Submitted to the
22nd International Conference of the System Dynamics Society
July 25-29, 2004 (Oxford, England)

Modeling the Dynamics of Health Care Services for Improved Chronic Illness Management

Gary B. Hirsch
Consultant, Creator of Learning Environments
7 Highgate Road, Wayland MA 01778 USA
508-653-0161
GBHirsch@Comcast.net

Jack Homer
Homer Consulting
3618 Avalon Court, Voorhees NJ 08043 USA
856-810-7673
jhomer@Comcast.net

Abstract

System dynamics models have shown considerable value for planning and designing chronic illness management (CIM) programs. Such planning should take into account the need for sufficient capacity in health care delivery systems to accommodate the additional workloads created by CIM programs. Without the necessary capacity, feedback effects may undercut the effectiveness of these programs or even threaten their viability. After discussing, in the form of general propositions, the interplay between CIM and delivery system capacity, we present two relevant SD applications. One uses a model of a diabetic patient population to evaluate the contribution of clinical care specialists, community nurses who assist patients in a CIM program. The other application uses a health care simulator to demonstrate the need for coordination between improvements in care delivery and the implementation of CIM programs. The paper concludes with a discussion of future work that can be done in this area with system dynamics modeling.

Background

Improving the management of patients with chronic illnesses has become a subject of intense interest in health care research. A recent report from the Institute of Medicine, *Crossing the Quality Chasm*, emphasizes the importance of organizing care around “priority conditions” that include the most prevalent chronic illnesses (IoM 2001). A number of studies have demonstrated favorable impacts of programs designed especially for particular chronic illnesses. For example, Edward Wagner and his colleagues at Washington State’s Group Health Cooperative have demonstrated better control of diabetes, less health care required over time, and better outcomes as a result of careful monitoring of patients (Wagner et al. 2001). Others have demonstrated dramatic improvements in populations of patients with heart failure (e.g., West et al. 1997; Rich and Nease 1999). Programs to improve chronic illness management (CIM) mostly focus on individuals who are already known to have a chronic disease like diabetes or heart failure. Through the use of medications and changes in nutrition and physical activity, these programs attempt to slow the progression of the disease, and the frequency of acute episodes and deaths resulting from the disease. More comprehensive programs may also include diagnostic screening, as well as the medical management of patients who are at risk for developing chronic illness.

Although the programs are typically expected to reduce overall clinical workloads as the frequency of acute problems decreases, aggressive CIM programs initially involve more work for providers, and particularly primary care providers, as they endeavor to help patients bring their conditions under control. Without expanding the capacity of practices to provide such primary care, the temporary increase in workload may feed back to undermine improvements in the care of chronic patients as well as making access to care more difficult for all patients. For example, an overly busy medical practice may only be able to provide episodic, fragmented care rather than the coherent focus needed to effectively prevent and control chronic conditions.

Expanding primary care capacity for CIM is thus an important topic, but one that has received only limited attention to date. Discussions of health care capacity expansion have tended to center on operational approaches to improving clinical efficiency for dealing with acute episodes, rather than for the routine activities of CIM. One such operational approach is Idealized Design of Clinical Office Practice (IDCOP), developed and spearheaded by the Institute for Healthcare Improvement in Boston (IHI 2003). The IDCOP approach strives for greater efficiency through better patient flow and eliminating unnecessary activities. Elements include same-day appointments to reduce the waste associated with appointment no-shows, as well as open access and continuous flow of communications based on patient needs rather than fixed appointment slots. If IDCOP is implemented successfully, it can bring down provider costs, improve provider and patient satisfaction, and free up provider time for activities other than acute care, chief among them CIM. But it does not directly address how CIM itself might be done more efficiently and effectively.

Group Health Cooperative’s Edward Wagner has described several key elements needed for improved CIM, including the use of multi-disciplinary care teams within and connecting clinical practices, systematic attention to patients’ individualized needs, and supportive information systems and payment mechanisms (IoM 2001). While all of these elements are meant to improve the quality or effectiveness of CIM, they are not particularly designed to reduce clinical resource requirements, and in fact may cost more than they save, at least in the short term. At present, there are many open questions about the magnitude of clinical resources needed for aggressive CIM and about how to pay for such resources.

How System Dynamics Modeling Can Contribute to the Conversation

Chronic illness management is a subject well suited to system dynamics modeling, in part because of the long delays that separate the stages of disease through its onset and progression, and in part because of the complex dynamics that affect human service delivery more generally (Levin et al. 1976). SD work has been done in the past on specific chronic illnesses and conditions (Hirsch and Killingsworth 1975, Luginbuhl et al. 1981, Hirsch and Wils 1984), and also on community health broadly speaking (Hirsch and Immediato 1998, 1999).

Recently, the two of us were involved as consultants in a project that addressed the subject of CIM directly. SD models were developed to support the planning of a program that is putting into practice, for an entire community, the ideas for quality improvement described in the Institute of Medicine report, as well as efficiency improvements through IDCOP (Homer et al. 2003). This "Pursuing Perfection" program is unprecedented in its scale—being applied to a population of over 100,000 people in Whatcom County, Washington State—and has attracted national and international attention. Its focus at present is on two chronic diseases, diabetes and heart failure, with hopes for future expansion to all other major categories of chronic illness.

The SD work described above has led us to a set of ideas about the dynamic interplay of CIM with issues of service capacity, demand, and provision. (See Hirsch and Homer 2004 for more on these ideas, as well as their connection to the critical issue of health care access.) We will have more to say about these ideas in the modeling examples presented below, but for now we will present them as a list of propositions:

- The clinical literature demonstrates that "ideal care", achieved by applying various tools of CIM at different stages of the disease chain, can significantly reduce rates of disease onset, progression, and complications, including life-threatening acute episodes. These tools include risk management, detection screening, and disease management. All of these tools must be applied, and applied effectively, in order to make significant and lasting headway against the health and monetary costs of chronic illness.
- CIM quickly achieves cost savings for the elderly portion of the population through the reduction of complications, but takes a much longer time to achieve cost savings for the non-elderly population through the reduction of disease prevalence. Some private insurers who cover a primarily non-elderly population may thus be resistant to the idea of reimbursing patients and providers for the added costs of a CIM program.
- CIM generates cost savings over a period of several years, but requires significant up-front costs that must be paid for somehow. If these costs are laid entirely on the shoulders of health care providers, they may be unwilling to participate in CIM, or to allocate sufficient resources to make it a success. Furthermore, if CIM is effective, patients are likely to have

fewer acute episodes and providers may suffer reduced revenues that create additional financial hardships and make it harder to repay these up-front investments. (Acute visits are more procedure-intensive than routine visits, and are typically reimbursed at a higher rate, particularly under fee-for-service arrangements.)

- CIM also generates increased workload for primary care providers for a number of years, as they endeavor to bring their patients' conditions under control. One of two things may happen if sufficient capacity is not created or freed up through efficiency measures to handle this increased workload. Providers will either cut back or eliminate their participation in CIM in order to balance their workload—thereby jeopardizing the CIM program—or will attempt to do more than they can with limited capacity.
- If providers try to do more than they can with limited capacity, several problems may develop. First, appointment backlogs will grow—for both the chronically ill and the acutely ill—and patients will encounter longer waiting times, more limited access to care, and perhaps a lower quality of care delivered by overworked providers. Second, the stress of overwork may reduce provider productivity, and may lead to and increased rate of provider retirements or departures. Thus, the stress of overwork may reduce capacity and further worsen the situation.
- Investments in clinical efficiency both free up provider time to handle the initial increase in workload due to CIM and reduce the unit costs of provider services. Reduced unit costs can help to offset the effects of (1) any unreimbursed costs of CIM, including investments in associated information systems, and (2) any loss of revenues due to the reduction of acute episodes that follows effective implementation of CIM. Without initial investments in clinical efficiency, providers implementing CIM may thus see reductions in income. These income reductions, in turn, may cause providers either to abandon the CIM program, or to further postpone the very investments in efficiency measures, such as IDCOP, which could permit continued profitability.
- Even with improvements in efficiency and capacity, it may be necessary to adjust reimbursement formulae so that savings generated by CIM programs are equitably distributed to providers and insurers making the investments in those programs. Otherwise, the dynamics of chronic illness may cause the benefits to fall disproportionately to certain players (e.g., the Medicare program for older people) and not to others. This was addressed in the Whatcom County work mentioned above, but is beyond the scope of this paper.

