

Addressing Methodological Issues in Simulating a Human Resources Problem across Multiple Levels of Interest

Chester S. Labedz, Jr.

Boston College
140 Commonwealth Avenue, Chestnut Hill, MA 02467
(401) 524-7711
labedz@bc.edu

George Stalker, Ph.D.

Dataleaf Technologies, Inc.
486 Concord Street, Carlisle, MA 01741
(978) 369-7735
georgestalker@dataleaf.net

Abstract

Problems of strategic human resource management, such as proposed reductions in prospective retirement benefits under an American pension plan, present complex challenges. Typically, a firm pursues strategic objectives such as Cost control through changes in focal policies and programs affecting its workforce. At such times a firm should be wary of long-term consequences among individual employees, since the firm, its programs, and its employees, comprise a three-level feedback system. Unintended consequences at the micro-organizational level may lead in turn to additional unpleasant surprises at the program or firm levels. This paper discusses the development of a simulation model combining the approaches of agent based and system dynamics disciplines in addressing a client's multi-level concerns. It articulates differences between the agent based and system dynamics modeling disciplines in the service of this task, and discusses issues of implementation and praxis that have governed our design at points of tangency between these approaches in such a hybrid model.

Key Words:

Strategic human resource management

Simulating multiple levels of interest

Hybrid Model methodology

Heterogeneous agent model

The Context of this Simulation Problem

In managing the human resources of an enterprise “strategically”, things too often don’t go as planned, go wrong. Unintended and undesired consequences arise, even when strategic action has been planned carefully. These adverse results may arise at the levels of the firm, the individuals who work for it, or the practices, policies and programs that translate between these organizational and individual levels. For example, in pursuing competitive success by attempting to control its costs of production, a firm may reduce elements of its compensation and employee benefits (collectively, “total rewards”) offerings. In some cases, such human resource interventions will assist in achieving the firm’s strategic objective; in other cases they will backfire, harming its competitive position over time. Depending on the outcome, the human resource change will be embraced as successful or, to the extent circumstances permit, altered or abandoned.

Unintended consequences that emerge in this way are often not cost-free to an organization. Fay, Hempel, Director and Huselid (1997) have identified administrative, adjustment and direct costs that are usually associated with human resource management interventions. Each of these costs will be incurred in the course of a change, and some of these sums will have been “wasted” if the initial intervention produces unacceptable consequences. Then, if a firm tries to ameliorate or eliminate the undesired outcomes through additional changes, it likely will incur even more administrative, adjustment and direct costs.

In several recent discussions, Becker and various colleagues (Becker, Huselid et al. 1997; Becker and Huselid 1998; Becker, Huselid et al. 2001) have extolled the importance of a systems perspective to successful strategic human resource management (“SHRM”). Application of systems theory in actual organizational contexts works to identify and then to affect “causal feedback loops” that exist among the components of an organization’s system (Senge 1990; Senge, Roberts et al. 1994; Sterman 2000). Over time, these feedback loops can either reinforce intended results or deliver accidental, “unintended consequences”. Failure to appreciate these “powerful connections” and “deadly combinations”, as Becker et al. style them, are “...the greatest single challenge facing traditional HR managers as they make the transition to becoming true business partners” (Becker, 1997: 236).

Strategic human resource management presents a dynamic problem involving feedbacks among three levels of interest. These are: (1) the firm’s business objectives and results (macro), (2) the diverse inputs and outputs of individual prospects and employees (micro), and (3) intervening HR programs, policies and practices (meso). Although we do not fully develop the theoretical model here, it brings together the following literatures, among others. The open system view of managerial job behavior or throughput offered by Campbell, Dunnette, Lawler and Weick (1970), and largely adopted by Wright and Snell (1991), frames the individual level core of the model. To explain macro level concerns, we draw upon and extend the literature relating to management of uncertainty and project risk. Finally, we specify human resource subsystems (Fitz-enz 1984; Dreher and Dougherty 2002) and HR practices as translation mechanisms through which macro-level operational and strategic concerns affect heterogeneous micro-level employee attributes and, in an open system view, are affected in turn by them. The model identifies three necessary HR subsystems: total rewards management, ability management and opportunity management.

In the United States, changes in employee retirement {"pension"} benefits can be especially susceptible to deadly combinations, fraught with risks of firm-level resource waste and individual-level disgruntlement. Federal law imposes a lengthy written notice period before reducing, suspending or eliminating prospective accruals of tax-qualified pension benefits by employees. Benefits already accrued by the pension plan's participants are protected against reduction under a separate tax law provision, the "anti-cutback rule." Thus pension changes, once made, cannot be perfectly undone even if unintended consequences later emerge. Even in a case in which an employer errs and then reacts by "throwing money at the [pension] problem" that its changes have set off, it may be difficult for the firm to respond to and reverse ill effects on employee morale and retention without even more cost and unpredicted response. The recent conversions of traditional defined benefit plans to so-called "hybrid" pension programs in more than 130 American companies, including such celebrated firms as IBM and Ameritech (Zink 1997; Geisel 2004; Johnston 2004), have triggered unintended consequences measurable in dollars, distracted management focus, diminished employee satisfaction and application of skills, knowledge and abilities, and increased governmental scrutiny.

In this spirit, research has been undertaken to examine the following question: Why may competitiveness-driven reductions in pension benefit offerings unexpectedly reduce human resource competitiveness over time? We have begun to study this problem through a model of human resource processes that links operational variables at the level of the firm and individual characteristics at the micro level through the human resources subsystems. A 5000-employee, multi-site American industrial firm that is considering prospective pension reductions has provided us with five years of historical data (reference modes) across the three levels of interest enumerated above. We have developed causal influence diagrams and basic stocks and flows models, as called for in system dynamics practice, through interviews with its vice president for human resources and eight of his professional colleagues. He has listed a set of unintended consequences that he fears may occur if the pension reductions are implemented. From these, we developed hypotheses that predict the emergence of unintended consequences over time, as the cascading effects of policy resistance to the pension reductions by groups of current and prospective employees, and then of responses by the firm in turn. Table 1 summarizes these hypotheses and identifies measures within the stocks and flows model of the variables of greatest interest. In a model that includes system feedbacks, these variables both are influenced by prior factors and in turn affect subsequent ones, so they are not merely dependent variables. In the style of traditional variance analysis of human resource interventions, however, we denote these variables of interest "dependent" in Table 1.

We do not yet report hypothesis-testing results from the simulation of this potential human resources problem nor answer the research question, as our modeling is ongoing. Instead, as its title suggests, this paper builds to discussions (in its final sections) of boundary issues that we have met, and tentative modeling recommendations that we advance, while developing a hybrid model to test. Leading up to these discussions, we contextualize our research question in the human resources literature and our modeling approach in prior hybrid model discussions.

Table 1: Preliminary List of “Dependent” Variables Related to Specific Hypotheses

Hypothesis	Summary	Measure
Affected Employees		
1(a)	Affected employees retire later.	Credited Service at voluntary termination.
1(b)	Increased health benefits Cost results from 1(a).	Gross per capita Cost of employee health benefits.
1(c)	Increased compensation Cost results from 1(a).	Per capita Cost of pay programs.
1(d)	Increased Cash contributions to Retirement plans results from 1(c).	Per capita employer cash contributions.
Junior Employees		
2(a)	Junior employees’ career frustration increases due to 1(a).	Voluntary departure rates, age- and service-adjusted.
2(b)	Reduced recruitment lure results from 2(a).	Recruiting success ratios.
2(c)	Compensatory increases in other Total Rewards offerings result from 2(b).	Aggregate per capita Cost of non-retirement Total Rewards.
2(d)	Ability Management Costs increase with workforce churning.	Ability Management program Costs (e.g., administrative, adjustment and direct expenditures related to training).
Prospective Union Transferees		
3(a)	Rate of union transferees, if grandparented in retirement plans, is unaffected.	Transfer rate, time per supervisory opening.
3(b)	Rate of union transferees, if not grandparented in their retirement expectations, decreases following the nonunion pension change.	Transfer rate, time per supervisory opening, subsequent to effective date.
3(c)	Rate of union transferees, if not grandparented in retirement plans, temporarily increases prior to the change.	Transfer rate, time per supervisory opening, prior to effective date.
Minority Employees and Candidates		
4(a)	Minority employees’ career frustration increases.	Minority recruiting success ratios, and voluntary departure rates, age- and service-adjusted.
4(b)	Compensatory increases in Total Rewards for minority employees result from 4(a).	Aggregate per capita Cost of non-retirement Total Rewards.

