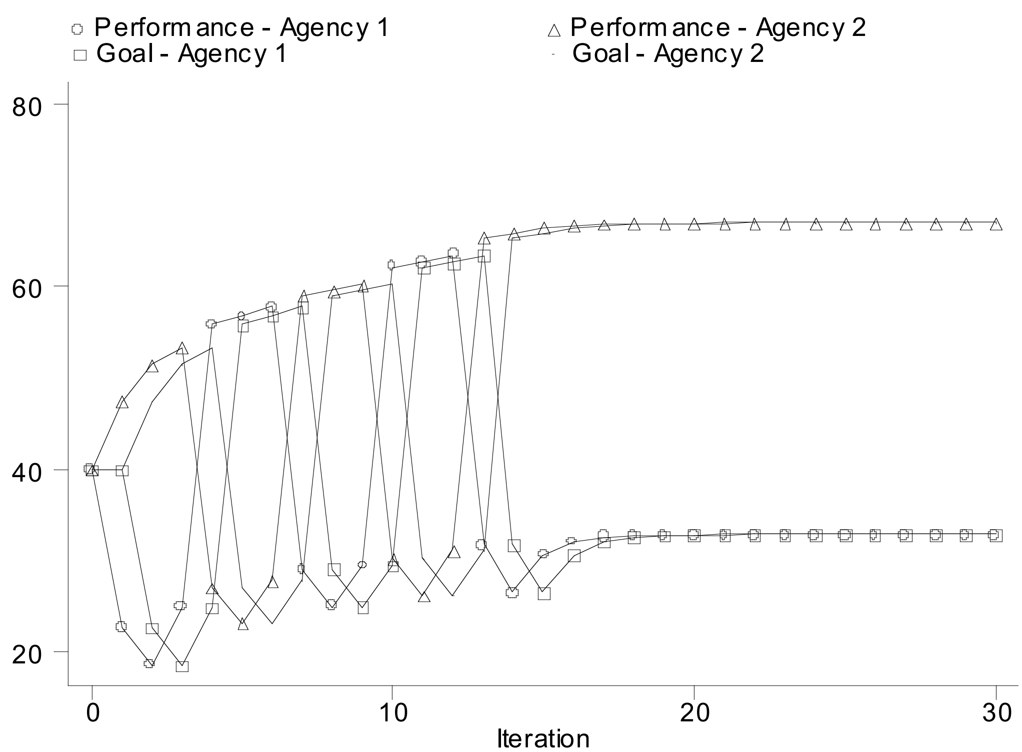

Figure 6: Competition: Performance and Goals

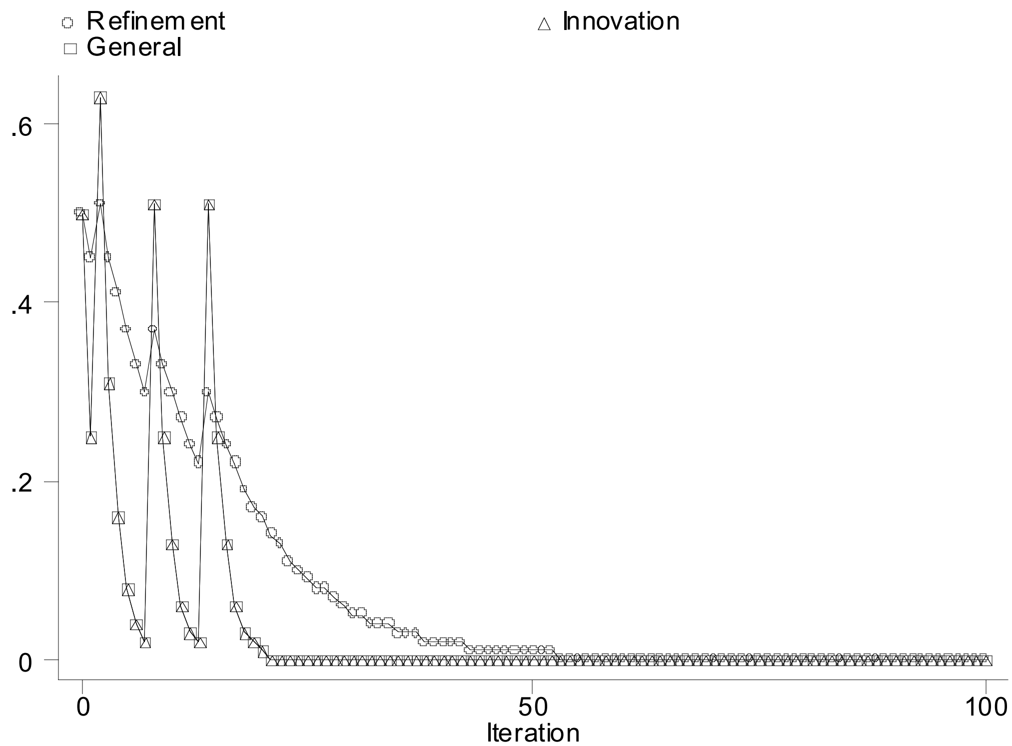

Figure 7: Competition: Propensities