These propositions are a helpful first step for thinking about the planning of CIM programs and investments in improved clinical efficiency or other types of capacity expansion. But they are of limited use for answering specific questions about particular patient populations, chronic illnesses, and providers or delivery systems. This is where SD modeling and simulation can contribute. The remainder of this paper is devoted to two examples. In each example, an existing SD model is applied to analyze the impact of different approaches to CIM and capacity or efficiency management in light of the above propositions. In the first example, the Whatcom County diabetes model is used to evaluate the importance of a community nursing resource, known as clinical care specialists, to the success of the Pursuing Perfection (P2) program. In the second example, the community health simulator known as the Healthcare Microworld is used to evaluate the importance of efficiency investments made in concert with or as prelude to a CIM program.

Example #1: Managing a Community Nursing Resource for Diabetes Management

The diabetes and heart failure models developed for Whatcom County identified two factors that could mitigate the ability of the P2 program to bring patients successfully under disease management: drug affordability, and the availability of clinical care specialists (CCSs) to keep up with demand for their services. In our previous writing we presented a simulation illustrating the potential impact of drug affordability, but did not explore the impact of CCS availability (Homer et al. 2003).

The CCSs are nurses, paid through P2 general program funds, to whom all participating physicians may refer patients for counseling, education, and navigation through the health care system.[1] Through such efforts, a CCS may help a patient achieve clinical control within the first few months of initial referral, and subsequently to maintain control. Patients at the later complicated stages of disease are more likely to need the assistance of a CCS for establishing and maintaining control than patients who are not yet suffering complications. In the case of diabetes, about 25% of people with early-stage uncomplicated diabetes benefit from the assistance of a CCS, whereas about 75% of people with complicated diabetes benefit. A person with diabetes who needs the assistance of a CCS will typically require about six hours of CCS time during the first few months to assist in bringing their blood glucose (and blood pressure and blood lipids, if those are also problematic) under control, but only about 1.5 hours of additional CCS support per year thereafter to help maintain control.[2]

Figure 1 presents an overview of the portion of the Whatcom County diabetes model, as recently modified, which addresses the CCSs and their impact on disease control.[3] The stock of CCSs may be increased by hiring and reduced by layoffs, changing (within budget-determined limits) in response to the demand for their services from new referrals seeking control and older referrals maintaining control (bottom of the diagram). At the left edge of the diagram is a

---

[1] The capacity of the physicians themselves was not modeled as a mitigating factor, because diabetes and heart failure each represent only a fraction of the workload of a typical primary care physician. Knowing these fractions, one may calculate the relative impact of a given increase in workload due to the P2 program, for either diabetes or heart failure patients, on a physician's practice overall. Estimates led us to conclude that these impacts of individual diseases were unlikely to be very great, and could be accommodated assuming successful implementation of IDCOP, which is an unquestioned component of the P2 program. Consequently, we did not pursue the issue of physician capacity further in the diabetes and heart failure models. However, we would want to revisit this issue if the P2 program were ultimately expanded to cover all chronic diseases, as is envisioned.

[2] Throughout the time of our study, there were two CCSs in Whatcom County, Nancy Stothart and Connie Golas, both of whom work with diabetes and heart failure patients. Nancy and Connie, together with the P2 program director Mary Minniti, helped us to understand the activities of the CCSs, and provided the parameter values for the CCS portion of the model. We would like to acknowledge their significant contributions to the modeling effort.

[3] For the purposes of this paper, we have made two significant changes to the diabetes model. First, the number of CCSs is now formulated as an endogenously adjusting stock, rather than as an exogenous time series. Second, we have introduced negative feedback from a CCS shortage to the decision by physicians to refer additional patients to the CCSs. Absent such feedback, the CCS referral backlog could, in the event of a persistent CCS shortage, increase without limit and without any response from physicians. In the newly revised model, it is still possible for a referral backlog to develop, but the demand feedback mechanism keeps it within reasonable bounds.

simplified view of the stock-and-flow structure of patients who, by achieving control, may flow from uncontrolled to controlled status.[4] Patients achieving control are of two types, those who do not need CCS assistance to achieve control (left side of the diagram), and those who do (bottom), both of which originate from patients seeking control (middle).

The number of patients seeking control is determined by the number of patients still uncontrolled, the fraction of patients whose primary care physician participates in the P2 program, the fraction who can afford the drugs needed to achieve control, and by the adequacy of the CCSs, as perceived by physicians, to assist newly referred patients seeking control. The adequacy of CCS hours for new referrals, in turn, is found by subtracting from total CCS hours available the hours needed to maintain existing patients, and then dividing by the hours needed to help new referrals to achieve control. If CCS hours are not fully adequate to handle new referrals, then some of these referrals will accumulate in a referral backlog stock (bottom of the diagram), where they continue to add to CCS demand until CCS time becomes available to see them.

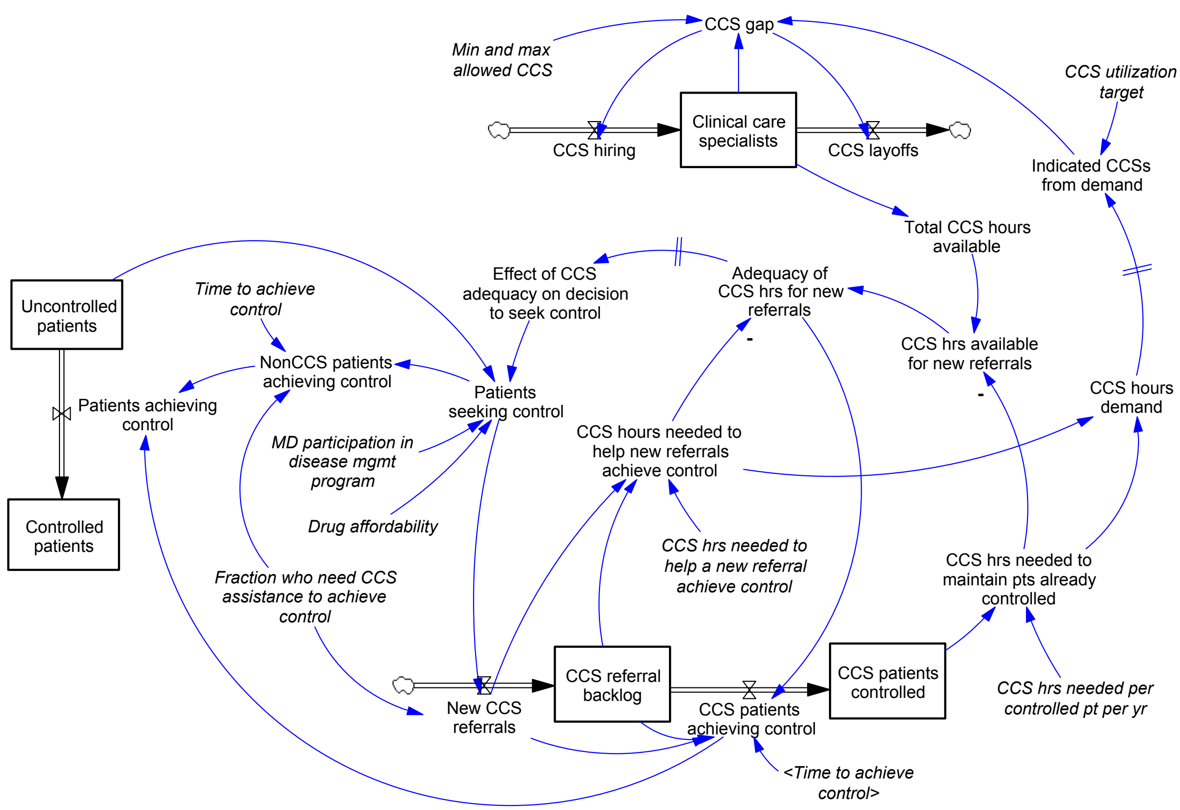

Figure 1. Overview of model structure for clinical care specialists and their impact on disease control

[4] The model includes uncontrolled and controlled patient stocks for each stage of the disease, with flows connecting one stage to the next. It also includes at-risk population stocks, flows of aging from non-elderly to elderly categories, and age-adjusted outflows of death from every population stock in the model (Homer et al. 2003).

The structure in Figure 1 contains a few feedback loops of note, as follows[5]:

- In the most basic goal-seeking loop, uncontrolled patients seek and then achieve control.
- But as the CCSs bring more patients under control, the need to spend continuing time with these maintenance patients increasingly takes CCS time away from the intake of new referrals, and threatens to reduce the flow of patients achieving control.
- In response to that threat (and to the extent the budget permits), the number of CCSs is adjusted to meet demand so that referred patients may achieve control, rather than accumulating in an ever-growing backlog.
- However, if that response is insufficient, CCS referral backlog can build up as a result and perceived CCS adequacy will slip. Physicians, in turn, will start to cut back on the extent to which they put effort into helping diabetes patients to seek and achieve control.