Review of the Available Approaches to Simulation

The system dynamics approach models problems that appear within complex feedback systems. The basics of the approach are well stated in Parts I and II of Sterman's (2000) and chapters one through three of Coyle's (1996) texts and will not be repeated here. Several norms are promoted to ensure appropriate system dynamics modeling. To limit the number of variables to a manageable few, to help ensure that the model can be implemented, and to permit the simulation to run to completion, the modeler is enjoined always to model a problem and never a system. He or she needs to develop reference modes that trace the development of the problem over time, and to collect reference data sufficiently into the past to permit the model first to reproduce the problematic behavior (Sterman, 2000: 89-90) System dynamics models (almost) always examine current problems; they usually don't "do" potential issues, although there are notable exceptions, as in the "world" models that have built on the pioneering work of Meadows, Meadows, Randers and Behrens III (1972) and Forrester (1973)..

System dynamics models generally focus on overall policies and dynamic behavior at the level of a system of interest (for example, a firm, function or process). In general, the approach has been used to concentrate on the strategic concerns of the top management of an organization, as Coyle (1996: 15) points out. Although systems models often include stocks of individuals and trace their rates of entry and exit, these behaviors are based on average rates of change and not on individual decisions to enter or exit a stock. As one example, Paterson's faculty promotion chain model, adapted by Sterman (2000: 486), contains such rates. It models junior promotions as perfectly-mixed first-order processes, but describes full professors' retirements through third-order delays in order to approximate their generally-extended tenures before retirement. In each of the three stocks of professors, its members have completely homogeneous characteristics, because this model does not require otherwise. By and large, systems models don't "do" individuals, either.

Of course, intended and unintended firm-level consequences do arise out of individual circumstances. Needs may vary from worker to worker in important ways, depending on their combinations of personal and professional characteristics. In one example, a particular systemic change in human resources practices or policies may cause some employees not to enjoy the levels of financial security, professional opportunity or job satisfaction that they desire or had expected, while others may be advantaged by the change. In turn, these "harmed" individuals may become less productive workers or even leave the firm. Even so, one co-worker's loss may be another's gain, as an incumbent's departure may lead to another's promotion. Different total rewards costs may be associated with different individuals, and so the firm's resulting cost and its worker complement may not be what it intended to achieve through its strategic intervention.

The successful strategic management of human resources activity across time calls for promoting certain reinforcing feedbacks and for balancing a wide range of concerns across the three levels of interest outlined at the start. In qualitative research we have already conducted, human resource specialists comprising three Delphi-method panels of discussants (Dalkey and Helmer 1963) identified such concerns and created causal influence and stocks and flows diagrams connecting them. The modeling client is concerned with unanticipated consequences arising from proposed changes in retirement benefits, so we focus in Figure 1 on firm and individual level elements whose relationships are mediated by the meso-level structure of the

retirement program design. Nonetheless, the structure describes with little editing any other total rewards offering, such as incentive compensation, medical care coverage, or cash or deferred savings arrangements under sections 401(k) and (m) of the Internal Revenue Code. We provide for this multiplicity of similar program structures through the subscripting feature offered by the Vensim systems modeling application (Ventana Systems 2003), release 5.2. We employ *italics* to indicate the micro-level elements, and regular font for macro-level ones.

This simplified causal influences diagram features three principal “stocks”: (1) the accumulated cost to the firm in its current fiscal year of any one of its total rewards programs, net of any cost control measures that shift gross cost from it; (2) an individual’s current level of satisfaction with that program offering; and (3) the “individual”. The last of these three is an unusual, multidimensional element in a system model, but because we believe it depicts values that are measurable at any point in time, it has the basic nature of a “stock”, and for our model purposes it is an important one. As we discuss below, we model individuals as consisting of multi-valued, mutually independent states, the observed values (not quantities) of which are traced across time through subscripts attached to each person within our employee database. Much more will be said about the individual stock in the Agent Features discussion.

Other model features should be noted. Although we depict here only the retirement plan elements, our hypotheses also implicate elements of the ability management and opportunity management subsystems; these two are “backgrounded” just for the present. Thus, Sterman’s textbook formulation (Sterman, 2000: 758) of a firm’s workforce vacancy-filling problem can be used with slight modification in conjunction with this retirement plan model. The underlined elements in our influence diagram (market factors, and individual needs, characteristics and attributes) indicate its two principal connections to the broader problem of strategic human resource management. “Market Factors”, for example, takes on each of the following meanings: external labor market conditions, experienced at the levels of both individual employee expectations and the Human Resources budget; competitive pricing pressures, experienced at the level of the firm.

We have indicated two principal feedback loops, market forces and individual responses and we have directionally signed most of the causal influence arrows, but not all of them. Our markings suggest that these loops may operate in either a balancing or reinforcing mode. Thus, the market forces loop may add to cost pressures whenever external market forces (e.g., benefits inflation or a “tight” labor market) provide reinforcing feedback to Cost. Yet this loop may reduce such pressures if labor market or industry conditions are deflationary, or if a firm is compelled to slash its costs to compete successfully or to survive. This suggests the existence (not pictured here) of “tipping point” structures (Repenning, Gonçalves and Black 2001) that affect SHRM efforts and successes.

In a corresponding fashion, at a personal level, the individual responses loop suggests that employees likely will react differently to cost control efforts based upon combinations of their individual characteristics and attributes. Thus, medical cost shifting to employees and dependents who require treatment will more greatly affect those family units whose members are treated more often, and pension reductions will more immediately affect older workers who expect to need their pension payments sooner. This “immediacy of impact” element serves to highlight that, although we have not marked them in Figure 1, in this model “there are delays, and then there are delays”. The challenge of the pension reductions being considered by the client makes this clear. First, the period needed for careful consideration, design,

to a change in a model parameter may not be the same as the aggregate reaction of the individuals that he or she “represents”. Further, “...trying to explain the behavior of a group by that of one individual is constraining. The sum of the behavior of simple economically plausible individuals may [in fact] generate complex dynamics” (Kirman, 1992: 118). Citing Geweke (1985) and Kupiec and Sharpe (1991) for their explicitly-acknowledged examples of the problem, Kirman reminds us that policy changes instituted by governments (or firms, we add) in fact often set out to affect subject individuals in different ways, so that a representative individual constructed before a change may no longer be representative after it. The exhortation that Kirman addresses to economists equally commends to macro-organizational theorists a more realistic view of micro-organizational behavior: “Only if we are prepared to develop a paradigm in which individuals operate in a limited subset of the economy, are diverse both in their characteristics and the activities that they pursue, and interact directly with one another will economics escape from the stultifying influence of the representative agent.” (Kirman, 1992: 123, 134). Holland (1998) makes a similar point in the complex systems literature.

We have noted above that pension changes, once made, cannot be perfectly undone even if unintended consequences later emerge, and that those consequences often involve financial cost to employers, to particular employees or both, as well as distracted focus and skills application by specifically-affected individuals in management, in human resource functions and within the workforce. The research aims to offer both explanatory and advisory value. As a result, our methodology has needed to “challenge the clouds”, as Sterman puts it, both by modeling a *prospective* systems problem and by attending to the importance of *heterogeneous* individual characteristics. In practice, firm-level consequences can easily arise out of individual effects in ways that evade the system dynamics modeling paradigm, as salient individual characteristics may not be randomly distributed among employees. Individual dynamics may not be a “wash” when rolled up to the firm level. Additionally, some individual characteristics require continuous visibility so that the human resource function can observe multiple workforce dimensions that are of strategic or regulatory interest, both in the simulation and in real life. Thus, the problem that we consider requires that we “do” both potentially-emergent issues and heterogeneous individuals, while still relying on the traditional strengths of the system modeling process.

