Chapter 10
Disappointing Outcomes: Can Implementation Modeling Help?

I. David Wheat and Eugene Bardach

Abstract This paper addresses questions about modeling the implementation requirements of a public policy proposal. Can modeling provide advance warning of problematic implementation requirements inherent in the design of a policy idea? Going further, can it suggest feasible redesign options to improve the chances for desired outcomes? Our methodology, system dynamics, is more than just a simulation tool; it also a method of scientific inquiry that fosters operational thinking about how to improve the functioning of complex social systems. Our model is motivated by a case often cited as the seminal work in the implementation literature: Pressman and Wildavsky's narrative of problems that undercut a US policy to combat persistent unemployment among minorities in Oakland, California in the late 1960s.

Keywords Implementation · Public policy · System dynamics

The classic case of big projects having little effect is the 'Oakland' fiasco famously analyzed by Pressman and Wildavsky (1973). Their book launched the implementation research agenda for the public policy discipline, guided by the hypothesis that 'separation of policy design from implementation is fatal' (Pressman and Wildavsky 1973, xxiii). We previously used the Oakland case to illustrate the benefits of interdisciplinary collaboration between scholars in the fields of public policy and system dynamics, and this chapter builds on that earlier effort.[1]

[1]The first version, "Public Policy Implementation Modeling: The Case of EDA in Oakland," was presented at the International System Dynamics Society Conference in Boston in July 2015. A substantially revised version was presented at the IJPA Symposium at the University of Palermo in May 2016.

I.D. Wheat (✉)
System Dynamics, University of Bergen, Bergen, Norway
e-mail: david.wheat@uib.no

E. Bardach
Goldman School of Public Policy, University of California, Berkeley, Berkeley, USA
e-mail: ebardach@berkeley.edu

E. Borgonovi et al. (eds.), *Outcome-Based Performance Management in the Public Sector*, System Dynamics for Performance Management 2,
DOI 10.1007/978-3-319-57018-1_10

Our approach is framed by two questions related to the implementation requirements of a public policy initiative. First, can modeling reveal those implementation requirements and the potential for disappointing outcomes? And, can a series of simulations under different assumptions about a policy suggest feasible redesign options to improve the chances for desired outcomes?

In the first section, we discuss distinctive features of implementation problems in the public realm and make the argument for a system dynamics (SD) approach to some of those problems. The second section provides a brief overview of SD-based qualitative implementation modeling, using a hypothetical policy issue to illustrate the method. The third section is a more detailed examination of a quantitative approach that utilizes SD simulation methods to explore economic development projects involving government and private sector partnerships, and the Oakland case provides our illustration. Finally, we conclude with take-away messages about the value of both qualitative and quantitative implementation modeling, and suggest ways that others might integrate their methods with the approach presented here.

10.1 Implementation and Policy Design

Implementation appeared in 1973, but the seamless web of policy design and implementation was recognized long before the 1970s; e.g., Carl Friedrich observed in 1940 that the 'formation' of public policy 'is inseparable from its execution' (cited in Wilcox 1978). For much of the twentieth century, however, the dominant paradigm encouraged a research demarcation between the formulation of policy and its implementation; the former involving politics and the latter involving 'mere' administration. See Wheat (2010) for a brief historical review of the paradigm shift that occurred in the 1970s after the much-publicized implementation failures of some Great Society programs.

Continuous policy resistance in the public arena accounts for some of the observed gaps between public program outputs and their impact. The political conflicts that have beset the adoption process do not disappear during the implementation process; in some cases, they may be aggravated. New conflicts may appear, lured out of hiding by issues that come up during implementation but had been suppressed or invisible previously. These conflicts, together with the problems of turning a policy over to existing public-sector bureaucracies and perhaps to a host of private-sector partners at the same time, guarantee a rocky implementation process. The results, frequently, are delay, erosion of policy goals, cost overruns, the intrusion of various interests seeking to capture economic rents, and a degradation of whatever future operational capacity was envisioned.

How might policy designers cope with the contingencies and probable setbacks of the implementation process? First, they must take some responsibility for implementation and avoid assuming it is someone else's job.[2] They can do this by anticipating implementation issues during the design process and crafting policies that would be reasonably robust against the difficulties of implementation. This means building in extra time for delays caused by busy or uncooperative bureaucracies, budget problems caused by overly optimistic financial planning, the sacrifice of certain goals to political and administrative compromises, and workarounds that lead to building a program out of components (such as a certain proportion of untrained or incompetent personnel) that are less well suited to the task than originally assumed. In other words, at the design stage, it may be possible to anticipate potential implementation obstacles and draft contingency plans for midcourse adjustments. Hence, the capacity to confront, assess, and make those tradeoffs might be built in advance.

Effective advance planning for such contingencies requires a systematic method. Richard Elmore's 'backwards mapping' approach can be useful: listing all the elements one would need to be working together once an operational system has been assembled, and then planning how to acquire them (Elmore 1979). One of us (Bardach and Patashnik 2016) recommends postulating certain failures (e.g., huge delays, complete program collapse, and bureaucratic resistance) and then writing, from some vantage point in the future, scenarios about how they occurred.

Here, we suggest simulation modeling as a useful implementation planning tool. We use the system dynamics (SD) approach because it is more than merely a quantitative tool for generating internally consistent projections. It is a method of scientific inquiry that helps develop an intuitive grasp of the functioning of complex systems. Compared to less formal approaches, it can help planners anticipate both intended and unintended effects of policy options. First, modeling insists on confronting implementation details often overlooked by policy designers. Secondly, many important details become visible only when the implementation of the policy at hand intersects with other systems within the larger governance context, e.g., procurement rules that severely limit management options or cause delay, local zoning ordinances that obstruct construction plans, and expenditure rules that preclude advancing payments to contractors before work is performed. Formal modeling forces designers to try to analyze what is admittedly a very uncertain field of forces. Thirdly, when bureaucracies become involved, it is often hard to know what will be happening within their sometimes opaque and unpredictable worlds. Certain general outlines can be theorized, but a lot depends on the details of personalities in government positions. Again, the modeling exercise insists on making explicit guesses about the relevant bureaucratic behavior. Finally, systematic modeling makes various value dimensions more visible than they might

[2] A light illustration of responsibility avoidance is Will Rogers' facetious suggestion during World War I that the best way to fight enemy submarines was to boil the Atlantic Ocean. When asked how that might be done, he replied, "I'm a policy man. I let others worry about implementation" (cited in Wheat 2010).

otherwise be. At the design and adoption phase, one naturally thinks about costs and effectiveness. But as one moves toward implementation, one has to think about delays, goal erosion, and rent-seeking. Formal modeling does not guarantee that unpleasant surprises can be avoided, but it enables policy designers to use a model as a training ground—practice implementation, experience setbacks, and test redesign strategies—in ways that might later prove useful to street level implementers; e.g., see Wheat (2015).