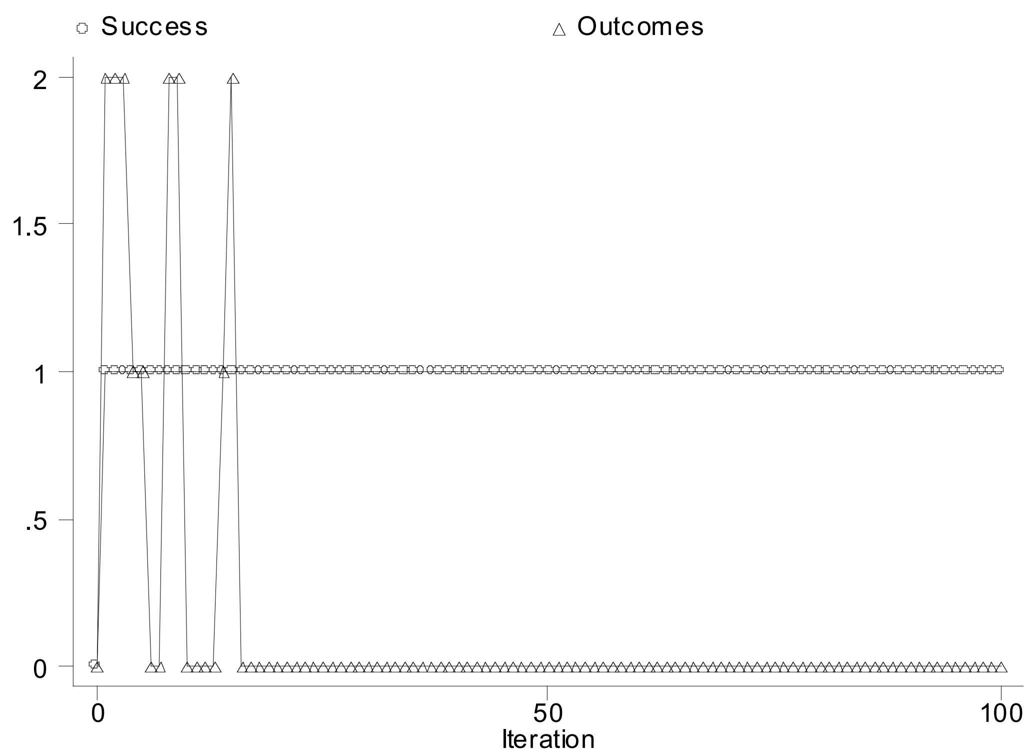

Figure 8: Competition: Success and Outcomes

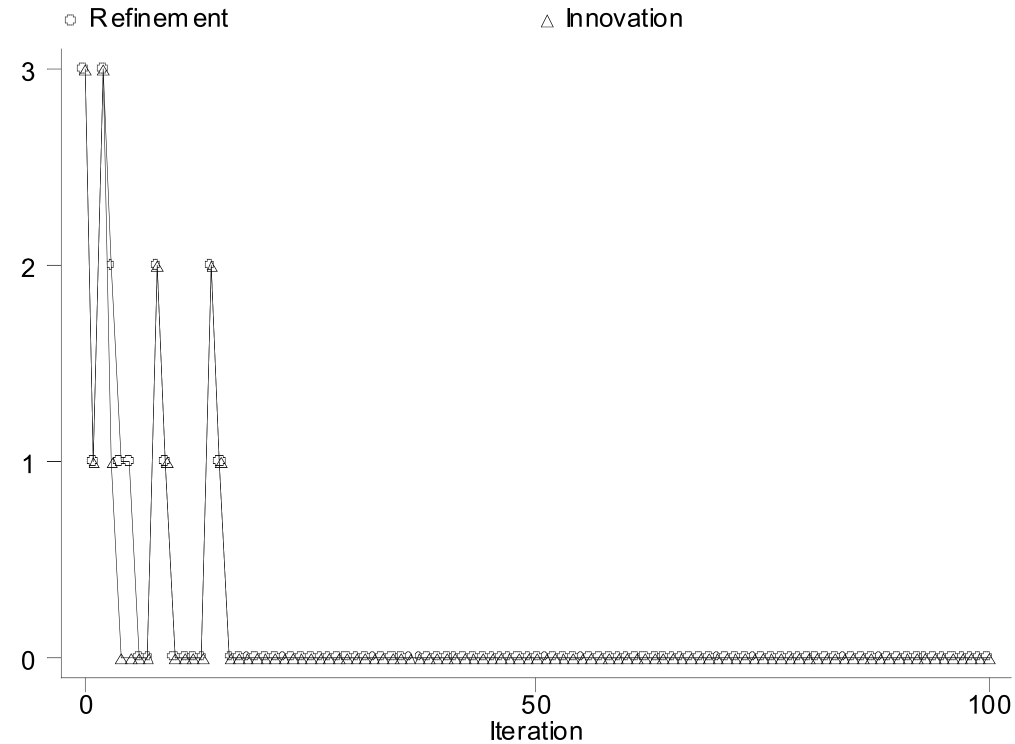