Agent modeling techniques have arisen independently of the system dynamics approach in recent decades. While the latter focuses on endogenous system behavior and the control of it, usually at a macro level of analysis, agent modeling has focused on the emergence of behavior at a micro level, “bottom up”. An agent model “consists of a set of agents that encapsulate the behaviors of the various individuals that make up the system” (Parunak, Savit et al. 1998) to “... look at [unpredicted] global consequences of [their] individual or local interactions...” (Scholl 2001). To date, agent based models have been used principally as an aid to theory development or refinement, tracing the effects of interactions of agents (Kirman, 1992) in creating the aggregate properties of a system through their repeated interactions. (Lomi and Larsen 2001: 4)

Because both of the approaches seek to explain dynamic, nonlinear social behavior, albeit from different directions, scholars (e.g., Scholl 2001) have begun to call for cross-study and joint research between them. Rahmandad and Sterman recently observed that, despite the practical relevance of integrating the two approaches, the literature is limited. “Some studies find that the dynamics of [agent and system dynamics] models are similar; others find they differ” (Rahmandad and Sterman 2004: 1). Rahmandad contrasts the dynamics of agent and system

dynamics models in considering the well-known SEIR model of infectious disease contagion. As he reminds us, “the assessment of modeling methods is always contingent on the purpose of the model. The purpose determines what variables are of interest, what level of precision is required” (Rahmandad 2004: 20).

For his purposes, it sufficed for Rahmandad to contrast the performance of independent systems and agent models of contagion, and this comparative line of inquiry has its merits. In contrast, several scholars recently have combined the two approaches. Thus, Größler, Stotz and Schieritz (2004) used system dynamics models created in Vensim to provide internal decision making schemata for supply chain agents that they modeled in the RePast agent software. Their approach provides for a single distinguishing attribute among manufacturers modeled as agents, i.e., their respective manufacturing capacities. Akkermans created a system dynamics model that incorporated supplier and customer agents. These agents were more homogeneous than not: “[t]he only thing in which the ten actors ... differ is the degree in which they emphasize the short-term or the long-term performance of their counterparts...” in making new contracting decisions over time (Akkermans 2001: 4). Geerlings, Verbraeck, de Groot and Damen (2001) modeled manpower planning in the Royal Netherlands Navy with a object-oriented simulation, tracing employee state transitions, that combines elements of the ability management and opportunity management subsystems that we mentioned above. Their purposes and expectations for outcome measures were more limited than ours: simply the timely matching of supply (sailors’ available competencies) and demand (aggregate competencies needed to crew a vessel). Individuals in this model were not required to have any other distinguishing characteristic (e.g., gender) or attribute (e.g., motivation), and none was tracked. This limited expectation legitimates a modeling process in which only shortages and surpluses of homogeneous sailors are of interest to the client.

To date, therefore, models have represented individuals exhibiting a limited set of dimensions (e.g., degree of job experience or developed competency, focus on a single decision-making criterion, or infectiousness) that change state. Systems models analyze problems and the underlying structures that generate them over time. The nature of the problem, as seen through the eyes of the modeler’s client, becomes a central focus of the model design and elements, since it is the client whose learning and policy-making needs the model seeks to illuminate. As stated above, adverse results may arise within SHRM at the distinct levels of the firm, the individuals who work for it, or the human resource practices, policies and programs that translate between these organizational and individual levels. In the specific case addressed in this research, the vice president of human resources seeks not only to model aggregate-level behavior (change in nonunion employees’ turnover rate) but also to understand how such turnover may affect, as just one example, the firm’s minority female engineering population, and then in turn other system elements.

The adequacy of any of these approaches (an agent or a systems model alone, or a combined approach) thus is contingent on the purpose of the model, the variables that are of interest, the rhetorical needs of the inquiry, and the levels of precision that are required with respect to each. For this investigation, in light of the individual data elements identified in our hypotheses, we believe that a hybrid model will serve better than a “pure agent” or a “pure systems” one. Further, the client’s interest in continuously observing many heterogeneous characteristics of its employees extends agent descriptions beyond models reported to date. The

balance of this paper explains our model choice and explores some of the advantages (and problems) encountered in developing it.

Having said this, it must be noted that our current model relies principally on a systems approach. For our client's purposes, and in light of its data limitations, we trace certain demographic characteristics but by no means all important parameters of its individual employees as agents. Thus, we observe the consequences on its composite workforce picture of the primary, independent decisions of some more-senior employees to delay retirement and of some mid-career employees to depart "for greener pastures" as a result. In a sense, each employee – with important individual characteristics – interacts independently with his or her employer.

Data limitations preclude for the present a more-developed agent modeling engine in which other local interactions of employees leads to macro-level consequences. Here is an example. In the Context discussion above, we spoke of an opportunity management subsystem which defines roles within an organization and of an ability management subsystem which regulates how those roles are filled. As a firm's current employees alter their tenure decisions based upon such influences as pension reductions, they alter the supply of ability that the firm enjoys. This altered supply leads in turn to a range of firm-level consequences. Among these may be increased or diminished firm-level performance, the rearrangement of job opportunities (through job redesign), and the development of some form of competition among individuals as they seek to match their abilities to the newly-demanded opportunities. This matching activity may lead to promotions or reassignments, and may identify redundancies or skills obsolescence. Thus those initial, altered employee decisions may both reshape the pattern of opportunities that the firm provides and prompt competitive behavior among individuals seeking to reposition themselves relative to their coworkers to supply abilities within the new demand pattern of opportunities. Our hybrid approach sets the stage for including such agent behavior and interaction in later modeling.

Strengths and Limitations of the Two Approaches

Differences between agent and systems modeling, or the differences we believe are present, have raised issues of design, implementation and praxis that have governed our actions in the model as we develop it. In the following discussion, we review the strengths and limitations of the two disciplines that have led us to address our practical goals through a hybrid model. References in the paragraph captions to “Table” and “Row” refer to entries in summary Table 2, which follows this discussion.

Treatment of multi-valued mutually independent states [Table 2, row 1, below]. In representing individual employees, who each is, what each chooses and is able to do, and what the firm chooses and is able to do with and through them, we assume with Kirman that there is no such thing as a *generally* average or representative employee. As a consequence, our model design intends to examine the emergence of unexpected consequences from systems-level HR policy change due to agent-level differences in employee characteristics.

Individuals are complex organisms. Bar-Yam (1997) estimates the descriptive complexity of human behavior on the order of magnitude of 10^{10} bits of data, or more. Individual employees have many independent multi-valued properties, whether personal characteristics or firm-specific attributes, that are relevant in the management of a firm’s human resources. The personal characteristics include one’s age, gender, ethnic origins, family status, general educational achievement, and military or veteran status. The firm-relevant personal attributes include the individual’s degrees of each skill and knowledge, whether technical or interpersonal, that is relevant to his or her current or prospective responsibilities within the firm, the nature and personal history of such responsibilities, the nature and extent of compensation and benefits which the firm has decided or contracted to offer and which he or she may have elected to accept, his or her demonstrated performance at the firm, and so forth. Together, these are the basis for the definition of the many data fields that a comprehensive human resource information system makes available to a firm.

To model the current HR problem, we forecast the need to trace more than a dozen salient characteristics on an individual-by-individual basis, and others at an aggregate level across segments of the work force. The individual data include variables measured at the nominal (gender, and ethnic, union and military (veteran) statuses), ordinal (performance rating, family status, organizational level) and interval (age, credited service levels, employer and employee contributions for each benefit offering, amounts of various compensation elements provided) levels. If there were only 10 individual-by-individual properties (and there are more), with, say, 5 values for each (though the real number of possible values may be two, or two hundred), and if all combinations were to be differentiated using stocks (or cells in stock arrays), then over 9 million stocks (or cells) would be required. All this stock structure would be needed at *each* step in the model flow, even for details that were not immediately necessary, lest essential information be lost to future steps.