Implementation analysis begins with a definition of the policy to be implemented. At a minimum, the definition should include (1) the nature of the policy mandate intended to accomplish something through the use of a government program, (2) an agency that will take the lead in the activity, and (3) some resources accessible to the agency. Typically, the lead agency will have to *assemble* program elements from other agencies, both public and private, into an operating system—the intended output of the system being, for example, a stream of subsidies or compliant behaviors.[3] This assembly process has three main streams. One is *technical*: the elements that need to be put in place to operate a program, such as personnel, organization, office space, manuals, training, clients, hardware, and procedures. Exactly what these elements are will depend on the particulars of the program. The second stream is *administrative*: authority to hire personnel, to expend budget dollars, to procure equipment, and so on. This stream supports activities in the technical stream, though perhaps with some friction and delay, because it proceeds somewhat independently, by its own logic and according to government-wide rules designed in large measure to prevent waste, fraud, and abuse. Thirdly, the *political* stream contains the support or approvals, in their great variety of forms, needed or useful for legitimating a government activity even after a general approval has been given for a policy or project. Given the US federal form of government, this often means that federal agencies seek general cooperation or acquiescence from their state and local counterparts (and constituencies they represent) and, in some cases, from private-sector partners.

10.2 Qualitative Implementation Modeling

The prospect of dealing with mathematical equations is not appealing to many who are engaged in the policy design process. This can cause resistance to using formal simulation models during that process. One way to lower that barrier is to begin with models that are qualitative rather than quantitative. A diagram of an SD model is a conceptual map that can be explored by policy designers without the cognitive burden of mathematics. Such a diagram is a qualitative model of a social or economic structure, including proposed structural changes, i.e., including policy options. It can be used for preliminary feasibility testing of policy proposals by

[3]Bardach (1977) develops the concept of an implementation assembly process.

encouraging analysts to envision policy outcomes—intended and unintended—and question how a policy would work in practice. When problematic feasibility issues are identified, planners can discuss ways to redesign the policy to improve the feasibility and raise the chances for successful implementation. The result is a revised conceptual model that reflects rejection or revision of initial options, hopefully with justifiable expectations of a more feasible plan for addressing the policy issues.

Qualitative feasibility testing begins by studying a diagram of a proposed policy and raising questions about it. The intent is to brainstorm political, administrative, and technical constraints that might impede the policy's adoption or prevent a policy from achieving its desired outcomes without negative side effects and then suggest ways to redesign the policy to improve its feasibility. This has proven to be an effective way to sensitize future policy designers trained in SD modeling. Students at the University of Bergen use this method in a master's level policy design and implementation course, while learning how to build implementation structure into their models and how to conduct feasibility analysis alongside cost-benefit analysis (Wheat 2013). Figure 10.1 displays a diagram that will be used to illustrate qualitative feasibility testing. The policy issue concerns regulation of over-fishing in a coastal region, and the model is adapted from Morecroft (2007).

The small inset diagram in Fig. 10.1 shows the historical downward trend in the fish stock, plus two alternative futures: continued decline or stability at a higher level. Symbols in the diagram illustrate the three building blocks of SD models: stocks, flows, and feedback loops. The boxes represent stocks (ships at sea and in the harbor, plus the fish population). Flow icons are the 'pipelines with valves' that control the rate at which material moves in and out of the stocks. Feedback effects are illustrated by arrows that form closed circles of mutual causation.

In this example, the policy feedback loop would regulate the number of ships at sea to achieve the desired fish stock. Government regulators would set a target for the number of ships at sea, based on estimates of the fish stock and a comparison with the desired stock. When the fish population is threatened by 'too many' ships at sea, some would remain docked in the harbor. When the situation improves, ship owners would be permitted to take more ships to sea.

After studying the model diagram, the policy design task is to identify political, administrative, and technical feasibility issues that might occur if such a policy were proposed or adopted. Below is a sample of the kinds of feasibility questions that inevitably arise during implementation analysis of the qualitative policy model displayed in Fig. 10.1.

Political Feasibility Issues

1. Does the public generally accept this kind of government regulation of business activities?
2. Will ship owners obey the regulations? Will they interfere with enforcement?
3. Will the government pay for ships sitting in the harbor?

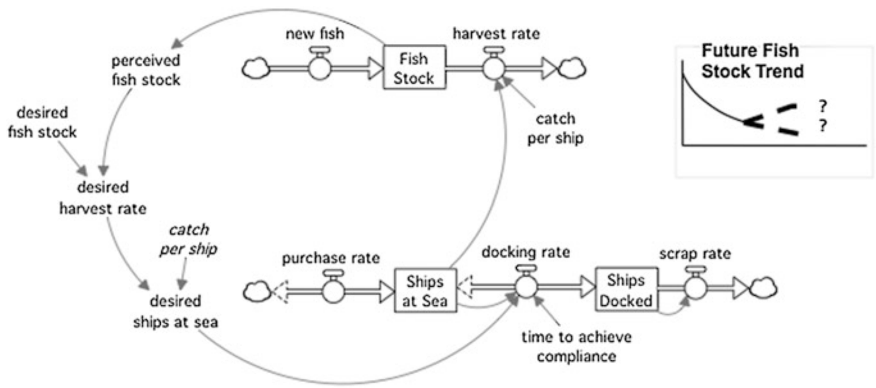

Fig. 10.1 Qualitative policy model of fishing regulation (simplified adaptation from Morecroft 2007, p. 347)

4. What groups are likely to oppose this policy?
5. If idle ships mean idle fishermen, how does that affect the local economy? Will there be pressure for government compensation?
6. What about the ships sailing under a foreign flag? What is the geographic boundary for this policy? Will this policy conflict with existing treaties or trade agreements?

Administrative and Technical Feasibility Issues

1. Which agencies are responsible for estimating (perceiving) ships at sea and the harvest rate? How reliable are their estimates, and what kind of delays should be expected?
2. Which agencies are responsible for estimating (perceiving) the fish stock? How reliable are their estimates, and what kind of delays should be expected?
3. Who will decide desired fish stock? Will the decision be based on an accepted scientific theory? Is there a 'scientific consensus' on the answer to this question?
4. Which agencies are responsible for deciding which ships remain in the harbor? How are those decisions made?
5. Do the agencies have adequate resources (funds, personnel, technology, experience) to do their various tasks?

Brainstorming feasibility questions in the context of a specific policy design is a sensitizing activity. It raises awareness of the potential for policy resistance during both the adoption and implementation stages, and it emphasizes that 'in a system, you can't do just one thing.' The *designed output* of the policy might be a precisely worded set of regulations aimed at a *single desired outcome*. Yet, the exercise reveals the potential for *multiple actual outcomes*, some of which could lower political support during the policy adoption stage or undermine achievement during implementation.

Qualitative implementation modeling may be sufficient to enable planners to redesign policy proposals in order to reduce chances for disappointing outcomes, or to narrow the number of promising policy options to a feasible subset. For some complex issues, however, simulation modeling can add value to the qualitative approach by quantifying cause-and-effect relationships implicit in a policy idea and projecting the likely behavior that would emerge over time. Moreover, the range of policy outcomes may be particularly sensitive to uncertain assumptions in the minds of policy designers, and simulation modeling enables testing the behavior of a model under various assumptions.