Using the modified Whatcom County diabetes model, we have performed a series of three 20-year simulations that illustrate the dynamics just described. The run names and assumptions are:

- No Program: P2 program not implemented. The controlled fraction of known diabetics remains fixed at baseline values—40% of non-elderly diabetics, 46% of elderly.
- CCS Variable: P2 program implemented starting year 1 of the simulation, allowing for substantial increase in controlled fraction. One CCS is hired initially, but more may be hired as needed.
- CCS Fixed: P2 program implemented as in CCS Variable, but budget allows for no more than one CCS.

Results of these simulations are presented in Figures 2 to 5. Figure 2 shows how, in the CCS Variable run, the number of CCSs (blue line) rises to 3 by Year 5, then comes down to 2 for Years 6 to 12, then increases again to 3 for years 13-19. This stepwise behavior is driven by "Indicated CCSs" (thin green line), which describes the number of full-time equivalents that would be required to satisfy recent demand given a target of 75% average CCS utilization. Indicated CCSs for the CCS Variable run grows through Year 4, subsides for a few years, then resumes its growth. Figure 2 also shows the corresponding outputs for the CCS Fixed run. The number of CCSs in this run is, of course, fixed at 1. But the indicated number of CCSs based on demand is even greater in the CCS Fixed run than that it is in the CCS Variable run; this perhaps unexpected result reflects the growth of a substantial referral backlog in the CCS Fixed run.

Figure 3 presents graphs of new CCS referrals and referral backlog that clarify the demand picture. New CCS referrals rise through Year 4 and then quickly subside in both the CCS Variable and CCS Fixed runs. The first four years of this pattern (including the annual peak-and-decline) reflect the assumed stepwise adoption of the P2 program (Homer et al. 2003), during which time a large cohort of uncontrolled patients are brought under disease management. These first four years of working through an existing stock of potential demand are followed by a more steady-state situation in which the gradually growing incidence of diabetes (due to population growth and aging) leads to a continuing flow of new candidates for disease management.

[5] Although the loops are all easily described, some are difficult to discern in Figure 1, which is more of an influence diagram than a causal-loop diagram. For this reason, and to reduce clutter in the diagram, no feedback loop symbols are included in Figure 1.

Figure 3 shows that new CCS referrals are less in the CCS Fixed run than in the CCS Variable run starting in Year 3, increasingly so as the run develops. The reduction in new demand in the CCS Fixed run occurs because CCS adequacy is compromised, and physicians respond negatively to this situation by cutting back on their commitment to the P2 program, as

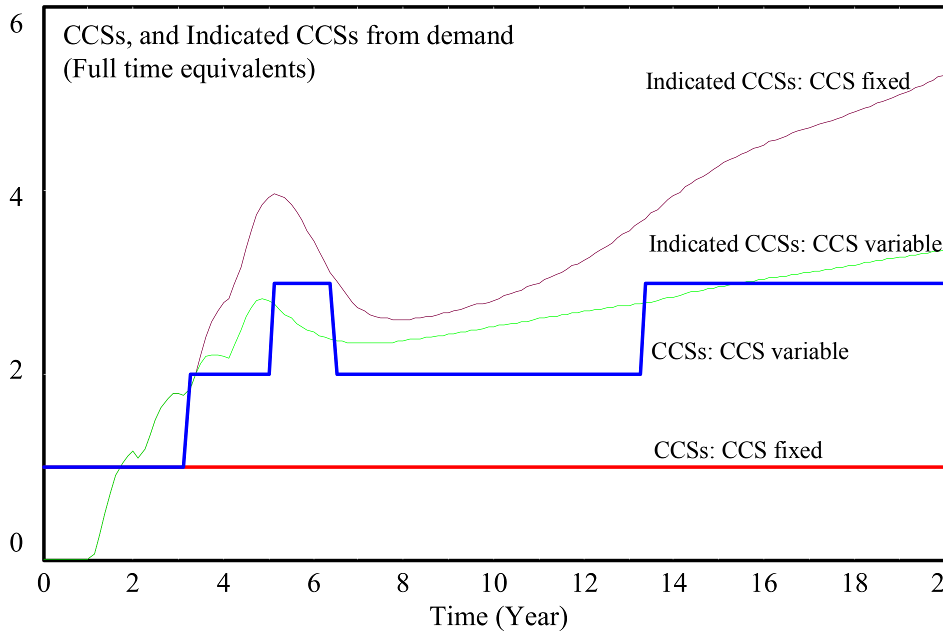

Figure 2. Comparison of Fixed capacity and Variable capacity simulations, in terms of Clinical care specialists and Indicated CCSs from demand

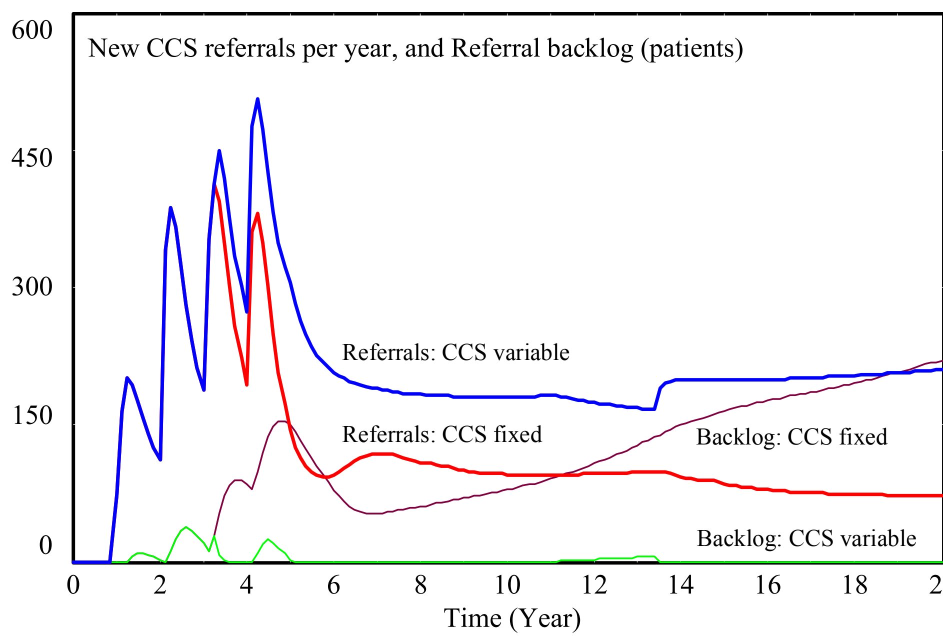

Figure 3. Comparison of Fixed capacity and Variable capacity simulations, in terms of New CCS referrals and CCS referral backlog

described above. The growing inadequacy of CCS time in the CCS Fixed run is reflected in Figure 3 in the growth of a substantial referral backlog. Some referral backlog develops during the initial years of P2 implementation (Years 1 to 4) in both runs. But whereas it is rather quickly resolved in the CCS Variable run through the hiring of additional CCSs, it becomes a persistent and growing problem in the CCS Fixed run. This is why the Indicated CCSs from demand in Figure 2 grow larger in the CCS Fixed run than they do in the CCS Variable run: Although the flow of new demand is actually reduced in the CCS Fixed run, the demand backlog is becoming ever greater, and an increasingly large number of CCSs would be required to eliminate that backlog.

Figure 4 shows the contribution of CCS efforts to the fraction of known diabetic patients under control. In the No Program run, this fraction remains at about 43% throughout the 20-year run, rising (but less than a percentage point) due to an increasing elderly fraction of the population. In the CCS Variable run, the fraction under control rises rapidly during the years of P2 adoption, settling at 71-72% by Year 7 and thereafter. This amount of improvement in control reflects assumptions about the potential impact of the P2 program and the effect of drug affordability. Under the CCS Fixed run, the controlled fraction rises exactly as in the CCS Variable run for the first two years of P2 adoption, but then diverges from this pattern due to the growing inadequacy of CCSs to handle new referrals. This inadequacy, as we have seen, affects both the seeking and the achieving of control. The controlled fraction reaches a peak of only 60% in Year 5, and then begins a long decline, falling to 47% by the end of the run.