In order to be useful to its audience, a model must simplify (Serman, 2000). This norm properly cautions us in the face of such a munificence of relevant employee attributes. We draw two conclusions. First, a system dynamics modeling tool (such as Vensim) must use other, more efficient data structures, each with a small set of properties or characteristics, to represent an

employee population. Second, stocks and flows alone will not serve if the model needs to trace such micro detail in order to serve fully the client's purposes.

In our hybrid HR model currently under development, a two-dimensional array (implemented as a stock, but semantically very different from one) is used as an employee database, supporting agent modeling processes. Cells in this two-dimensional array are coded representations of employee identifiers and data field values, not quantities (as one would find in a subscripted stock). Such a 'database' representation is very natural within our modeling tool HAM (Heterogeneous Agent Model) and it has a great advantage in simplicity. This is because – with the addition of each new employee characteristic – only a single new subscript value is required, rather than one new dimension plus as many new subscript values as the characteristic has possible values. Of course, a database structure can be copied into a traditional system dynamics environment as a subscripted stock, once all database values are recoded as integers – and we do this, for computational reasons, in our model – but the result is semantically unusual and inconvenient to update, since previous values must be backed out and all unchanged values must be reassigned with each iteration.

Instead of this cumbersome mechanism, therefore, we are developing and will employ HAM, a rule-based, transformational platform to trace heterogeneous agents having a number of independent (and often dynamic) properties over time. It is described in more detail in the Control Component discussion below. This opportunity to handle multi-valued independent states under the agent modeling paradigm without extreme model size and complexity motivates us to include some agent components in modeling the emerging workforce problems which the client anticipates.

Treatment of period-to-date financial data accumulators [Table 2, row 2]. Firms operating in financially developed economies in the world today use an annual budgeting cycle, as well as annual, quarterly, and monthly forecasts, for the control of budgeted resources. The process is nearly universal, and – during periods of strategic stability for the organization – can be carried out in a very principled and mechanical way. Thus, this review and budgeting cycle is a very good candidate for inclusion *inside* an HR systems model.

However, even though period-to-date financial accumulators fit well in the computational context of stocks and flows (and hence in the classic systems paradigm), they are not, in fact, classic stocks, because they are zeroed out at accounting intervals. A workaround is generally possible, even in an environment like Vensim. Typically a separate variable is created there for each period-to-date financial accumulator, which – at the end of an accounting period – receives the contents of the stock and creates a reversed, canceling component to the 'flow' into that stock, which, in effect 'zeros' the accumulator.

Such a financial modeling construct is a system dynamics modeling idiom, but it is not really in the spirit of that paradigm. From a practical point of view, it has the disadvantage that the zeroing of the accumulator occurs after the first assignments of a time period, and before the last. Where possible, therefore, we have set up period-to-date accumulators in the HAM segment of our HR model, with the quantities based on preliminary agent transformations defined in HAM, and then we furnish those values up-front, as data, to the Vensim-based portion of the model.

Thus the opportunity for including more facile and realistic financial accumulators afforded by the agent paradigm motivates us to employ a model with at least some AB

components in designing an HR-workforce study. We note that this issue is somewhat less motivating than the others we tabulate here; a workaround commonly used in systems modeling environments for period-to-date accumulators (see above) gives fair results.

Support for sorting operations based on multiple characteristics [Table 2, row 3]. Many HR-managed processes (like presenting offers to employment prospects in an order determined by a suitability heuristic) depend on sorting and ordinal positioning. An ordinal viewpoint is not natural to the differential equation paradigm, except where ordering resides in a *single* sequence of states, or in a queue. The HR commonplace, whereby individuals are ranked according to numerous different ranking heuristics, and motivated according to numerous influences, does not really have a place in the systems conceptual framework.

In implementing our HR model, we have responded to this limitation in two different and complementary ways. Where phenomena are being modeled in which stochastic variation is involved (perhaps through the right of individual refusal), and in which there is no fixed ‘floor’ or ‘ceiling’ in dollars or headcount, the entire process can often be implemented probabilistically in the agent paradigm: the same characteristics that were to have created an ordering will now affect probabilities. A model for participation in an elite training program might go along those lines.

Likewise, where the management of a given process is explicitly based on orderings of employees or prospects according to one or multiple criteria, or when a certain quantity of resource must not be exceeded in dealing with elements of varying cost, actual sorting is carried out as part of an agent process. A model for making job offers to prospects based on test scores, with or without any preference in “furthering diversity” within a workforce, might operate along these lines.

So, the opportunities afforded for unconstrained rank ordering, sorting, and comparison under the agent modeling paradigm motivate us to employ a model with at least some agent modeling components in designing an HR-workforce study.

Support for test-terminated iterative procedures within a single time period [Table 2, row 4]. In many HR management situations (such as some of those described immediately above), managers proceed in an iterative way, until some predetermined condition or deadline is reached, or until designated resources are expended, or until some goal is met. Thus, for example, a department or a firm may *iterate* through a list of heterogeneous prospects, each available (if at all) at a different price, and each more or less likely to accept an offer; and continue this iteration until a certain amount has been expended or a certain headcount condition has been reached, and then stop.

With certain simplifying assumptions, this is a classic situation for stock and flow modeling. However, without those simplifying assumptions (mostly assumptions of uniformity and uniform distribution of characteristics among individuals), the problem above needs to be formulated in an iterative way, as something like: “do offer *until* a quota is met or earmarked resources are exhausted.” Since it is the client’s and the modelers’ wish to explore effects in workforce evolution and management that depend on the *non*-uniformity of individual employees and prospects, this formulation is the one we must address.

‘Do-until’ is a familiar programming control structure, but it is not present in the systems modeling paradigm, since a differential-equation-based model, with an infrastructure that

assumes continuous change, cannot offer application programmers an iterative clock shorter than a single dt time period. Thus, in systems modeling, iteration is pervasive and complex, but the only clock it exposes to the application for iterative control is the model clock.

In the current HR modeling project, this limitation is really more an implementation issue than a methodological one. We find that when we embed agent modeling structures in the Vensim environment we cannot implement test-terminated iterative procedures to operate on them, because Vensim – a differential equation-based platform – does not support ordered sequences of executable statements at the application programming level. Hence, the opportunities afforded for iterative control structures under the agent model paradigm motivate us to employ a model with at least some agent modeling components in designing this workforce study.

Sensitivity Testing [Table 2, row 5]. Whether or not workforce modeling is in place, HR and corporate management must determine, at an aggregate level, which program or funding changes will be small in their effects, and which large, in order to develop strategic initiatives. This is a general need.

A systems model does this very well at an aggregate level when program elements and responses can be modeled as stocks and flows: results are incisive and easily understood. Within an agent model process, the best one can do is make relatively large changes in one parameter or another, and then run the model many, many times. This difference – favoring the system dynamics approach as an HR modeling tool – strongly motivates us to employ a hybrid model, rather than a pure agent model. Of course, caution is required: any external agent modeling processes that normally perturb the Vensim implementation must be ‘frozen’ for the nonce. That is, in a sensitivity-testing context, agent modeling computations offered to Vensim as data *must* be held constant from run to run, and our agent modeling platform HAM *must* not interpose dynamic computations between the beginning and the end of a run. Since the systems model can be counted on to observe the effects of very small changes during sensitivity runs, agent modeling features can temporarily be ‘frozen’ without much loss of salience.

So, the opportunities afforded for sensitivity testing under system dynamics applications like Vensim motivate us to employ a model with at least some system modeling components in designing an HR-workforce study.