In short, while both types of models can represent the *structure* of a policy, only a quantitative simulation model permits analysis of the dynamic *behavior* that is expected to arise from that structure. In the next section, we demonstrate how quantitative modeling can aid the policy planner, and we use the Oakland case to provide a real-world context for a stylized simulation model of policies aimed at local economic development.

10.3 Quantitative Implementation Modeling

We approach the building of the simulation model from three directions. First, we rely on available empirical evidence which, in this case, consists of a well-documented case study of an implementation process to help ground our model in at least one actual instance. This provides structural and behavioral benchmarks against which to compare our model's structure and behavior. Here, our benchmarks are provided by Pressman and Wildavsky's case study of a US federal policy initiative to increase hiring of long-term unemployed persons in Oakland, California, during the 1960s.[4] Secondly, we rely on our general theoretical understanding of governance and political processes. For example, we assume that government agencies typically specialize by mission—turning out grants to businesses, for instance, or guarding the integrity of procurement decisions—and tend to emphasize the priority of that mission at the expense of other values that, when balanced properly against the mission priorities, might deserve higher weights than they receive. Finally, we conceptualize as stocks and flows the variables suggested by our theoretical and empirical foundation, and define the boundary of the model broadly enough to reveal an endogenous feedback structure that accounts for the behavior of the model.

[4]The Pressman and Wildavsky book is the sole source of facts about the Oakland case, although their case study has generated analyses too numerous to count (e.g., a Google search for "Pressman and Wildavsky" yields 15,000 hits).

10.3.1 The Oakland Story

Policy designers are habitual optimists. The world where the policy will be implemented is, by nature, less hospitable to the designers' wishes than they would like to believe. Things cost more, take longer, and are more subject to being hijacked by political interests who do not care much about the original policy objectives but do care a lot about their own policy, institutional, and career interests. With only occasional exceptions, therefore, the implementation phase of policy-making is disappointing. And the story of EDA in Oakland is not one of those exceptions.

The Oakland case is an old one.[5] Yet, it suits our purpose for two reasons. First, it is well known for its illustration of implementation issues that are uniquely problematic in the public sector, namely those requiring reconciliation of diverse public and private interests and coordination of multiple bureaucratic programs and procedures. Another reason is its special relevance to an outcomes shortfall: it was a jobs-for-hardcore-unemployed project that cost more than $10 million but created fewer than 100 jobs, far from the goal of 3000. Moreover, few if any of the jobs went to the target population. Another $13 million was scheduled for spending, but the plug was pulled on the Oakland project before the wasted effort could escalate even higher.

The seeds of the project had been planted in 1965, when the US Congress authorized and funded a government subsidy program for public works projects that would support local economic development designed to encourage hiring long-term unemployed persons, most of whom were racial minorities. The lead agency was the Economic Development Administration (EDA) in the US Department of Commerce, and EDA focused its resources on Oakland, California. A local public agency, the Port of Oakland (the Port), would receive the federal government funds and build an airplane maintenance hangar, which it would lease to World Airways (World). In effect, EDA was contracting with World through a public-sector intermediary. In return, World was expected to hire local unemployed persons for the short-term construction jobs and for the more skilled long-term maintenance jobs. The expectation was that EDA and World would jointly arrange for the training of job seekers and new hires. The needed technical elements to be assembled in Oakland were: (1) jobs, (2) qualified potential employees; (3) a way for government to enforce hiring commitments by recipients of the funds; and (4) training for a large fraction of the potential employees.

---

[5]As are the authors. One of us was literally present at the creation of the Oakland case study project led by Pressman and Wildavsky at Berkeley, having been a professor of public policy at the Goldman School of Public Policy since 1970. At that time, the other author was a student of public policy at Harvard's Kennedy School, thereafter serving on the White House staff. We have seen our share of gaps between policy efforts and outcomes, not only in academic research but also while in government staff positions and as consultants to governments.

Despite the availability of EDA funds amounting to $23 million in 1966, signs pointing to a disappointing outcome were evident early in the project, as various delays ensued. During lengthy contract negotiations with EDA, World objected to any provision that would permit EDA to reclaim funds contingent on post hoc approval of World's hiring successes; in the end, World would agree only to including a *plan* for hiring in the initial contract. In 1968, the Port estimated a cost overrun of nearly $5 million for the hangar project and asked the EDA to absorb it. EDA tried to use the occasion to leverage its demands on World to further the hiring and affirmative action goals, and continued to do so through early 1969, when it finally turned down the Port's request. Meanwhile construction did not go forward. On at least one occasion, World apparently threatened to back out of the project if the EDA put World at greater financial risk. Early in 1969, World told EDA that it was withdrawing its hiring plan in favor of one that promised less minority employment.

The worker training program never materialized. The program needed numerous approvals: by World, by units within the US Department of Labor and the US Department of Health, Education, and Welfare (HEW), by the California state Department of Employment Development, and by EDA. Reviews and negotiations went on for nearly 2 years, until HEW finally vetoed the plan in 1968 and World ceased participation in plans for worker training.

The contracts for architectural plans for the hangar were not let until mid-1971, nearly 6 years after the initial mandate, and fully 5 years after the EDA had made a big public announcement that it had a project on track that would produce 3000 jobs in the Oakland labor market. In the end, the number of new jobs totaled only 2–3% of that goal, and only a small fraction of that total went to the target group: long-term unemployment persons.

10.3.2 Behavior of the Model

The SD modeling process usually begins by studying a time series graph that displays historical patterns of behavior that a model will be designed to explain. However, despite several careful readings of *Implementation*, all we can say for sure is that the $23 million of EDA funds were not fully distributed during the 6 years from 1966 to 1972, a period within which most observers expected the investments to be made. Cumulative spending was closer to $10 million. And the number of *new* jobs created was nowhere near 3000; in round numbers, it was probably no more than 100, if that many. We want to compare these rough historical estimates at the end of 1971 with the simulated results generated by our model.

Comparing model behavior with even rough estimates of historical Oakland 'data' requires calibrating our generic model with numerical estimates or, in some cases, guesstimates of Oakland-relevant parameters. Given what we know about the Oakland case, we can safely assume that training capacity did not exist and that the total number of jobs actually created in the Oakland project was no more than 100. Therefore, in the model, we set training capacity equal to zero and the initial value of qualified long-term unemployed persons (those not needing to be trained) at 100. Some of the other parameters in the model are not necessarily Oakland-specific, although we attempted to base estimates on empirical economic data for Oakland during the 1960s whenever possible.

The results are displayed in Fig. 10.2, which compares the simulation results with our knowledge of cumulative spending and employment. The thin lines indicate simulation results and the wide bars represent the data estimates for cumulative EDA spending and employment.[6] We made no attempt to speculate about the unknown historical pattern; thus, the bars show the best guess total at the end of the project.