The adverse effect of fixed CCS capacity is not quite as bad as it appears in Figure 4, because the denominator—the number of known diabetics—is greater with the P2 program than without, due to the community screening component of P2 (see Homer et al. 2003). Nonetheless, the adverse effect is significant, and what is most notable is that the inadequacy of CCSs adversely affects control not only for patients who need a CCS, but for all patients who might be better managed under the P2 program. In other words, an inadequacy of specialized CIM capacity targeting a portion of the candidate patient population may spill over broadly to affect provider commitment to the program, and thereby undermine the value of the CIM program for all candidate patients.

Figure 5 concludes the analysis, showing the annual number of deaths in Whatcom County due to complications of diabetes, and the impact of the P2 program and CCS efforts. Due to the increased controlled fraction, deaths are reduced relative to the No Program scenario within a short time after P2 program adoption in both the CCS Variable and CCS Fixed scenarios. But the effectiveness of the program in reducing deaths becomes significantly less under CCS Fixed than it is under CCS Variable, due to the eroding extent of control as described above. For example, in Year 8, deaths are reduced by 35% under CCS Variable, but by 21% under CCS Fixed; in Year 20, the corresponding reductions are 40% and 19%.

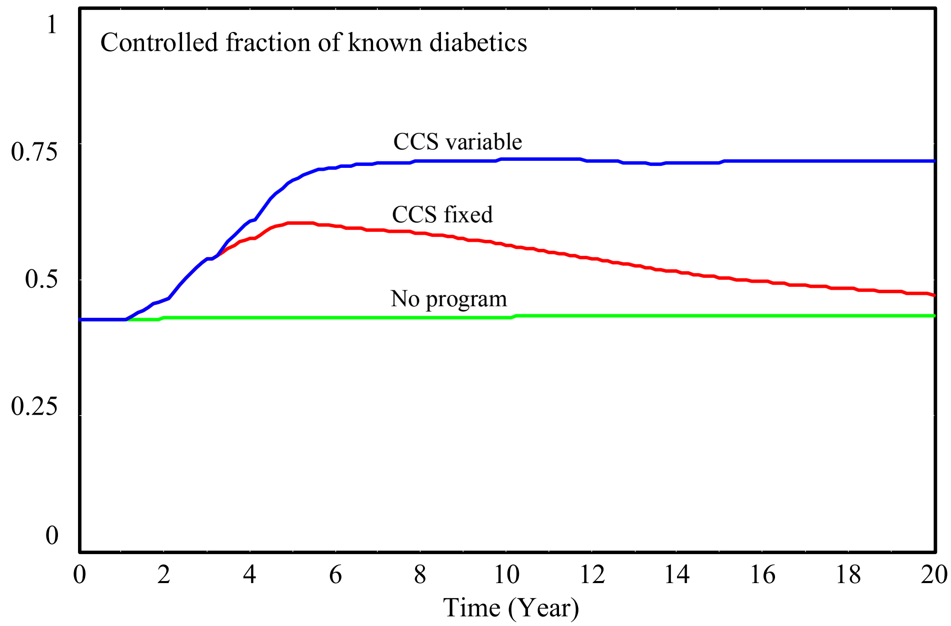

Figure 4. Comparison of No program, Fixed capacity, and Variable capacity simulations, in terms of Controlled fraction of known diabetics

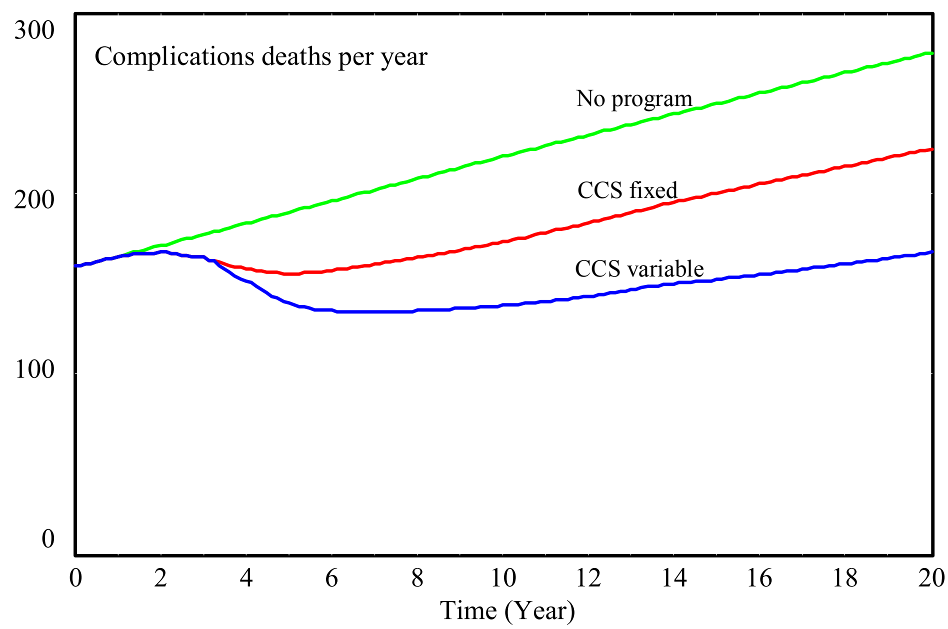

Figure 5. Comparison of No program, Fixed capacity, and Variable capacity simulations, in terms of Complications deaths

Example #2: Evaluating Health Plan Investments in Clinical Efficiency and Chronic Illness Management

We will now present a series of simulations that demonstrate the importance of coordinating improvements in CIM with enhancements in a practice's overall efficiency. The "Healthcare Microworld" used for performing these simulations was developed for a consortium of healthcare organizations brought together by Innovation Associates and the New England Healthcare Assembly in the mid-1990s. (See Hirsch and Immediato 1998, 1999, for a more complete description.) These organizations sponsored the work because they wanted their managers to get a better understanding, via simulated experience, of the many changes going on in health care at that time. The Microworld was developed over two years through an intensive collaboration with teams from the 12 organizations that sponsored the work. Though not based on historical data from any one organization or community, the Microworld was based on extensive data from the field and the collective experience of 100+ people collaborating on its development. The rich set of connections in the Microworld help to show the pitfalls of implementing strategies as if they are simply additive rather than carefully coordinated in their implementation.