HR processes which are managed as flows [Table 2, row 6]. For representing processes which a firm perceives and manages as *flows* – such as the filling of vacant positions where greater individual homogeneity is acceptable – anything other than a flow model will falsify the real nature of the process and the way it is governed. The system paradigm models these processes quite well, and there are many examples. As we stated above, Sterman’s textbook formulation (2000: 758) of a firm’s workforce vacancy-filling problem can be used with slight modification in conjunction with our retirement plan model. In it, a stock of vacancies together with an average time factor needed to fill them combine to create a flow of hiring into the second principal stock, that of active employees. Qualitative research that we have conducted suggests that a model segment like the one shown above well presents the view of human resource practitioners who are responsible for hiring. It is natural within the systems paradigm, or within a hybrid model where a systems modeling component passes on a ‘hiring instruction’ to an agent modeling component. From the perspective of American law regulating collective bargaining, a firm’s union-represented employees often present a relatively

homogeneous employee population. Greater uniformity of relevant attributes arises from pay scales, pension benefits and other total rewards and ability management provisions that are based by bargaining agreement on seniority, not performance or employee choice. Thus, a stocks and flows model applies especially well in our view for union-represented employees at a relatively high level of aggregation acceptable to the client.

If instead we try to realize such a mechanism in a pure agent modeling environment, we are obliged either to use stochastic processes that will approximate the desired flows, or to implement software monitors that shut down agent modeling mechanisms when completion conditions are met. In our view, neither of these possible means is completely representative of what really occurs: that is, neither really represents the element of direct control based on quantity as cleanly as the system model formulations.

Table 2 summarizes the foregoing discussion. The first four rows of the table reflect reasons why agent characteristics are desirable in the model, while the last two reflect reasons why system dynamics model characteristics are important.

Table 2: System/Agent Modeling Divergences That Motivate the Use of a Hybrid Model for the Pension Problem

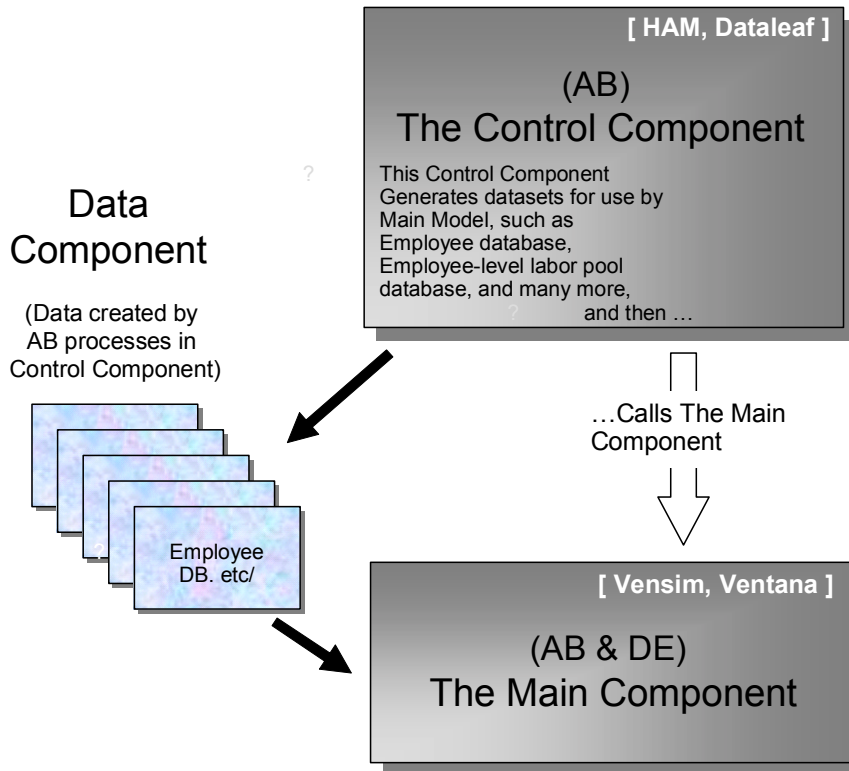
System Characteristic(s)	Agent Characteristic(s)	Implementation Issue(s) for this HR Simulation	Methodological Issue(s) for this HR Simulation
1. Complexity grows exponentially with the number of relevant independent properties (or sequences of states) of individuals	Individual human complexity is quite large, and includes both personal and firm-relative characteristics.	Procedural complexity for systems models can be substantial, since employee properties of interest must be carried through multiple subscribed stocks.	To meet the client's needs, modeling is necessary at multiple levels of analysis. For modeling here, individuals have >12 independent multi-valued properties.
2. Processes requiring period-to-date eliminators are not natural to the paradigm. Elimination occurs over time, through flows draining stocks.	Period-to-date accumulators are natural because accumulators can be zeroed at any time without constraint.	Workarounds are not in the spirit of the differential equation-based paradigm.	In HR, period-to-date accumulators are an essential component of much planning, including performance to fiscal year budgets.
3. Sorting individuals by multiple properties or characteristics is extraordinary.	Tracing of multiple properties of individuals and their combinations is facile for modeler.	Sorting by multiple criteria is generally unsupported in system dynamics.	The ordering of employees according to complex heuristics is an HR and client commonplace.

System Characteristic(s)	Agent Characteristic(s)	Implementation Issue(s) for this HR Simulation	Methodological Issue(s) for this HR Simulation
4. Test-terminated iteration (e.g., “hire heterogeneous prospects to fill a billets quota”) is not supported within a single time period.	Iterative procedures are supported, within or across time periods.	Only limited workarounds in a differential equation environment such as Vensim are possible.	Iterative processes involved in satisfying cost, schedule and quality targets are fundamental to and widespread in HR management.
5. Sensitivity testing is efficient and performance is good at an aggregate level of analysis.	Sensitivity testing is relatively slow and cumbersome.	All agent-based input to a system dynamics model should be held stable in successive runs for sensitivity testing	HR models often seek to optimize employee placement and use of firm resources.
6. Processes that are managed as flows (e.g., filling job openings with homogeneous candidates) are represented naturally.	Processes that the firm manages as flows are clumsy in the agent paradigm. Stochastic processes to approximate the desired flows, or software monitors to shut down agent modeling, are needed.	Using a system dynamics mechanism to pass directives to an agent mechanism for execution is a viable implementation approach.	HR processes such as hiring are often <i>managed</i> as flows within a firm.

Agent and Systems Features in the Model under Development

In view of the issues discussed above, we employ a hybrid model based in Vensim but also employing agent model functionality and individual characteristics as required to address the client's concerns. This hybrid model is illustrated in Figure 2.

Figure 2: Hybrid Agent/Systems Model for HR Workforce Project (Simplified)



The hybrid system consists of three main components: Control, Data and Main. The nature and role of each is described below:

- The top level Control Component, written using the HAM platform, performs a number of tasks in the hybrid model, which include (1) calling the Main Component, (2) executing the necessary row-column transformations on the client's real HR data, thereby preparing this data to be passed to the Data Component for use in the Main Component, and (3) creating statistical profiles for prospective new hires. Heuristics implemented in the Control Component are agent heuristics. HAM is a rule-based, transformational platform intended to trace heterogeneous agents having a number of independent (and often dynamic) properties over time. It is designed to import such real-world data from the client site into a system simulation. Currently under development, HAM has a rule taxonomy and order-of-evaluation conventions that are specifically tailored to workforce modeling.

- The Main Component is a detailed systems model that is written in Vensim and implements the heart of the study. It contains both system and agent heuristics, although Vensim is chiefly designed and intended to support the former.
- A Data Component, consisting of about forty external tables which contain simulation setup data, client budgetary data, HR program parameters and reference modes, regional economic trend data industry norms and so forth, all realized as external Microsoft Excel spreadsheets that are read directly by the other two components. In addition to these Excel spreadsheets, the data component also includes a copy of the client's HR database augmented and formatted by the Control Component and read by the Main Component.

Paradigm Boundary Issues in the Hybrid Model

We next examine the ways in which differences between the systems and agent paradigms have governed our actions at points of tangency between their respective components. Once again, references in the paragraph captions to “Table” and “Row” refer to entries in the summary table that follows the discussion.