Figure 9: Competition: Number of Searches

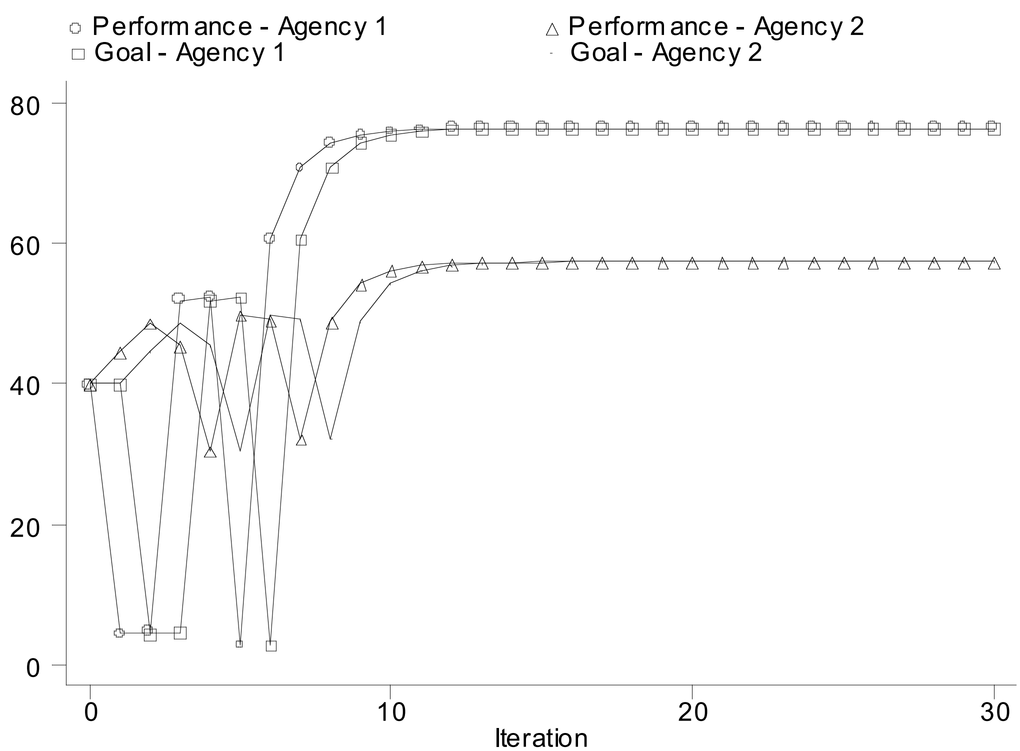

Figure 10: Asymmetric Competition: Performance and Goals