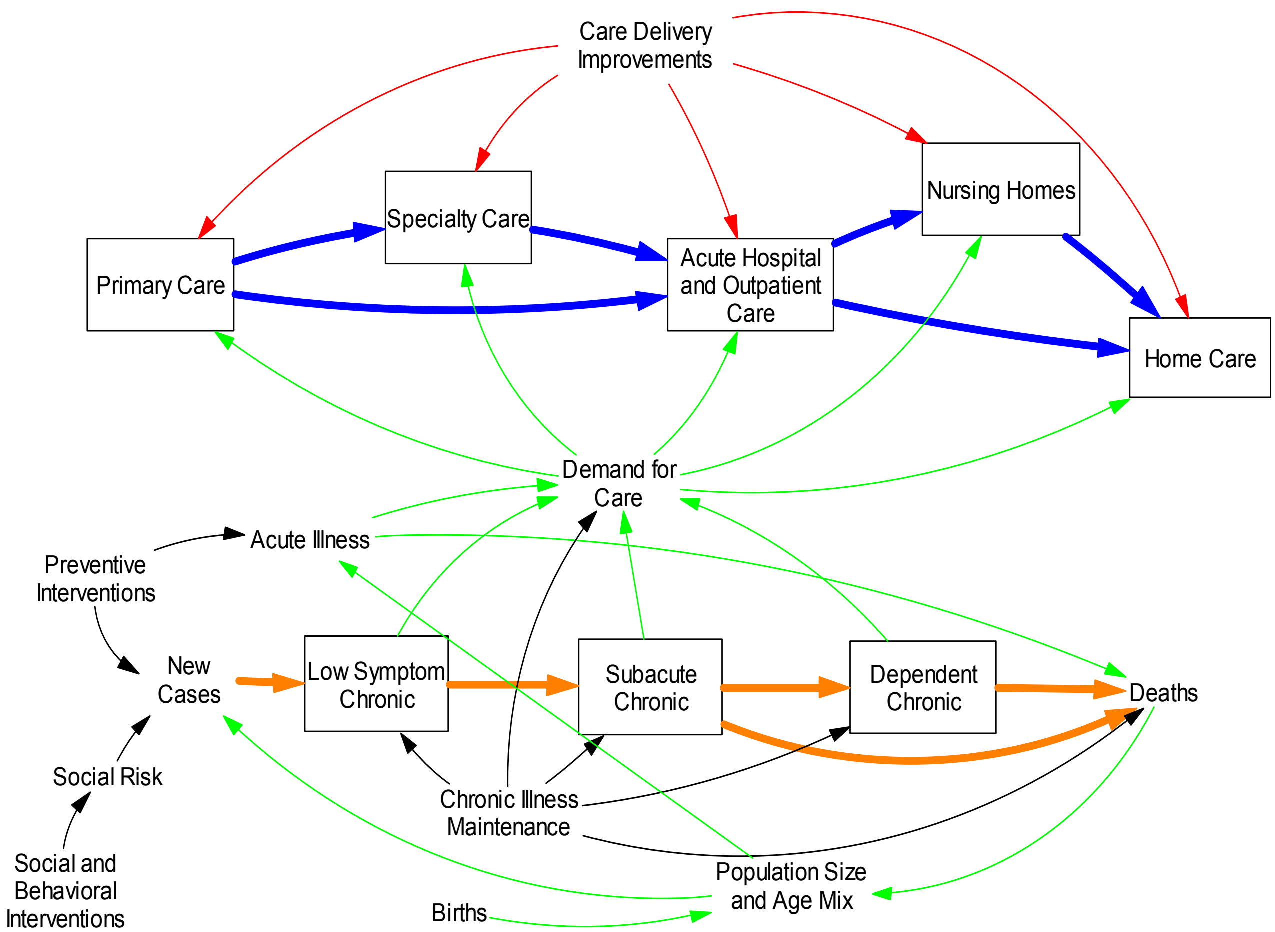

Figure 6. Overview of model underlying the Healthcare Microworld

Figure 6 provides an overview of the high-level relationships in the model underlying the Microworld. Demands for care generated by both acute and chronic illness arrive at the different sectors of the health care system. Patients receiving primary care may also be referred on for specialty care and those receiving specialty care may need to go to the hospital for inpatient care or outpatient procedures. Patients discharged from the hospital may need long-term care in nursing homes or home care delivered by visiting nurses and therapists.

The demand for care is determined by the size of the population being served and rates of acute and chronic illness. Health improvement programs can affect rates of illness. These include preventive programs that can reduce the incidence of acute illness and new cases of chronic diseases. Chronic illness maintenance programs (also referred to as disease management) slow the progression of chronic illnesses and reduce the frequency of acute episodes in the same manner as the controlled status in the diabetes model described earlier. Longer-term interventions such as job training can reduce social risk that ultimately leads to illness.

Users of the Microworld implement strategies by allocating limited resources among a number of interventions including those that improve the functioning of the delivery system (e.g., investments in clinical information systems, process improvement projects) and health improvement programs that can eventually reduce the demand for care. Successful interventions of either kind, delivery system or health status improvement, can lead to improved financial performance and more resources to invest later. Ineffective interventions and competition from other health plans can cause funds available for investment to dwindle. Figure 7 shows how these kinds of improvements interact with each other.

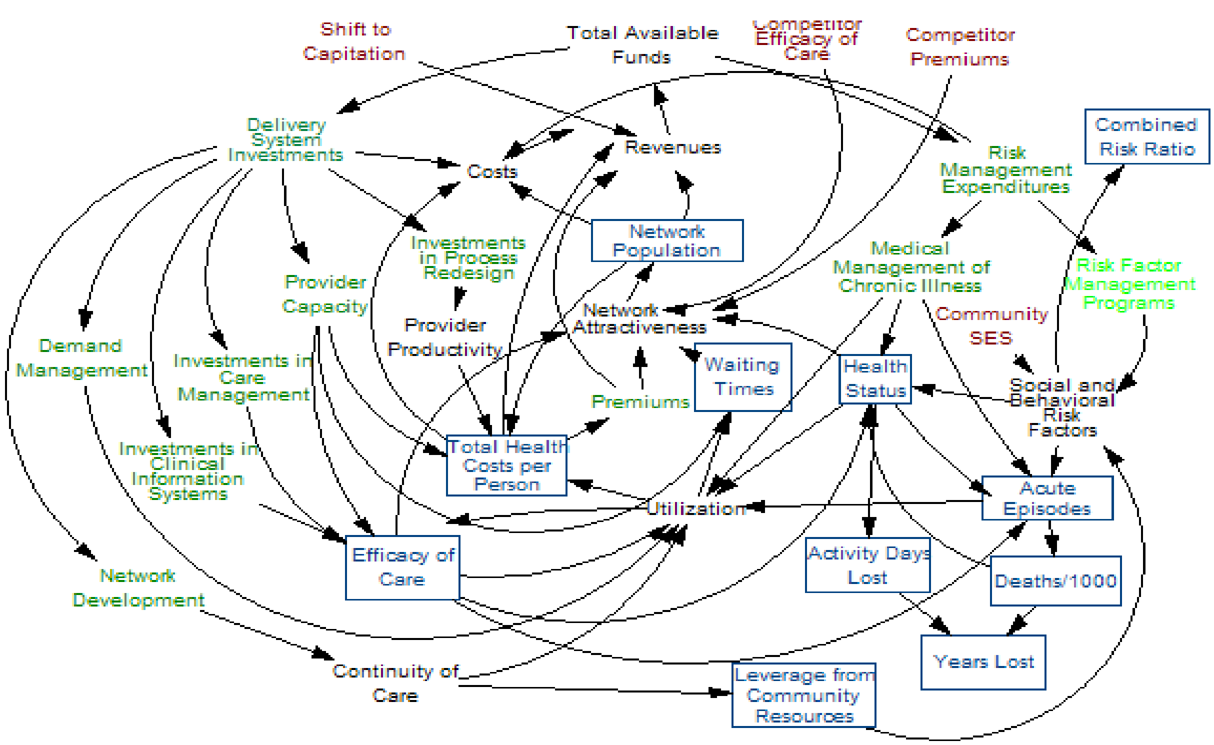

Figure 7. Interrelationships of delivery system and health improvement interventions

As shown in Figure 7, investments in care delivery help to lower cost, which allows the Plan to charge lower premiums and make itself more attractive to patients. Lower cost and higher net income also allows the Plan to make additional investments. Investments in such things as clinical information systems also enable the Plan to increase the efficacy (quality and effectiveness) of care, which helps to attract more patients, lower utilization and cost, and improve patients' health status. Investments that increase capacity and productivity help to keep waiting times down and also make the Plan more attractive. Investments in health status through prevention, chronic illness management, and other programs serve to lower utilization and cost. They also make the Plan more attractive by allowing it to boast of healthier patients.

Simulations with the Microworld provide some useful insights about the effects of delivery system constraints on chronic illness management. Its representation of chronic illness is not as finely developed as in the model of diabetes described earlier since it is attempting to reflect all chronic illness in aggregate. For example, there is no separation of the chronic population into segments whose disease (like diabetes) is uncontrolled and controlled or who receive "usual" and "ideal" care as was done in a model of heart failure (see Homer et al, 2003). Representing all chronic illnesses in aggregate also makes it difficult to represent the many important differences that distinguish chronic illnesses from each other. However, by representing all chronic illnesses in their totality, the advantage offered by the Microworld is that it allows us to match demand generated by chronic illness with the capacity of the system and understand the potential effects of too little capacity. The following paragraphs describe a set of explorations with the Microworld that would lead to such an understanding.

After exploring a number of different options, users of the Microworld may identify a strategy consisting of measures such as process improvement and investments in better clinical information systems that, by improving clinical efficiency, reduce costs per encounter and increase perceived quality of care. This measure is implemented at the beginning (1997) of the ten-year simulation. The cost reductions enable the plan to lower premiums by 20% and still remain profitable. The simulations suggest that lower premiums and higher perceived quality would let the plan expand its market share and build a strong financial base; see the green graph lines in Figure 8, specifically those for Network Population and Net Income. The health plan has the financial resources necessary to make planned investments in clinical information systems, investments which enable the plan to increase efficacy of care and gain additional savings by treating patients more efficiently. Waiting times for appointments (number of days wait to see a clinician) go up as new members flock to the plan, but remain within an acceptable range.

One might think that since much of the plan's expense is related to the care of chronic illness, a strategy that includes more stringent CIM could improve results by slowing the progression of illness and reducing the frequency of expensive acute episodes in those patients. To test this hypothesis, another simulation may be done that adds more aggressive CIM to the efficiency measures that produced such good results with the previous strategy. In this simulation, patients with chronic illnesses receive several more visits per year and benefit by having fewer acute episodes and slower progression of their illnesses to more severe stages. Just as the previous efficiency measures, the aggressive CIM is implemented at the beginning of the simulation, so that, in theory, it may have the greatest impact. The results of the combined strategy may be seen as the blue graph lines in Figure 8.