Typical significance of numbers in system and agent models [Table 3, row 1, below]. In systems models, numbers are typically aggregate quantities or rates, such as “number of engineers” or “rate per month of voluntary terminations of employment”, or auxiliary factors. System dynamics modeling norms (Sterman, 2000: 866) and the Vensim software impose a dimensional consistency test to help to ensure the appropriate measurement of parameters. In agent models, numbers (or combinations of numbers) instead typically stand for properties of or relations between agents, such as “currently active full-time female engineer” or “supervisor of (other specified individuals or positions)” (Carley and Krackhardt 1998). Numbers also may serve as pointers or indices within an agent database, as in “next-most-qualified candidate for job offer or promotion”.

Here is an example of the complexity inherent in the interface of the system and agent paradigms. In our HR-workforce model, there is a flow-oriented component, similar to the modified Sterman model discussed above, which models the process of filling non-union positions. Each period it delivers a quantity, “Junior Hiring,” which specifies how many individuals should be added to the employee population from the relevant labor pool. This value, which is a simple quantity, doesn't flow anywhere. Instead, it enters a subsystem of agent components, then interacts with an indicator pointing to the first available prospect in a labor pool array, directing that certain cell values change from 0 to 1 in an employee database array which represents each individual as an entire row. All this occurs within the Main Component, which is a systems model written in the Vensim environment. As it does in all transactions involving external databases, Vensim therefore scrutinizes it for unit-of-measure and “model check” errors and issues warnings, even though the paradigm boundary is an appropriate one and the joint model runs as intended. We conjecture that this will be typical of such paradigm boundaries in any but the most tactfully assembled hybrids, at least in any development environment where unit checking is an important discipline for the modeler.

In our development effort, we treat this ‘symptom’ as a significant caveat. For example, we refrain from sensitivity testing centered on variables that participate in these transitions

between a Vensim model component where numbers are focused on rate and quantity and an agent component where numbers rarely have those meanings.

Typical semantics of arrays in system and agent models [Table 3, row 2]. Within a systems model, the overall array of values typically consists of a charily-increased number of stocks plus the rates of change and other parameters that affect them. (Sterman, 2000: 217). In modeling heterogeneous agents, however, the focus on heterogeneity generally makes relevant to the human resources practitioner a larger array of variables, ones that are not as interdependent. We have argued that the structure presented in Figure 1 above is able to describe total rewards programs in addition to our focal program, pension benefits. Thus, we can present in the next paragraph an example dealing with employer-provided medical benefits that ties the element of Cost Control Efforts in Figure 1 to a larger array of heterogeneous employee variables.

Generically, such cost control efforts fit one of three types: cost shifting, cost sharing and cost reduction. Cost shifting efforts systematically transfer a portion of gross medical coverage cost to employees through increases in “employee contributions” taken from their paychecks. These amounts often increase with the size of an employee’s “family”; that is, whether he or she elects single, two-person “couple”, or family coverage. Thus the demographic variable “family status” affects both the gross cost and cost shifting amounts for medical coverage, while the employee’s pay level affects the economic impact he or she personally feels, and the resulting actions he or she may individually take. Cost sharing efforts instead transfer a portion of gross cost -- through episode-by-episode out-of-pocket co-pays and deductibles -- for consumption of health care by covered individuals. Now, in addition to the individual variables already discussed, the ages of the worker and his or her dependents become relevant, as younger individuals may be actuarially expected to seek medical treatment less often than older ones, and so experience less incidence of cost-sharing. Employer cost reduction efforts attempt to shift gross cost back to the external environment through, for example, contracting for discounted rates from physician networks, local hospitals and prescription drug outlets. Our array of relevant heterogeneous agent characteristics now grows further with considerations of covered individuals’ residences (related to proximity to contracted and excluded health care suppliers) and network relationships (continuity of contract ties to providers like their “family doctor” and availability of nearest or “best” hospitals). Thus, macro- or systems-level cost control efforts will have varied impact through such heterogeneous variables on individual employees and their dependents, and their varied micro-level responses (e.g., changes in benefit elections, disgruntlement, quits) will feed back to workforce availability, motivation and productivity.

Turning back to our pension model under development, there is yet another kind of array: the database. Databases are a natural feature of a HAM model, which makes extensive use of dynamically-generated SQL. But even in the Vensim-based portion of our hybrid model, where databases are a decidedly *unnatural* usage, we are obliged to implement an employee-and-labor-pool database in an array. In this array, each person corresponds to a row, so the row subscript designates the employee or prospect. The second subscript is really a container for non-stock, non-flow properties like race, gender, department, date of hire, employee status, etc.

Since Vensim forbids the modeler to add or remove rows from such arrays dynamically (and if such are viewed as arrays of stocks, as they are under the systems paradigm, it should not *be* possible!), individuals must be moved from one state to another (e.g., from prospect to employee, from employee to retiree) by altering the coded contents of a property cell. We believe that differences in the semantics of arrays in the two approaches impose a significant

caveat on the development of hybrid systems. For example, in the current project, we refrain from sensitivity testing centered on ‘database’ arrays, whose uses and semantics vary sharply from what is implicit in the Vensim development environment, because the numbers in those arrays aren’t really quantities.

Resistance to exogenous perturbation during a run [Table 3, row 3]. To a degree, systems models view past and future time periods *synoptically* as they seek to describe behavior produced endogenously within a system (Sterman, 2000: 95). To accomplish this, external perturbations of a system during a given run must be prevented or minimized, and Vensim is carefully written to prevent them.

This places rather strict limits on one potential mode of collaboration between separate agent and systems platforms (such as our control and main components), namely, for the agent platform to *review and update* the system platform dynamically after each *dt* time period. Clearly, indiscriminate updating of this sort would violate system modeling assumptions (see Sterman, 2000: 206-8), and compromise some (though by no means all) of the more valuable capabilities of Vensim. From an implementation point of view, we have had to secure a level of period-by-period collaboration without excessive compromise through the “game” functionality offered by Vensim. Using such an arrangement, the HAM control component ‘games’ the main systems component, supplying heterogeneous agent modeling services at gaming intervals.

In the interest of conserving space, we do not update Figure 2 here. An updated version of our hybrid model would replace “...Calls the Main Component” in Figure 2 with a closed loop of arrows marked “Repeated Gaming Calls”. In a simulation run of this sort, the HAM system interacts with the “play” (i.e., the Vensim run) exactly as a human game player intervenes when it is his or her “turn” in an ongoing strategy game. Of course this mode of collaboration precludes the use we can make of certain system dynamics modeling capabilities, like delaying, smoothing, forecasting, and sensitivity testing, and it exacts a significant performance penalty as well, so we will limit this mode of collaboration to instances where it is clearly necessary. This mechanism does importantly provide a kind of disciplined collaboration between two very different paradigms.

Level-of-analysis handoffs between system and agent components [Table 3, row 4]. In our hybrid HR workforce model, developed to model unintended effects of pension changes, we have noticed an asymmetry. Aggregate hiring and job-loss mechanisms tend to be defined in systems terms, specific retirements and other departures from the firm’s employment need to be defined in the agent portion of the model, and promotion and succession elements are somewhat a “mixed bag”.

Although a complete explanation is beyond our scope here, this seems a natural result of the general theoretic model described near the outset of this paper. The opportunity management subsystem of HR activity enumerates job requirements and stores “open” positions until these are filled through the sourcing, recruiting, and qualifying activities that occur in “early” stages in the functioning of the ability management subsystem. The stock of vacancies (i.e., requirements less current supply) is reduced as actual persons (agents), who more or less satisfy the demands of the posts for specific skills, knowledge and abilities, individually choose to accept offers and assignments. We conjecture that this may be a general characteristic of hybrid human resource management models. A particular job opening does not require that candidates meet an entire set of agent-level characteristics or attributes, although satisfaction of certain criteria may be

required. Thus an employer may wish to hire or promote someone who is a minority candidate (to “promote diversity”), or who has a status (citizenship, for a post that requires security clearance) or a qualification (current nuclear engineering certification, for power plant employment), and so forth, but in many other respects an agent that doesn’t exist yet can’t be modeled as an agent. It is only when he or she begins employment that a firm may come to know characteristics like date of birth, family status and personal savings propensity, and to assign attributes like pay rates and job assignments, that will come to affect an individual’s career and (for purposes of our problem) his or her retirement decision at some future date.