The top frame indicates that the simulated cumulative spending after 6 years is similar to the 'data' we have (about $10 million). Likewise, the bottom frame shows a simulation result that is consistent with the upper bound estimate of new jobs (100) actually created by the Oakland project.

The simulation experiment described above, while pertinent to the circumstances in Oakland, does not permit exploring the full range of behavior our model can generate, primarily because we assumed zero training capacity. We will now reverse that assumption and observe how strategic interaction between government agencies and private-sector institutions can generate a range of plausible behaviors when training capacity is optimal. The interaction in the model can be aggregated and summarized as the degree of company *cooperation* with the government. In this context, full cooperation includes a shared goal for total project employment and the time period during which that goal should be achieved. That would mean, for example, company acceptance of a target capital-labor ratio that would be lower than the company's normal target. In our model, that has implications for a company's willingness to adopt the government's 5-year employment goal and the short-term employment targets; and the latter has immediate impacts on hiring. The desired pace of company investment may also conflict with the government's deadlines. These sources of conflict do not necessarily have to be activated; they can remain dormant and, if they do, we will call that 'company cooperation' with the government. Conversely, a lack of alignment between the goals of the company and the government constitutes lack of cooperation.

The company's response to government sanctions is also indicative of the degree of cooperation. If the company falls behind the government's desired hiring rate and pays a penalty in terms of slower cost reimbursement, cooperation means that the company acknowledges its failure and the legitimacy of the penalty and does not

[6] In the model, *LTU Employed* refers to long-term unemployed persons actually hired, and that is the variable graphed in Fig. 2. However, we should emphasize that whatever the actual employment total in Oakland, only a fraction of that number included the target population, and this discrepancy is not specified in our model. In addition to assuming no training capacity, the simulation results in Fig. 2 also assume weak cooperation between World and EDA, the interpretation of which is explained in the text.

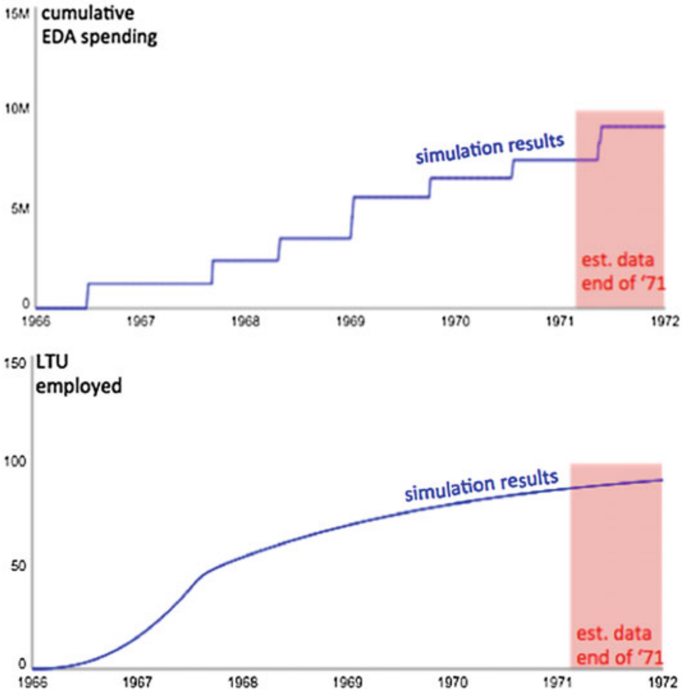

Fig. 10.2 Model behavior and estimated Oakland data (historical pattern of data unknown)

retaliate in any way. In the model, retaliation by an uncooperative company takes the form of slowing the hiring rate. Figure 10.3 displays the results of three simulation runs, each with different assumptions about company cooperation. Note that the simulation continues beyond the 5-year government subsidy program; thus, this should viewed as a generic test of model behavior that has nothing to do with the details of the Oakland case even though the horizontal axis still refers to that time period in history.

As before, *LTU employed* refers to total project employment. The *Target LTU employment* refers to the company's goal, which matches the government's goal only when there is full cooperation. The best-case scenario (top frame) requires optimal training conditions (capacity to train 500 persons per year, at least 20% enrollment potential each year, 100% training success, and no dropouts) *plus* full company cooperation. That scenario generates employment that approaches the government's goal, but it takes more than a decade to do so, despite tacit company acceptance of the government's hiring schedule. Failure to keep pace with that schedule results in government sanctions (delays in cost reimbursement), but the full cooperation assumption assures no retaliation in this scenario and, eventually, the government's desired employment level is reached. With weak or nonexistent company cooperation (middle and bottom frames), employment stabilizes below

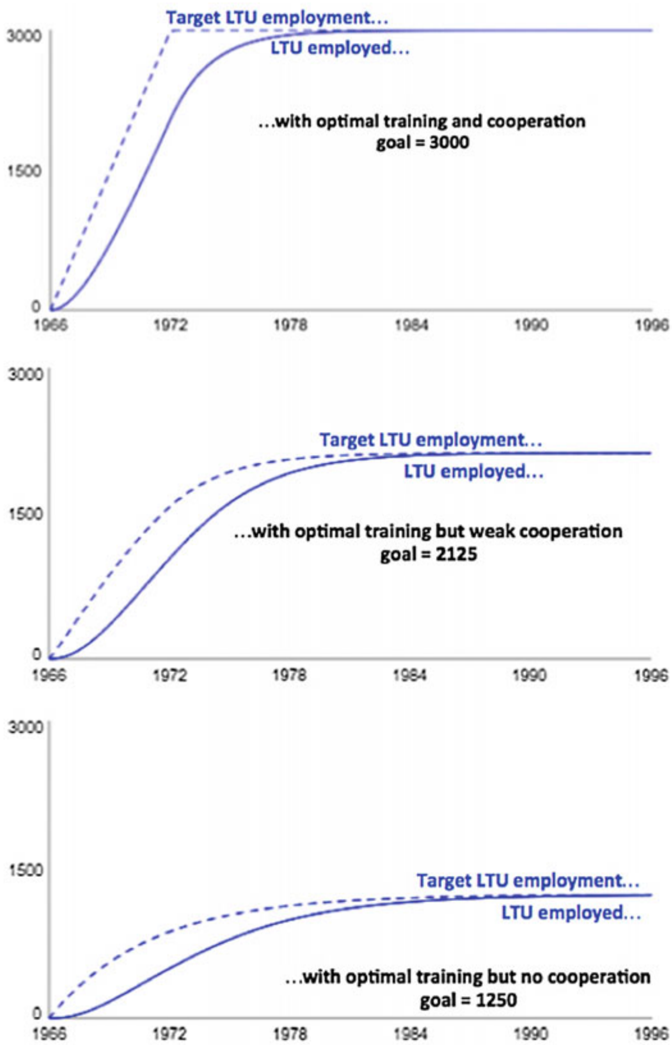

Fig. 10.3 Growth toward rising goals

government's goal even if there is optimal training capacity. Despite the quantitative differences in Fig. 10.3, there is a similar qualitative behavior in all three frames: goal-seeking patterns for both the target and actual employment levels. Employment rises toward a rising employment target. Next, we examine the structure of model, seeking the source of these persistent dynamic behavior patterns.