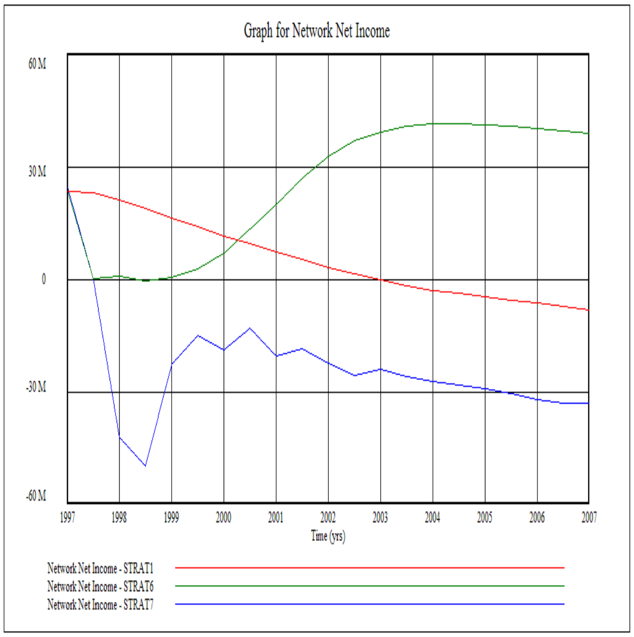

Network Population

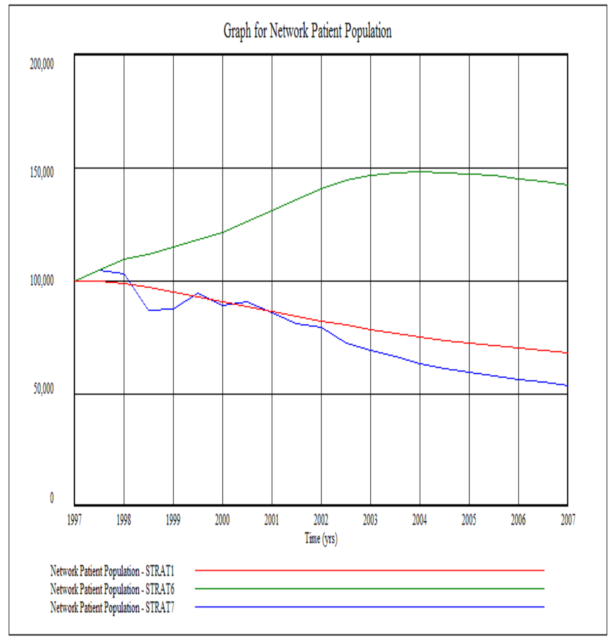

Network Net Income

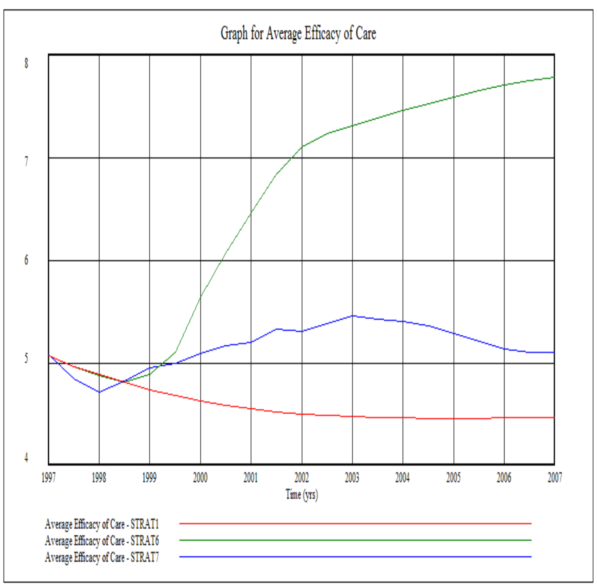

Average Efficacy of Care

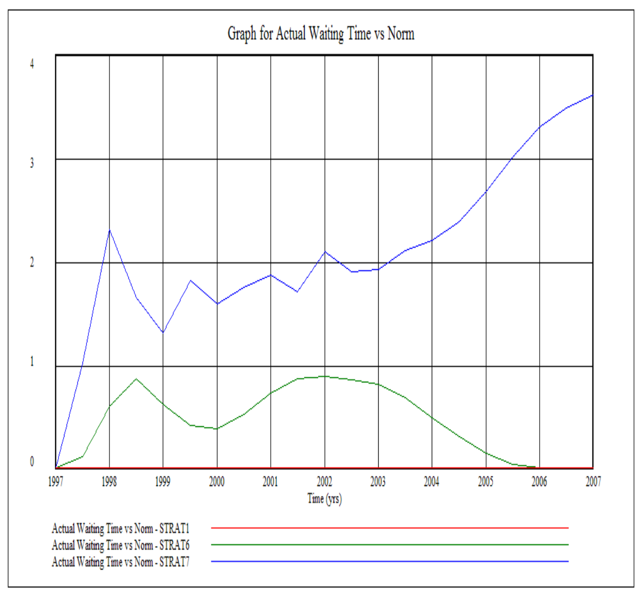

Average Waiting Time vs. Norm

Figure 8. Comparison of results of three simulations: Do nothing (red), Improve efficiency (green) and Improve efficiency with chronic illness management (blue)

Clearly, the impact of the combined strategy is the opposite of what was expected. Network Population and Net Income actually fall below levels produced by the Do Nothing strategy. Efficacy of Care fails to increase very much because the health plan's poor financial condition precludes the necessary investments in clinical information systems.

Why are the results so poor when what looks like a good feature is added to an already successful strategy? The results for waiting times for appointments suggest part of the problem. Doing more stringent management of chronic patients requires more visits that add to the plan's workload. While these would be easily accommodated later after efficiency measures are implemented, implementing chronic illness management right at the beginning of the simulation drives up workloads and waiting times and drives away members who experience reduced access to care. This exodus of members initiates a downward spiral in which the financial consequences of having fewer members causes the plan to lower compensation and do other things that make it lose clinicians even faster than members. As a result, waiting times remain high and encourage more members to leave the plan.

Figure 9 shows in greater detail what happens when the health plan implements intensive CIM without first building sufficient capacity. Average waiting time for the network (red line) increases along with the ratio of demand to capacity for its primary care providers (green) and specialty providers (blue). Though a brief initial peak subsides as the system adjusts (e.g., with shorter visits), the loss of members due to long waiting times has adverse financial effects that, in turn, lead to a loss of providers. This loss of capacity allow the ratio of demand to capacity to rise even though members are leaving.

The black line represents "Out of Plan" costs as members faced with long delays go to outside providers, at great cost to the health plan. These out of plan costs subside as members begin leaving the plan entirely, but grow again toward the end of the simulation as its physicians leave and capacity "melts down". The ratio of demand to capacity (gray line) decreases only for acute care simply because the number of acute beds and outpatient visit capacity are assumed to remain constant, unlike primary and specialty care where poor financial conditions can lead to an exodus of providers.

The graphs in Figure 10 further help to explain what is going on. They show population per physician as well as the ratio of demand to capacity for primary care and specialty physicians. As shown in the left-hand graphs in Figure 10, the population per physician (red lines) increases for both primary and specialty care even as the total subscriber population is plummeting. This indicates that providers are leaving even faster than their patients. The right-hand graphs indicate the increases in waiting times for primary and specialty care increase accordingly. Waits increase to 2.1 months for primary care as the number of primary care physicians decline from 40 initially to 11 at the end of the simulation and to 6 months for specialty physicians as their number goes from 60 initially to 19 by the end. Cost per capita decreases for both primary and specialty care because the remaining physicians are each caring for a very large number of patients. However, this is a false "efficiency" and cannot offset the negative effects of the long waits on numbers of members and, in turn, on revenues.

This simulation demonstrates direct as well as second-order effects of failing to expand capacity before initiating chronic illness programs. Direct effects include poorer service for all patients including those with chronic illnesses. The second-order effect of providers leaving suggests the more subtle, but potentially more severe dangers inherent in operating health care systems beyond their capacity, even temporarily in an attempt to "jump start" a CIM program. Increases in capacity really need to precede chronic illness programs that may, in the short term, increase the demand for care.