In our model of emerging unintended consequences of pension reductions or other strategic human resource management interventions, this view that system processes adjust job opportunities while agent processes adjust to these is central. A system model can supply information, based upon mean employee responses to workplace stimuli like employee benefits changes, of changing aggregate employment numbers. It conveys modest information at best about how such changes may affect a host of other characteristics (demographics) and attributes (motivations and behaviors) of the workforce that are nonetheless very relevant to a firm’s business. Traditional agent models permit the modeler to trace the emergence of workforce responses at the individual level. These agents usually vary across a limited number of state conditions without reflecting their heterogeneity across other important characteristics and attributes. Our research suggests that a hybrid approach incorporating heterogeneous agents may offer distinctive information that neither a “pure” system dynamics nor a traditional agent modeling approach can supply.

Table 3 summarizes the foregoing discussion.

Table 3: Characteristics Affecting the Points of Tangency between Paradigms

Systems Model Characteristic(s)	Agent Model Characteristic(s)	Implementation Issue(s) for this Project	Methodological Issue(s) for this Project
1. Numbers are quantities, rates at which quantities change, or auxiliary factors.	Numbers are not usually rates or quantities; they may be nominal- or ordinal-level properties, pointers, indices, tuples, probabilities, etc.	None. However, we expect unit-of-measure and model check exceptions to arise at the boundaries.	Avoid sensitivity testing centered on boundary transition elements.
2. Multi-dimensional arrays typically are collections of related stocks.	A multi-dimensional array typically represents either a set of relationships between individuals, or a database.	Unit-of-measure exceptions are likely if agent and systems components exist in the same modeling platform.	Boundary variables should not be made the object of sensitivity runs.
3. The systems model in Vensim protects its validity by strictly limiting changes made during a model run.	The agent model simply runs. Extending agent modeling beyond emergence examination, perturbations warranted by real-world processes are permitted.	‘Gaming’ must be used as a mechanism for agent components to insert results in the system simulation at every time period.	Gaming must be used in such a way as not to violate systems modeling assumptions.
4. Individuals entering the model are enumerated by flows.	Individuals exiting the model are selected through stochastic processes.	Open positions are specified within stocks, then filled by agents with individual characteristics; reversed process for terminations.	Agent and system components populate and depopulate the model in complementary ways.

Discussion: Some Useful Functionalities for Hybrid Models

We recognize that future developments in modeling software will be complex, as the field of possible applications will be increasingly so. However, we do believe it is now possible to identify certain specific features that, if realized in modeling software platforms in the future, may have significant utility for hybrid modeling projects across a much wider field of application than ours.

In conclusion, we have listed a few of these features below. Some we have introduced experimentally as part of the functionality of the proprietary agent modeling platform, HAM. Some, for clarity, we describe as they might appear if implemented – in the most rudimentary possible form – in a standard system modeling platform such as Vensim. Again, a discussion and then a summary table follow.

Accumulators [Table 4, row 1, below]. In system dynamics models, “stocks” have certain familiar characteristics: their contents persist from one cycle to the next, and – except initially – their values are not replaced by a modeling rule or expression; they are only incremented and decremented. This is entirely suitable for stocks that reflect a uniform physical flow, or a business quantity like the number of employees in a given department.

In some real world situations, especially in business processes, there is often seen a slightly different kind of stock, one that has a crucial role in guiding the business: the period-limited accumulator. Usually it is a quarter-to-date or year-to-date financial total. In principle, period-limited accumulators are not restricted to problems such as ours. (One might imagine a stock that represents a population of adult insects of some species in some particular location, which goes to zero each winter not through an exhaustion of its food supply but instead due to a hard freeze.) Still, we believe the greatest potential benefit of these accumulators lies in the system dynamics component of a system-agent hybrid model simulating a managed business process.

In the context of one common system dynamics platform, Vensim, such an accumulator would probably appear as a third, optional argument to the INTEG() function. If present, it would specify a time interval at which the stock that it governed would be re-initialized. Thus a value of 12 would create a year-to-date accumulator, and 3, a quarter-to-date accumulator, assuming a monthly time step.

‘Table’ variables [Table 4, row 2]. When agents have a large number of independent properties – in the case of employees, these are properties like date-of-birth, gender, department, job, etc. – one relatively simple data structure that can capture them and their properties is a multidimensional array, understood as a database “table.” For example, a fairly simple “table” of employees would be a two-dimensional array with a row for each employee, and a column for each property – such as department, date-of-birth, union status, incentive compensation eligibility, or some other.

Tables of this sort are quite natural under the agent modeling paradigm, but less so in system dynamics. In the system dynamics component of our current analysis we implement an array of employees as a doubly-subscripted stock, or level variable, as our model requires that some agent information be accessible in that component; but in many ways this is an unsatisfactory compromise. Among the disadvantages are:

- When it is necessary to assign a new value for a particular property of a particular employee, the value must be backed out and added in again. This multiplies the number of formulas alarmingly if we allow this component of the model to become too complex
- All cell values must be coded as numbers. This makes the entire data structure very opaque inside Vensim; this is uncongenial, and can lead to errors. In our project, we use the agent modeling HAM platform to control the system dynamics component. This means that HAM also can be employed to keep track of the integers that we use to code everything from department to date-of-birth.
- Rows (that is, employees) cannot be added or deleted dynamically. This is a reasonable limitation in the system dynamics paradigm, but it means that we need to dimension our array with enough rows to handle new hires, and to manage the recruitment of new prospects with cumbersome pointer variables that are implemented as levels, although semantically that is not what they are.
- Unit-of-measure discipline suffers in the vicinity of such an array in environments like Vensim that support such discipline. For example, in the context of hiring and termination, our employee array can hardly avoid appearing to the system modeling software as an array of people; yet none of the quantities in any of the cells are people at all.

The “table” variable type provides a mechanism that would facilitate designs like ours, incorporating data needed to support agent modeling processes in a system dynamics environment. To give a concrete example, such a “table” variable in a Vensim-like context might resemble a subscripted stock, except in the following ways:

- The first subscript (and the first subscript only) would be an integer, not a symbol; ‘rows’ could be added dynamically.
- Data type and unit-of-measure would be defined at the level of the last subscript governing a particular cell, not at the level of the entire table. String, real, and integer data would be permitted.
- Value assignment to cells, whether individually or in groups, would be through direct assignment, not through incrementing and decrementing. Cells that were not assigned values in a given time step would retain their previous values, and no error messages would be generated.

Test-terminated iteration [Table 4, row 3]. It is idiomatic, in business processes, to manage some short-term initiative or task by doing something “while” or “until” a particular condition pertains. In the realm of Human Resources, for example, one often hires *until* one’s department’s job requisitions are filled; or one offers training billets *while* one has something left in the training budget. At first sight these activities seem to resemble ordinary flows subject to carrying capacities, but in fact they are somewhat different. A firm does not make offers to job prospects uniformly, or at random, but according to various more or less complex evaluation metrics; and not all prospects accept offers, nor are all of those who do so paid alike.

In the business world, the process is still fairly straightforward. Management makes offers to prospects according to whatever legally acceptable heuristic it chooses, negotiates compensation, and (sometimes) hires. This continues until all job slots are filled. We believe

the process should be managed in modeling as it is in life. However, with system dynamics platforms, test-terminated iterative processes may not be supported for multiple iterations occurring within a single time step. Of course, it is generally possible to *call* internally iterative procedures from inside a system dynamics modeling environment. In Vensim, such a procedure can be realized as a DLL. Unfortunately, procedures of this sort are often limited as in Vensim to returning a single value, which provides insufficient information to update the model completely. It is as if one knew exactly how many training dollars one had expended, but not for which courses, nor for which employees, nor for how many employees.