10.3.3 Structure of the Model

The full model consists of four sectors: hiring, training, spending, and reimbursing. Before examining the detailed stock-flow-feedback structure in each of those sectors, we present a high level view of the feedback structure responsible for the goal-seeking behavior pattern displayed in Fig. 10.3. A simple set of feedback loops ties together three sectors of the model: hiring, company spending, and government reimbursement. The feedback loop diagram in Fig. 10.4 displays the source of the goal-seeking dynamics in the model.

Feedback loops are distinguished by their positive or negative polarity. Positive feedback loops have self-reinforcing effects. There is no normative connotation in the ‘positive’ label; behavior that feeds on itself can cause growth or collapse and, depending on one’s values, can be virtuous or vicious. To avoid a misunderstanding, positive loops are often called *reinforcing* loops, denoted in feedback loop diagrams by the letter R. In contrast, negative feedback loops have self-adjusting effects. Their goal-seeking structure counteracts tendencies for a system to grow or collapse. Sometimes called *counteracting* loops, they are denoted by the letter C.

The feedback loop diagram in Fig. 10.4 reveals two counteracting loops, C1 and C2, that are responsible for the goal-seeking behavior in the full model, and a reinforcing loop R1 that has the potential to weaken loop C1 and hinder its goal-seeking tendency.

The hiring loop C1 functions in a way that closes any gap between target and actual employment. The faster the hiring adjustment time, the quicker the gap is closed. Previously, we discussed the potential for government to seek leverage over the company’s hiring process by slowing the reimbursement process. When actual employment fails to keep up with the government’s scheduled employment goal, the reimbursement time increases. The company’s retaliation option is to slow the hiring process even further. That is the essence of loop R1 when activated by sanctions and retaliation; it can frustrate both the government and the company and, in so doing, weaken the net hiring loop C1.

Target employment depends on the company’s stock of physical capital (infrastructure, equipment, tools, etc.) and the desired capital-labor ratio. If investment exceeds depreciation (not shown), the company’s capital increases and the target for employment increases proportionately. Growth in the capital stock is controlled by loop C2, which closes any gap between actual and target capital. To the extent that the company aligns its operating strategy with government’s policy goals, the desired capital-labor ratio, the target for capital, and the pace of adjustment—and, therefore, target employment—would reflect the government’s goals. Lack of company cooperation would reduce alignment with government’s goals, lower the target employment, and reduce the hiring rate in loop C1. These strategic interactions between government and the company are exogenous in the current version of our model. The degree of goal alignment can be varied by the user of the model and the impact of different assumptions can be observed in the simulation results.

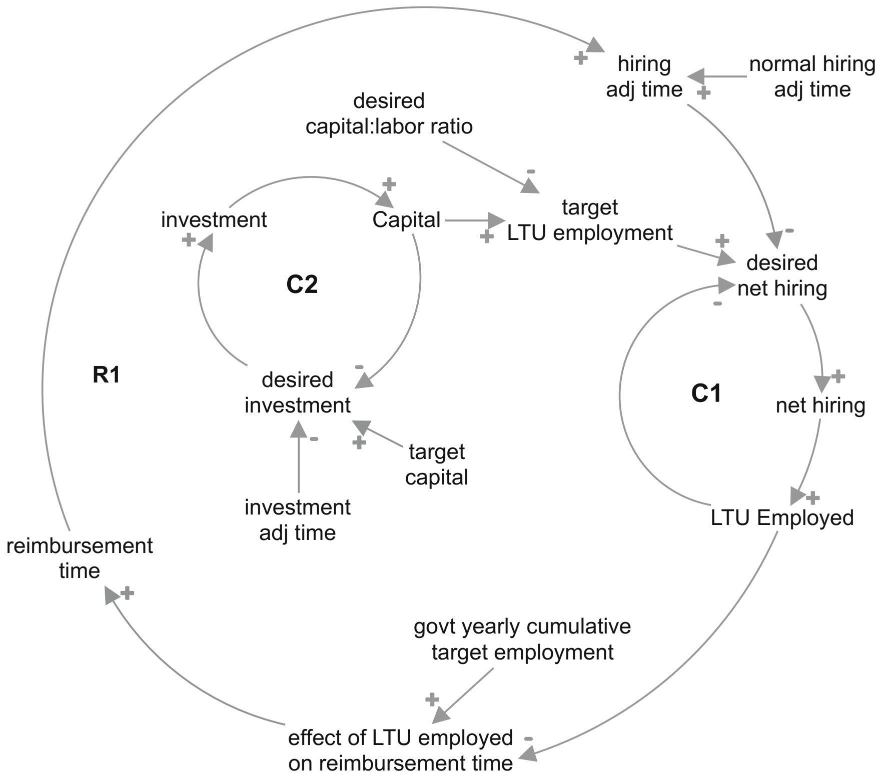

Fig. 10.4 Feedback loops responsible for goal-seeking behavior extracted from full model in Fig. 10.10

Our hypothesis that loops C1 and C2 are responsible for the goal-seeking behavior is supported by two simulation experiments with the full model.[7] Figure 10.5 displays the model's behavior when loops C1 and C2 were deactivated or 'cut' during the simulation. In the left frame, cutting loop C2 stops investment and the growth of the capital stock which, in turn, stops the growth in *Target LTU employment*. In the frame on the right, cutting C1 stops the growth of *LTU Employed*. The employment target is not part of that loop and continues to rise to its own goal, unaffected by the deactivation of loop C1.

In the remainder of this section, we examine the details of the model's stock-and-flow structure and gain additional insight regarding the source of dynamics in the model. Figures 10.6, 10.7, 10.8, and 10.9 display close-up views of the four sectors in the model, and the full model is displayed in Fig. 10.10. Although the 'EDA in Oakland' case motivated the model, we have adopted

[7]For this test, a training program is activated so that the stock of qualified applicants is large enough to accommodate the desired hiring rate.