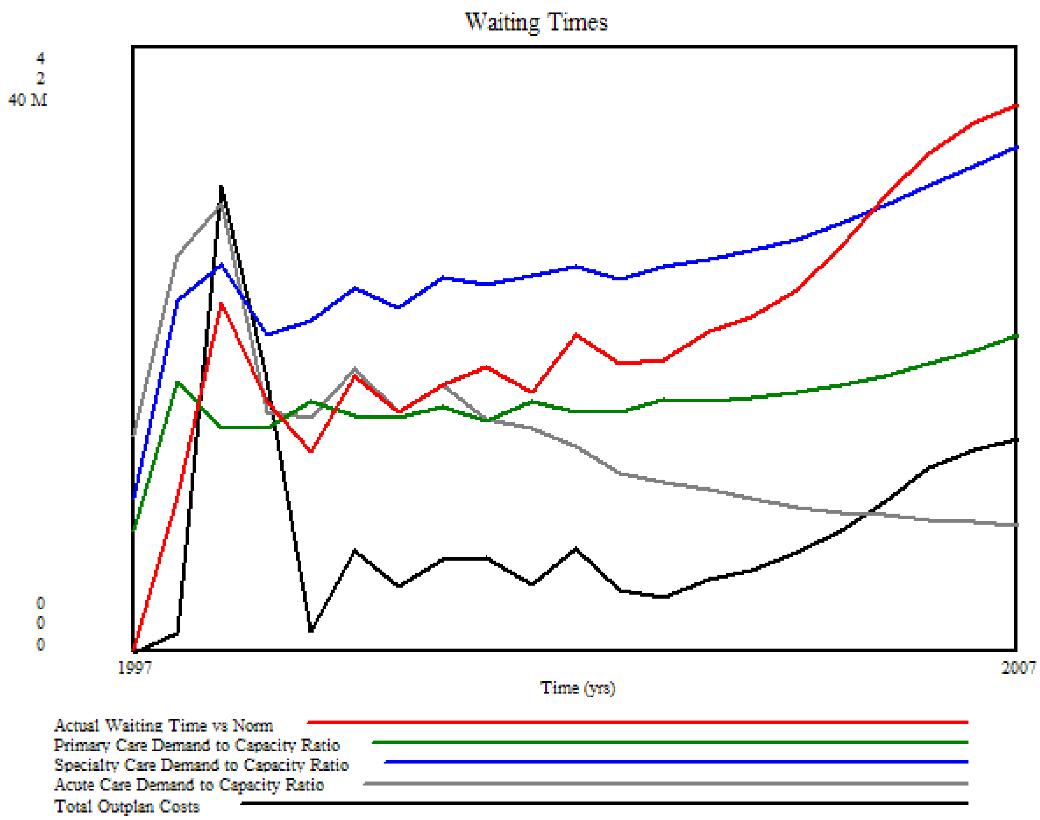

Figure 9. Causes and consequences of longer waiting times for simulation in which chronic illness management (CIM) is implemented without sufficient capacity

Figure 11 shows the feedback loops that drive these adverse effects. R1 is the loop that would drive the expected improvement by reinvesting savings due to acute care costs into further improvements in CIM. However, B2 and B3 undercut the effectiveness of care through increased waiting times and reduced access as increased demand temporarily exceeds capacity. Reduced access limits the impact of the CIM programs. Reduced provider revenues as patients leave in response to the longer waits for care exacerbate this situation by further reducing capacity as providers depart or fail to make the needed investments (B4, B5). The decline in capacity is worsened by an additional reinforcing loop (R2) as providers "burn out" in this high stress situation and more of them leave.

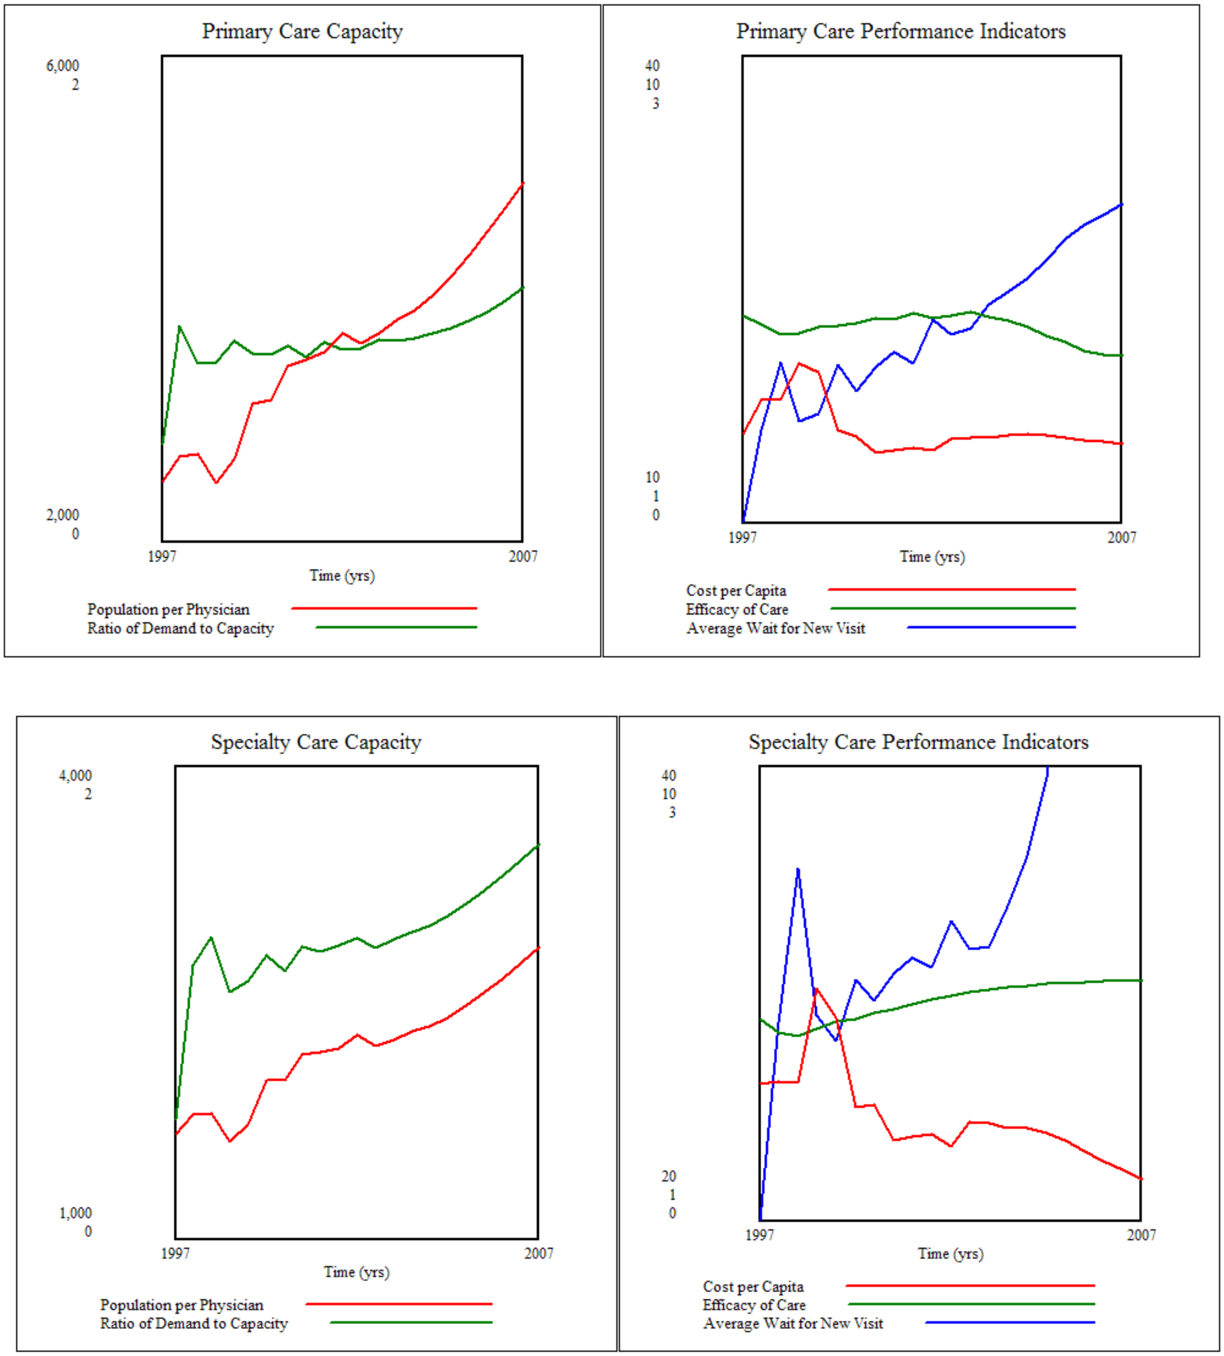

Figure 10. Population per physician, ratio of demand to capacity, and performance measures for simulation in which CIM is implemented without sufficient capacity