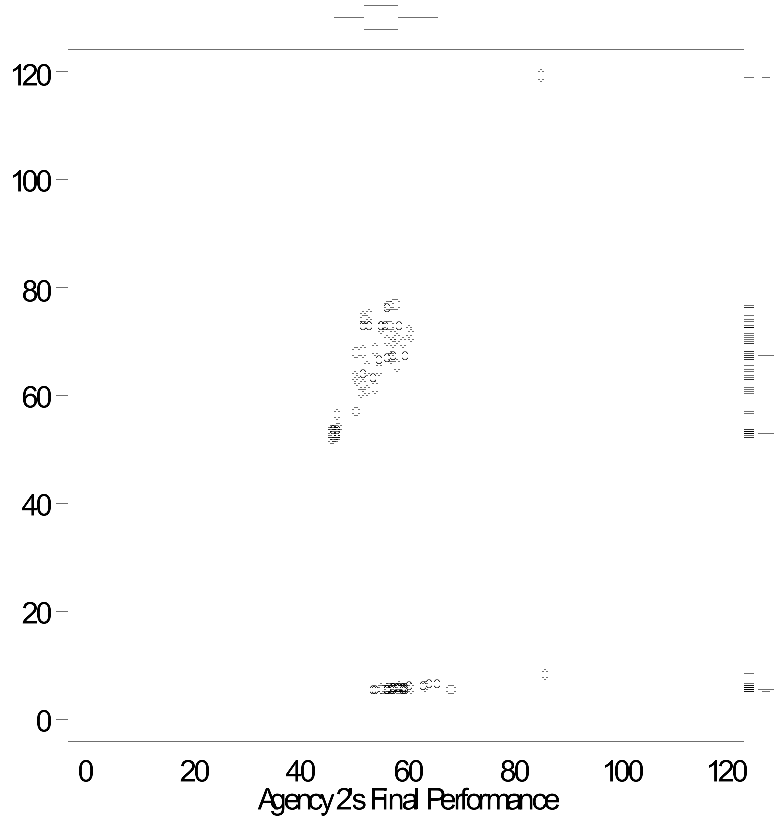

Figure 11: Asymmetric Competition: Performance of the Two Agencies

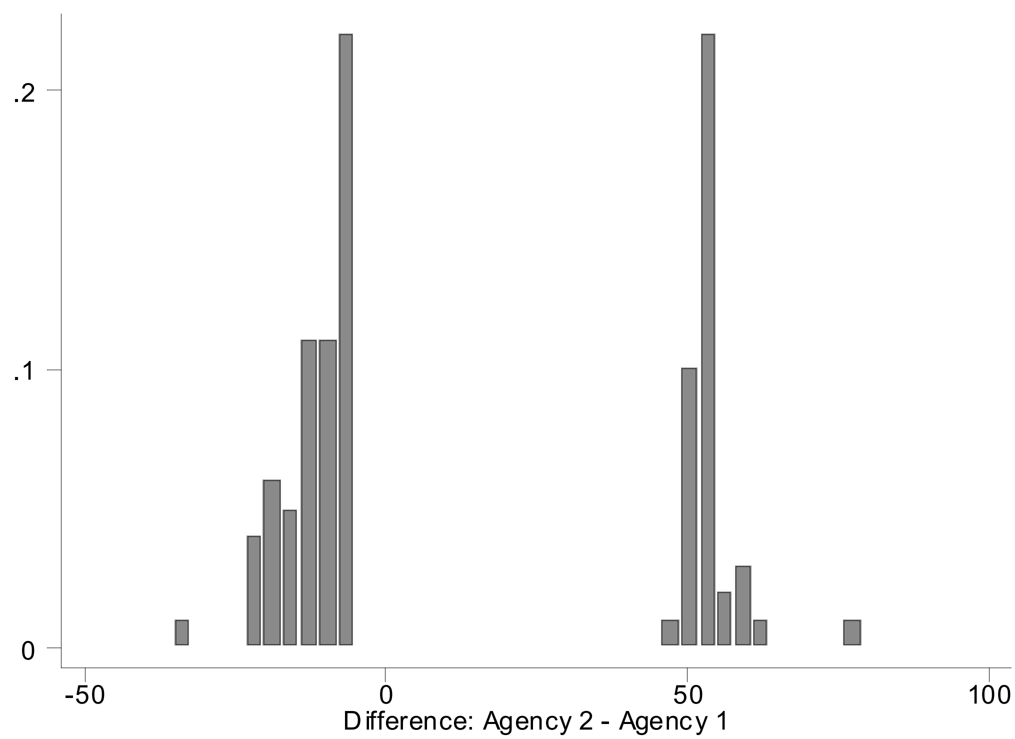