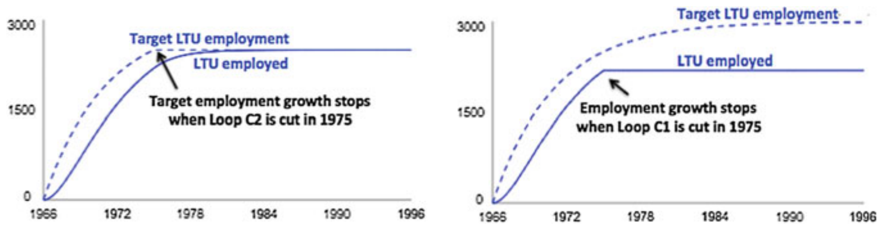

Fig. 10.5 Cutting counteracting feedback loops stops goal-seeking growth

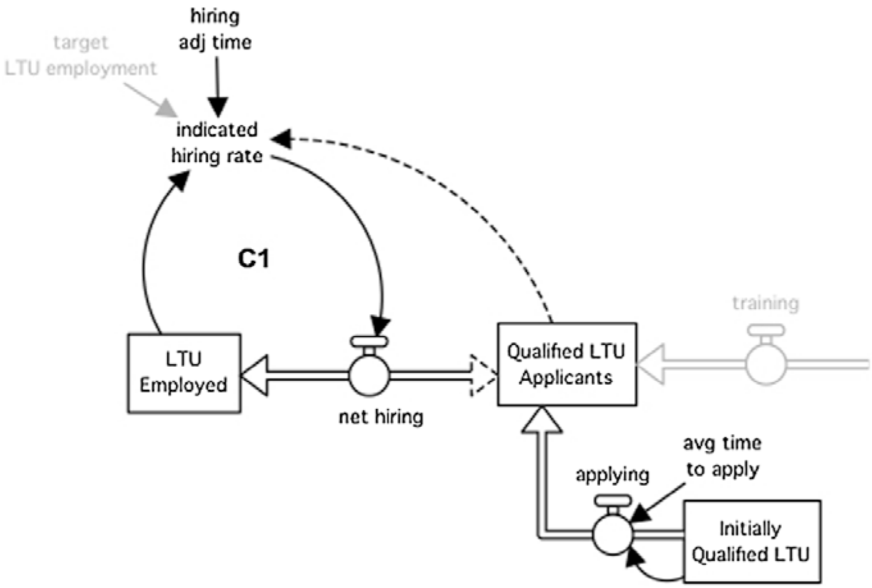

Fig. 10.6 Hiring sector

generic names such as ‘government’ for EDA and ‘company’ for the various private interests, the largest of which was World Airways. We have also selected round numbers for parameter values such as delay times and various coefficients in the model affecting spending, training, etc. All parameter values can be modified by users wanting to test the effects of different assumptions. The generic approach facilitates adapting the model for other policy design research tasks, and using it as a ‘method of inquiry’ tool for policy designers.

Figure 10.6 displays the stock-flow-feedback process that governs hiring in the model. As long as *target LTU employment* exceeds *LTU Employed*, *Qualified LTU Applicants* are being hired. When *net hiring* is negative, layoffs occur. The faster the *hiring adjustment time*, the sooner actual employment rises to meet the target. This is the same counteracting loop C1 displayed in Fig. 10.4. Here, however, the stock-and-flow structure specifies how the process operates; what Richmond (1994)