If one were to incorporate a mechanism for test-terminated iteration in a system dynamics platform, it might be restricted to manipulating only “Table” variables as described above, and other variables local to the iterative procedure itself.

Explicit handling of uncertainty [Table 4, row 4]. In agent-based models of the type we describe, uncertainty is pervasive and affects tens of variables. This can be true to such a degree that stochastic mechanisms make it difficult to determine *where* in all the intermediate computations uncertainty plays a significant part.

We believe that a useful complementary approach is to *propagate explicit measures of uncertainty* throughout the model, rather as it is done in expert systems and decision support tools. In HAM, we have provided for two alternate running modes for any model with “Monte Carlo” components. These modes are *stochastic*, in which values are randomly selected according to some distribution, and *uncertainty-propagating*, in which an Uncertain Number – a data object with a number component, a distribution component, and a variability component – is used. Arithmetic operations on Uncertain Numbers are defined so that they always yield Uncertain Numbers. A global switch determines whether HAM will operate in stochastic or uncertainty-propagating mode for a given run.

A strategy whereby measures of uncertainty are explicitly propagated through the model has limitations of course. For example, stochastic processes that determine whether or not an agent enters the model cannot be replaced – even temporarily – by any other mechanism. Thus, when operating in “uncertainty mode,” our HAM platform leaves all such stochastic processes in place, although it does “freeze” them so they will give the same results in successive runs. Likewise, any system that propagates uncertainty measures requires precise mechanisms to specify dependence and independence of model components..

Prospective Contribution of this Hybrid Right Workforce Model

As outlined above, the nature of the problem to be simulated as seen through the eyes of the client is central to model design, since it is the client whose learning and policy-making needs the model seeks to illuminate. To date, workforce management models have depicted individuals as exhibiting a very limited set of dimensions. In the situation at our data site, which is actually to be simulated in the next stage of this research, the client seeks not just to have aggregate-level behavior of the workforce “rolled up” but also to understand how such turnover may unexpectedly affect its distribution of important demographic characteristics and individual competencies. We intend to report our findings in a subsequent paper.

Table 4: Useful Functionalities for Hybrid Models

Functionality	Possible Benefits and Examples	Realm of Applicability
1. Accumulators	Representation of re-initialized accumulators such as fiscal year spending. Example: Annual Closing, Financial Books (Figure 1).	System dynamics: business financial modeling:
2. 'Table' variables	Compact encapsulation of agents with multiple independent properties in a system dynamics environment. Example: Table of Individual Needs, Characteristics and Attributes (Figure 1).	System dynamics with agent modeling elements: not area-specific.
3. Test-terminated iteration	Accurate representation of various business process directives in an environment where agents are heterogeneous. Example: Execute prioritized job offers to fill open requisitions	System and agent platforms: "do-until" processes, especially managed business processes.
4. Explicit handling of uncertainty	Propagation of uncertainty information throughout the model, as an alternative to stochastic mechanisms. Example: HAM functionality.	Agent modeling: not area-specific.

Bibliography

- Akkermans, H. A. (2001). Emergent supply networks: System dynamics simulation of adaptive supply agents. 34th Hawaiian International Conference on Systems Science, Wailea, HI.
- Bar-Yam, Y. (1997). Dynamics of complex systems. Reading, MA, Addison-Wesley.
- Becker, B. E. and M. A. Huselid (1998). High performance work systems and firm performance: a synthesis of research and managerial implications. Research in personnel and human resource management. G. R. Ferris. vol. **16**.
- Becker, B. E., M. A. Huselid, et al. (1997). "HR as a source of shareholder value: research and recommendations." Human resource management **36** (1): 39-47.
- Becker, B. E., M. A. Huselid, et al. (2001). The HR scorecard: linking people, strategy, and performance. Boston, MA, Harvard Business School Press.
- Campbell, J. P., M. D. Dunnette, et al. (1970). Managerial behavior, performance and effectiveness. New York, McGraw-Hill.
- Carley, K. M. and D. Krackhardt (1998). A PCANS model of structure in organizations. International Symposium on Command and Control Research and Technology, Monterrey, CA.
- Coyle, R. G. (1996). System dynamics modeling: a practical approach. London, Chapman & Hall.
- Dalkey, N. and O. Helmer (1963). "An experimental application of the Delphi method to the use of experts." Management science.
- Dreher, G. F. and T. F. Dougherty (2002). Human resource strategy: a behavioral perspective for the general manager. New York, McGraw-Hill/Irwin.
- Fay, C. H., P. S. Hempel, et al. (1997). Costing human resource initiatives, School of Management and Labor Relations, Rutgers University.
- Fitz-enz, J. (1984). How to measure human resources management. New York, McGraw-Hill Book Company.
- Forrester, J. W. (2nd ed. 1973). World Dynamics. Cambridge MA, Productivity Press.
- Geerlings, W. S. J., A. Verbraeck, et al. (2001). Manpower forecasting: a discrete-event object-oriented simulation approach. 34th Hawaiian International Conference on Systems Science, Wailea, HI.
- Geisel, J. (2004). Hybrid pension plan growth has stalled. Business insurance. **38**: 3.
- Geweke, J. (1985). "Macroeconomic modeling and the theory of the representative agents." American Economics Association Papers and Proceedings **75** (2): 206-210.
- Größler, A., M. Stotz, et al. (2004). A software interface between system dynamics and agent-based simulations: Linking Vensim and RePast. System Dynamics Society, Oxford, UK.
- Holland, J. H. (1998). Emergence, from order to chaos. Reading, MA, Addison-Wesley.
- Johnston, D. C. (2004). I.B.M. Agrees to Settle Part of Giant Pension Case. The New York Times. New York.
- Kirman, A. (1992). "Whom or what does the representative individual represent?" Journal of economic perspectives **6** (2): 117-136.
- Kupiec, P. H. and S. A. Sharpe (1991). "Animal spirits, margin requirements, and stock price volatility." Journal of Finance **46** (2): 717-731.
- Lomi, A. and E. R. Larsen, Eds. (2001). Dynamics of organizations: computational modeling and organization theories. Menlo Park, CA, AAI Press.

- Meadows, D. H., D. L. Meadows, J. Randers, and W. W. Behrens III. (1972). The Limits to Growth: A Report for the Club of Rome's Project on the Predicament of Mankind. New York: Universe Books.
- Parunak, H. V. D., R. Savit, et al. (1998). Agent-based modeling vs. equation-based modeling: a case study and users' guide. Multi-agent systems and Agent-based simulation (Multi-agent systems and Agent-based simulation conference, 1998).
- Rahmandad, H. and J. D. Sterman (2004). Heterogeneity and network structure in the dynamics of contagion: Comparing agent-based and differential equation models. System Dynamics Society, Oxford, UK.
- Repenning, N., P. Gonçalves, et al. (2001). "Past the tipping point: The persistence of fire-fighting in product development." California Management Review **43** (4): 44-63.
- Scholl, H. J. (2001). Agent-based and System Dynamics Modeling: a call for cross-study and joint research. 34th International Conference on Systems Science, Wailea, HI, Institute of Electrical and Electronics Engineers, Inc.
- Senge, P. M. (1990). The fifth discipline: the art and practice of the learning organization. New York, Doubleday.
- Senge, P. M., C. Roberts, et al. (1994). The fifth discipline fieldbook: strategies and tools for building a learning organization. New York, Doubleday.
- Sterman, J. D. (2000). Business Dynamics: systems thinking and modeling for a complex world. Boston, McGraw-Hill Higher Education.
- Ventana Systems (2003). Vensim: the Ventana simulation environment. Harvard, MA, Ventana Systems, Inc. Release 5.2.
- Wright, P. M. and S. A. Snell (1991). "Toward an integrative view of strategic human resource management." Human Resource Management Review **1**: 203-25.
- Zink, A. (1997). Sort of a 401(k), sort of a pension. Crain's Chicago Business. **20**: R3.