Figure 12: Asymmetric Competition: Relative Performance of the Two Agencies

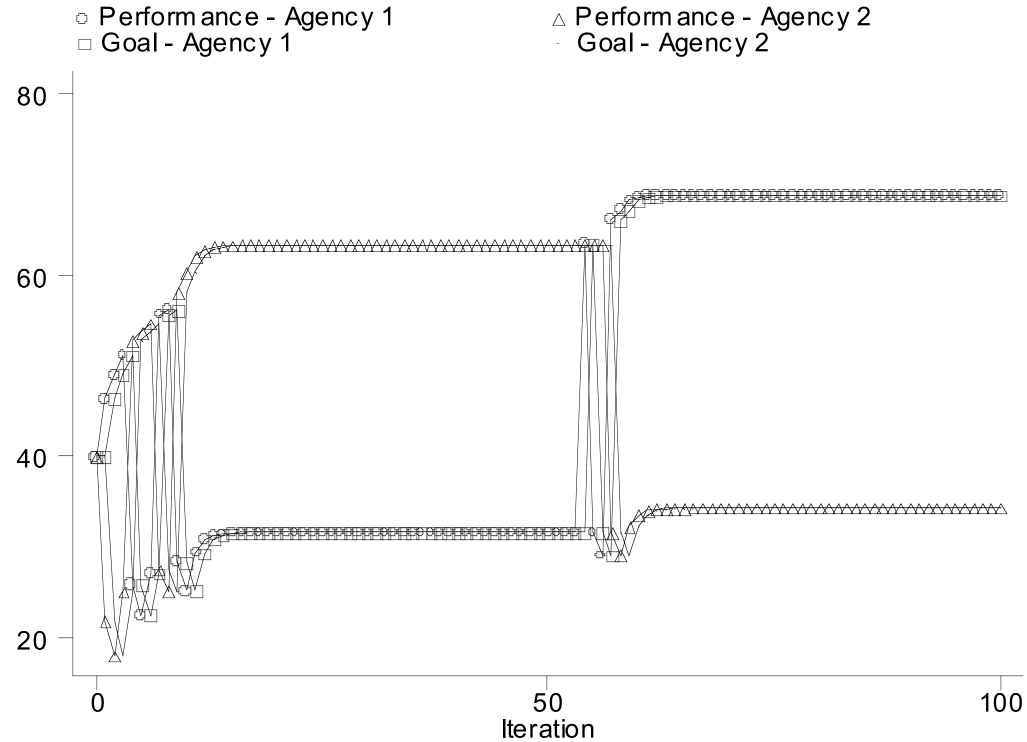

Figure 13: Competition and Imitation: Performance and Goals