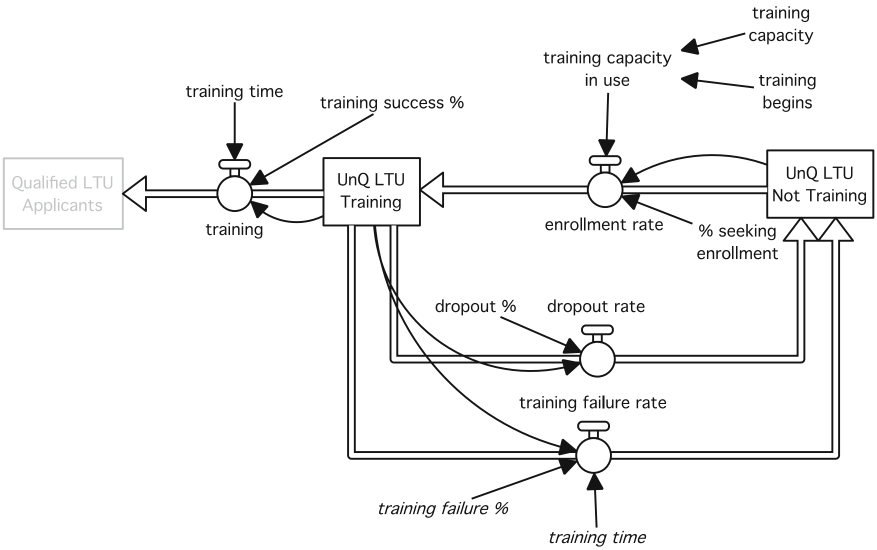

Fig. 10.7 Training sector

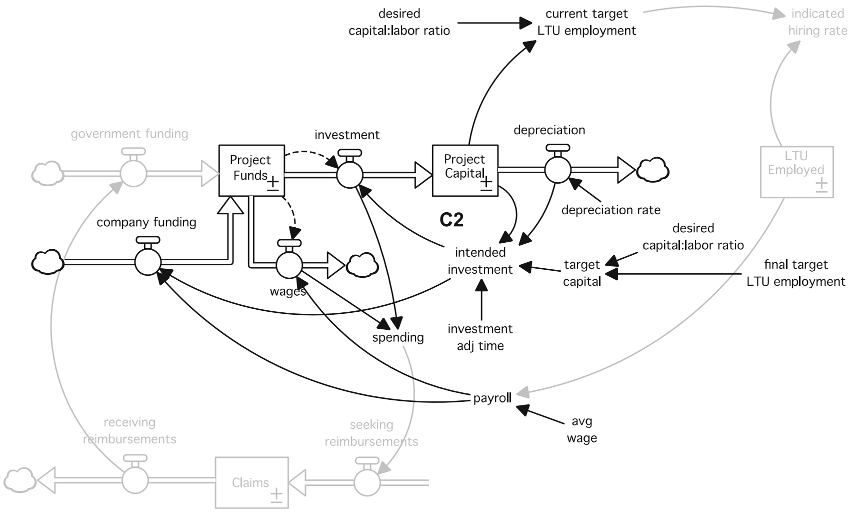

Fig. 10.8 Spending sector

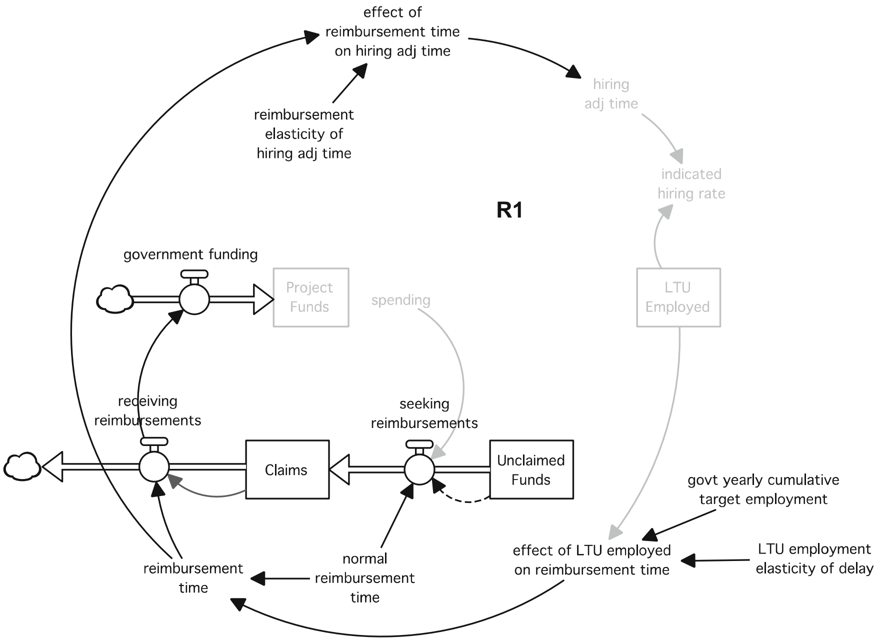

Fig. 10.9 Reimbursing sector

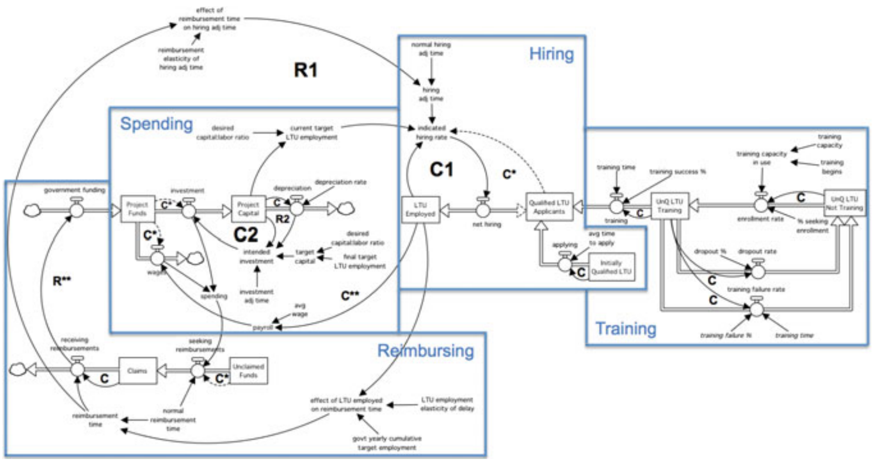

Fig. 10.10 Simplified view of full model

calls the ‘physics’ or ‘plumbing’ of the system. Significantly, it reveals the real-world constraints on hiring. Hiring requires a stock of *Qualified LTU Applicants* (initially zero) that depends on an inflow of applicants from the *Initially Qualified LTU* (assumed to be 100) or those successfully completing their *training*. As long as there are qualified applicants, loop C1 operates freely. Otherwise, the dashed link signals the absence of qualified applicants and the *indicated hiring rate* is zero, making loop C1 dormant.

Figure 10.7 displays the training sector of the model and reveals its connection to the hiring sector, via the *training* flow. Although training never materialized in Oakland, this sector is an essential component of any model of a job-creation policy because it raises critical policy design questions. The annual *training* rate depends on the number enrolled in a training program (initially zero), the time it takes to train them, and the fraction successfully trained; i.e., those truly qualified and available for employment. Those failing to be trained rejoin the ranks of the unqualified LTU not enrolled in a training program (6000 initially, based on rough estimates for Oakland in 1965). In addition, there are dropouts. The annual *enrollment rate* depends on the capacity of the training facilities and the percentage of LTUs enrolling each year. In the Oakland story, *training capacity in use* is zero, which prevents enrollment and training and (in Fig. 10.6) hiring. In other cases, training capacity may exist but insufficient enrollment, high dropout rates, or ineffective *training* may limit growth in the number of qualified applicants. Each of these leverage points should be highlighted during the policy design stage to activate contingency planning.

The spending sector is displayed in Fig. 10.8, along with its connections to the (dimmed) hiring and reimbursement sectors. Company *spending* is the sum of *investment* and *wages*, and the total drives reimbursement *Claims*. To jump-start the process, *company funding* is needed, but *government funding* replenishes the *Project Funds* stock as reimbursements are received. The dashed links to *investment* and *wages* slow those outflows if funds run low, and no *spending* occurs if there are no funds at all. *Investment* adds to *Project Capital*, in response to feedback loop C2 that gradually adjusts the current capital stock to its target value.

Both the capital target and adjustment time are influenced by strategic interaction between the company and the government. With full cooperation from the company, the *desired capital-labor ratio* and therefore, *target capital*, will reflect the *government’s final target LTU employment*. With company resistance, the target will more likely resemble the company’s capital-labor ratio preference. Likewise, the degree of company alignment with the government’s project deadline determines the time period over which the capital stock is adjusted (in the model, the particular strategic reactions are exogenously controlled by the user, and the controls are not shown in Fig. 10.8). For private companies, demand for labor is usually derived demand; i.e., it depends on the demand for the goods and services that labor can produce. Here, we simplify the labor demand structure by assuming the company regularly adjusts its target for employment based on the level of installed capital and the (exogenously determined) *desired capital-labor ratio*.

The *current target for LTU employment* then influences hiring, and changes in *LTU Employed* affect the *payroll* and the next round of *spending*.

The final part of the model to inspect is the reimbursing sector, displayed in Fig. 10.9. This sector governs the reimbursement process after the company submits a project spending claim. This sector interacts with both the spending and the hiring sectors (both partially displayed and dimmed). In the Oakland project, EDA distributed funds only to reimburse company spending after the fact. One could imagine other possibilities, but that is not an uncommon way that governments distribute grants; thus, it is the procedure we assume here. We also assume the government slows the reimbursement process during periods of negotiation when the company fails to meet government's annual hiring targets (estimated as a linear trend from the beginning to the end of the project). As discussed previously, feedback loop R1 implements the company's retaliation when reimbursements are late. The effect of the loop is to lengthen the hiring adjustment time, further slow the employment of LTUs, and reinforce a vicious mutual effect on the government, the company, and the long-term unemployed persons waiting to be hired.[8]

Figure 10.10 displays a simplified version of the full model, with several parameters and one flow (company funding) deleted for clarity. Close scrutiny reveals 16 feedback loops, 13 of which are counteracting, and only those could account for the goal-seeking behavior generated by this model. The four denoted as C* (with a dashed link in the loop) are dormant unless their relevant stocks approach zero.[9] Six of the remaining counteracting loops have an implicit purpose of draining their stocks to zero; none could be pushing employment up toward a goal. For example, the training sector's counteracting loops constrain hiring; the cumulative net inflow to *Qualified LTU Applicants* represents the maximum number that could be hired but that number does not drive the hiring rate. That leaves only loops C1 and C2 as the source of goal-seeking dynamics, with loop R1 weakening the employment adjustment impact of loop C1, as confirmed by our previous analysis (Figs. 10.4 and 10.5).

[8]The strength of loop R1, assuming it is activated, depends on assumptions about the reaction functions influencing the government and the company. For example, we assume the government increases the normal reimbursement time by 3% when LTU Employed is 10% below the government's target level (elasticity = −0.3). We assume the company slows the hiring adjustment to match the slowdown in the reimbursement process (elasticity = 1.0).

[9]The reimbursement loop R** aggregates two loops—one stemming from *wages* and the other from *investment*. However, R** never becomes a closed loop unless the C* loops are active, in which case *Projects Funds* would be zero. If R** raised *Project Funds* above zero, that would make the C* loops dormant and immediately deactivate R**. The *Project Funds* stock constrains spending on *investment* and *wages* but it does not drive those outflows. Similarly, the potential C** payroll loop has no effective feedback effect on *LTU Employed* because the loop is only closed when *Project Funds* is at or near zero. We include R** and C** in our total feedback loop count, but they could not be responsible for the model's goal-seeking behavior.