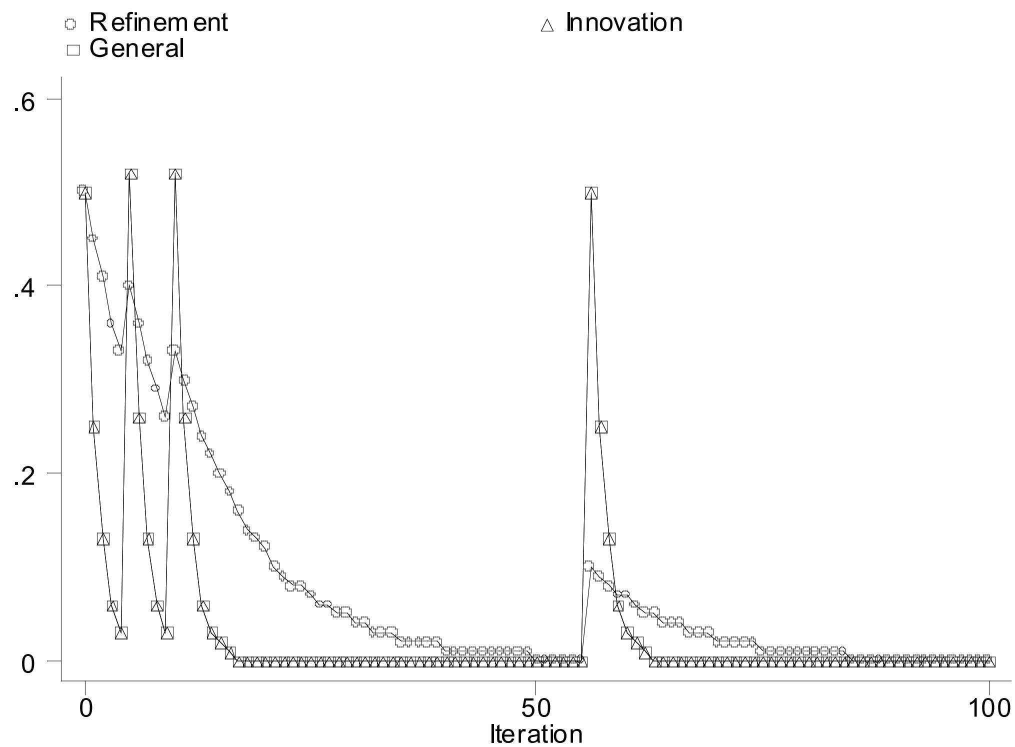

Figure 14: Competition and Imitation: Propensities for Agency 1

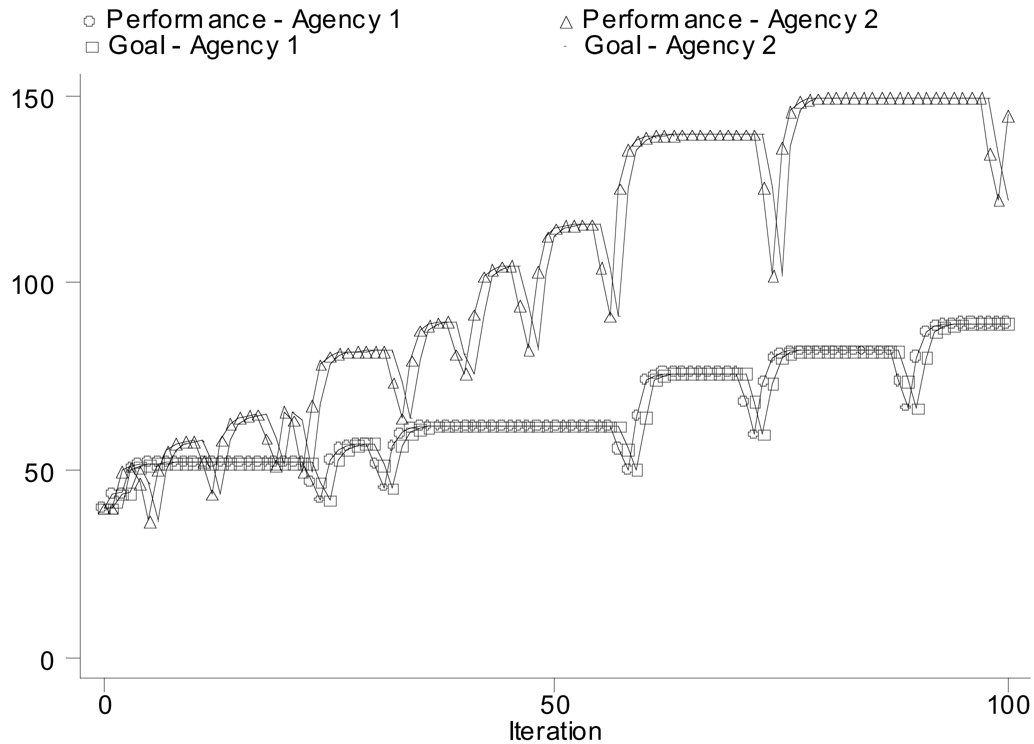

Figure 15: Random Punishment: Performance and Goals

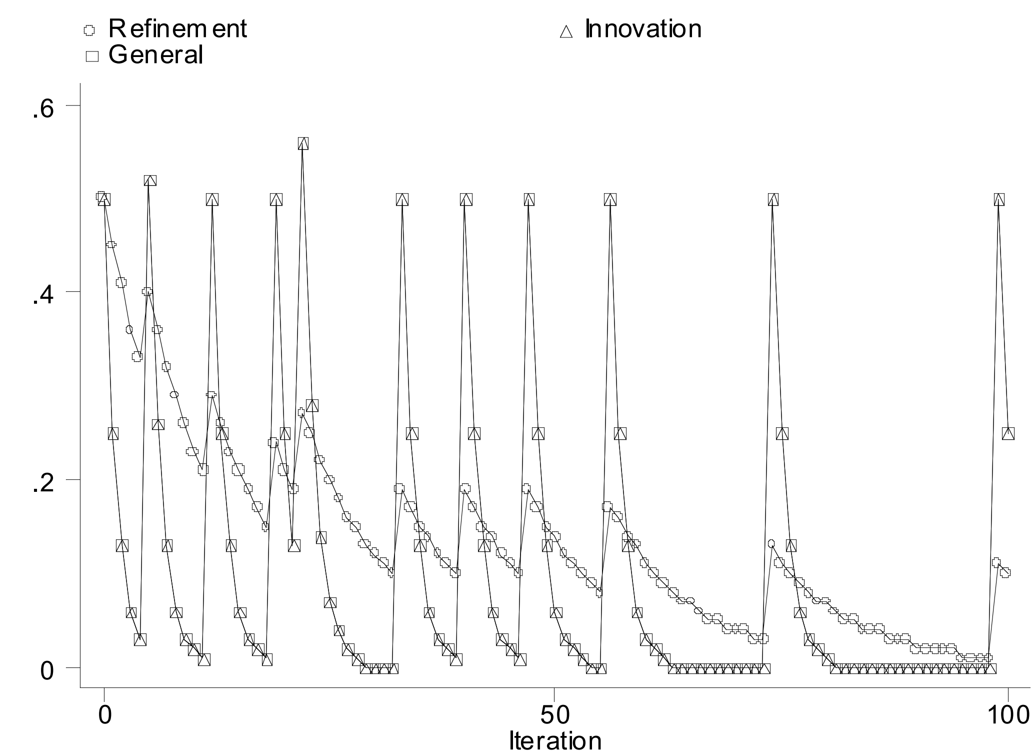

Figure 16: Random Punishment: Propensities for Agency 2