10.4 Conclusion

Qualitative modeling can sensitize policy designers to the technical, administrative, and political feasibility issues that can impair policy initiatives with time-delayed destructive elements. Quantitative simulation modeling can add value to qualitative maps by revealing the dynamics of complex systems, and experimenting with a simulation model provides vicarious experience in policy design and can hone the skills of policy designers.

The questions raised by the qualitative fishing regulation example (Fig. 10.1) illustrate how implementation difficulties can be predestined by the original policy design. And, in the Oakland example, the diagrams in Figs. 10.6, 10.7, 10.8, 10.9 and 10.10 could be used to generate questions about make-or-break issues such as the training program or to anticipate the likelihood and implications of divergent company and government goals or the company's reaction to government sanctions and the likely impact of that vicious circle on the pace of employment. A collaborative effort to sketch a causal model of how a policy is expected to work is likely to generate critical questions about policy ideas. A policy design tool that provokes this kind of thinking and communication promises to be useful to those with responsibility for envisioning outcomes.

Quantitative simulation models encourage planners to view feasibility issues in the context of activity streams that flow over time, interact in unexpected ways, and generate outcomes that may not be intended. The Oakland model, for example, demonstrates how millions of dollars could be spent before it becomes apparent that no training program would materialize. Witnessing a stream of spending that does not produce jobs could energize efforts to make sure that obstacles in the way of training would receive early and continuous attention. Simulation experiments also reveal (in Fig. 10.3) that a training program is a necessary but not sufficient component of a jobs-creation project. Without company cooperation, the employment potential could be well below the government goal even with optimal training capacity. Moreover, formal methods of quantitative model analysis can identify the structural reasons for dynamic behavior (Figs. 10.4 and 10.5) and provide valuable clues about how to redesign a process to achieve a better outcome. For example, simulation results reveal how the company retaliation feedback effect (loop R1) can undermine the hiring process (loop C1). If apparent during the policy design stage, such results could foster debate about the potential for certain types of sanctions to be counterproductive, and a model could enable tests of alternative ways to sanction. Even without further testing, the simulation results could raise the debate about sanctions to a higher level of specificity about how they would work, the reactions they might provoke, and the expected impact on outcomes. Simulation results in our example also underscore the critical importance of alignment between company and government goals regarding employment targets and desired levels and timing of investment, and reveal the naiveté of simply assuming that subsidies would result in company operations that followed government guidelines instead of standard business guidelines and procedures.

We encourage policy designers to look for synergy in the joint use of these approaches with other good methods; for example, the failure scenario writing exercise described in Bardach and Patashnik (2016). Designed to brainstorm ideas for disaster avoidance or damage control, that exercise can assess the feasibility of a policy option. When used in combination with qualitative feasibility testing, it would encourage mental simulation of unintended consequences. In addition, qualitative feasibility testing specifies implicit mechanisms in a policy, and that can enrich the scenario writing process by spotlighting the specific resources that must be assembled to facilitate implementation. The value is not in a model per se; the value is in how the modeling process can shape the mental models of the participants in advance and thereby influence their strategic thinking, their contingency planning, and their design of the content and transmission mechanism of a particular policy.

We acknowledge limits to implementation modeling. Modeling is no panacea for policy failures in public institutional settings characterized by conflicting views and shared powers. We do not think that everything about a policy that *might* be modeled *should* be modeled. Certainly, not all implementation-relevant factors are included in the model inspired by the Oakland story. Some of the limits are deliberate. Like a highway map that omits local streets, the details of a simulation model reflect its purpose, and a high-level model of a job-creation program will permit later addition of contextual details. Other limits are problematic. For instance, when considering how to model discrete as well as continuous patterns of political conflict among officials who share powers within and across governmental units, an argument can be made for an agent-based approach. Yet the more aggregated system dynamics approach is better for mapping endogenous feedback structure and encouraging operational thinking about how complex systems work and how they could be modified to work better. In this example, a methodological compromise may be justified and is certainly possible.[10]

We envision an accessible inventory of generic but insightful causal models that can be adapted for practitioners in the policy design arena. Developing such models requires closer collaboration than currently exists between the modeling disciplines and the public policy research disciplines, something we have encouraged (Wheat 2010; Wheat and Bardach 2015). A desirable by-product of such collaboration would be a new instrument in the research toolkit that policy analysts could use to improve understanding of gaps between policy inputs, outputs, and outcomes.[11]

---

[10]For example, *AnyLogic* (anylogic.com) software supports both agent-based and system dynamics modeling. Moreover, one of our colleagues at the University of Bergen, Pål Davidsen, is using features of *Stella Architect* (iseesystems.com) to represent individual agents interacting within a system dynamics model.

[11]The Oakland model is available for online simulation at https://sims.iseesystems.com/david-wheat/oakland/#page1. Readers wishing to use *Stella Architect* to study model equations and experiment with alternative formulations are encouraged to request a fully editable copy of the model from the authors.

References

Bardach, E. (1977). *The implementation game: What happens after a bill becomes a law*. Cambridge: MIT Press.

Bardach, E., & Patashnik, E. (2016). *A practical guide for policy analysis*. (5 th ed.). University of California, Berkeley Washington: CQ.

Elmore, R. F. (1979). Backward mapping: Implementation research and policy decisions. *Political Science Quarterly, 94*(4), 601–616.

Morecroft, J. D. (2007). *Strategic modeling and business dynamics: A feedback systems approach*. Chichester: Wiley.

Pressman, J. & Wildavsky, A. (1973). *Implementation: How great expectations in Washington are dashed in Oakland, or why it's amazing that federal programs work at all; this being a saga of the economic development administration as told by two sympathetic observers who seek to build morals on a foundation of ruined hope*. Berkeley: University of California Press.

Richmond, B. (1994). Systems thinking/system dynamics: Let's just get on with it. *System Dynamics Review, 10*(2–3), 135–157.

Wheat, I. D. (2010). What can system dynamics learn from the public policy implementation literature? *Systems Research and Behavioral Science, 27*(4), 425–442.

Wheat, I. D. (2013). *Teaching policy design: A case study of unintended consequences when the EU regulates hospital doctors' hours*. International System Dynamics Conference. Cambridge, USA.

Wheat, I. D. (2015). Model-based policy design that takes implementation seriously. In E. W. Johnston (Ed.), *Governance in the information era: Theory and practice of policy informatics*. New York: Routledge.

Wheat, I. D. & Bardach, E. (2015). *Public policy implementation modeling: The classic case of EDA in Oakland*. International System Dynamics Conference. Cambridge, USA.

Wilcox, R. F. (1978). Review of the implementation game: What happens after a bill becomes a law. *Policy Sciences, 10*(2–3), 225–227.