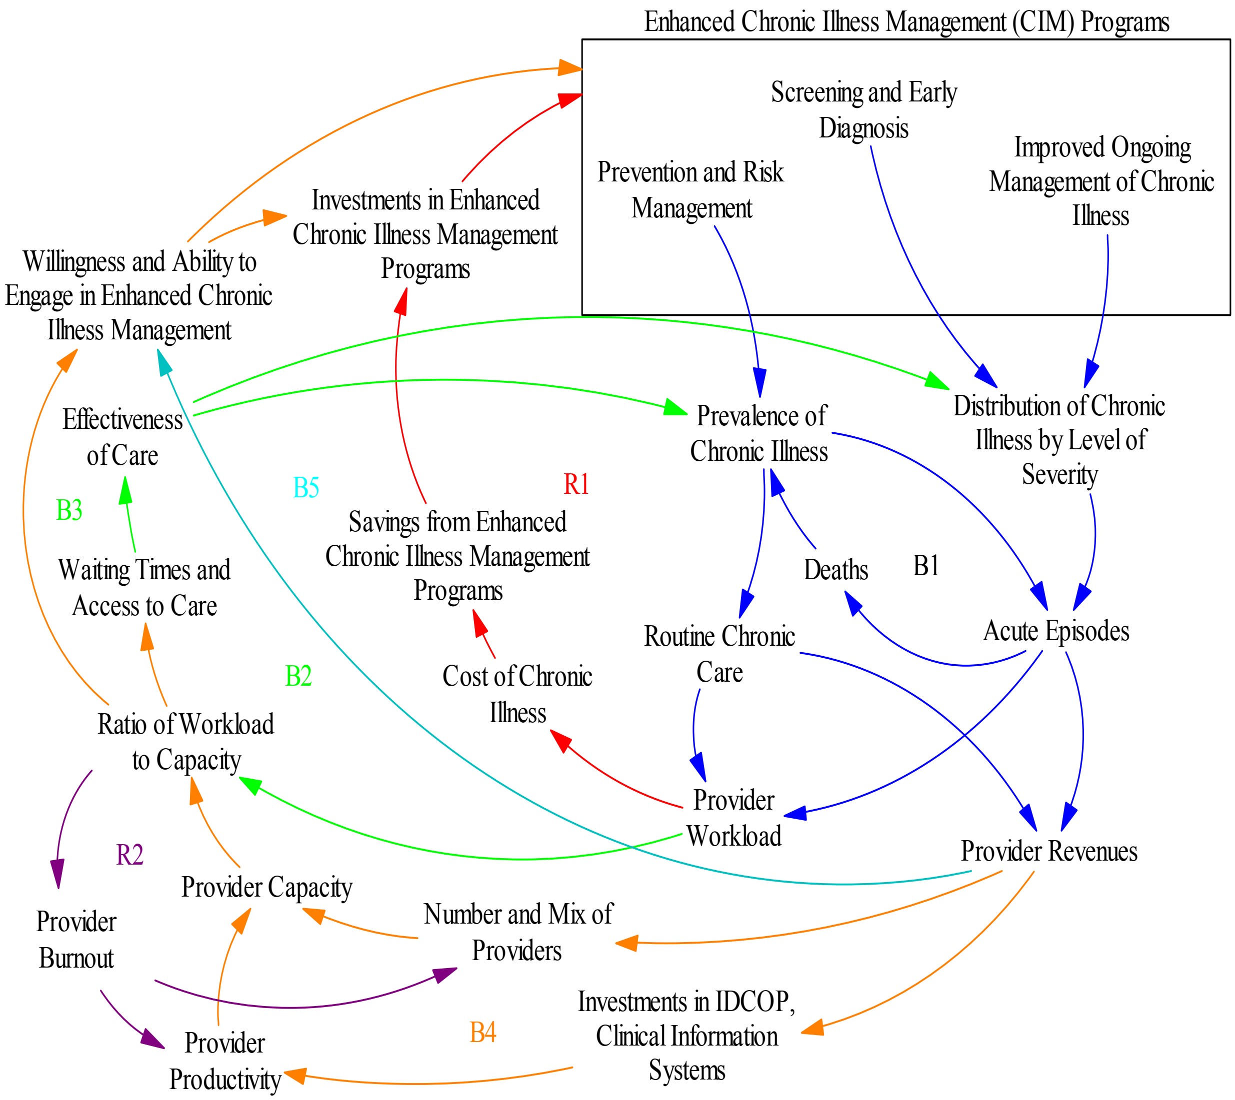

Figure 11. Feedback loops leading to adverse effects of aggressive CIM, if CIM is implemented before delivery system efficiency and capacity have increased sufficiently

Was implementing chronic illness management a bad idea? There's nothing wrong with more stringent CIM itself. It was just implemented prematurely without regard for the impact it would have on the rest of the system. The extra workload it entails could not be handled without major improvements in capacity. CIM has to be implemented in a manner that is coordinated with other improvements. Figure 12 compares the results of the disastrous simulation (blue) with another one (gray) in which the implementation of CIM is simply postponed for a few years.

As shown in Figure 12, this delay allows the improvements in cost and quality to have their impact and create the capacity to accommodate the increased workload of taking better care of patients with chronic illnesses. Waiting a few years also enables the plan to build up its membership and financial reserves, and to make the investments that improve efficacy of care, though not quite as much as in the simulation with efficiency improvements alone (green). Waiting times for appointments jump briefly around 2001 as CIM is implemented, but the higher capacity that has developed by that time renders this jump only a temporary peak and waiting times fall to an acceptable level. The downward spiral seen in the disastrous strategy does not occur. Waiting to implement CIM seems to make a big difference and turns an unsuccessful strategy into an effective one.

The results in the last strategy, with delayed implementation of CIM, are quite acceptable from a business standpoint even though they do not seem to be nominally as good as the one with efficiency improvements alone. Why implement CIM at all? To understand the benefits fully, we must look beyond health care cost savings and profitability and examine patient outcomes as well. The payoff can be seen in the graphs in Figure 13 that display the health status impact of the strategies. Aggressive CIM, implemented a few years after the efficiency measures, produces much lower rates of mortality and activity days lost (compare gray to green). These results have real value for members and their employers and suggest the benefits of implementing a combined strategy. But the benefits are not realized when aggressive CIM is implemented early in the simulation (blue), pointing again to the need for a coordinated, phased strategy.

These results from the Microworld are intended to be illustrative of the value of combining chronic illness and health care delivery perspectives rather than being predictions of what might happen. As indicated earlier, the Microworld was not based on historical data from particular organizations. However, it was based on generic data from the health care field (e.g., from the National Center for Health Statistics in the US, from household studies in 25 communities) and the collective experience of 100 or so health care providers who were involved in its development. Its output was also found to be plausible by several hundred people who used it in a number of workshops. Working from this rich base of data and experience, the results suggest that a model of the delivery of chronic illness care can be very helpful for planning.

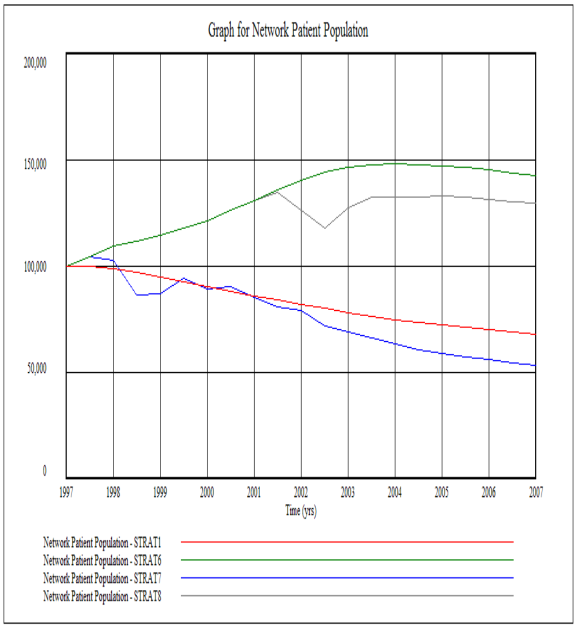

Network Population

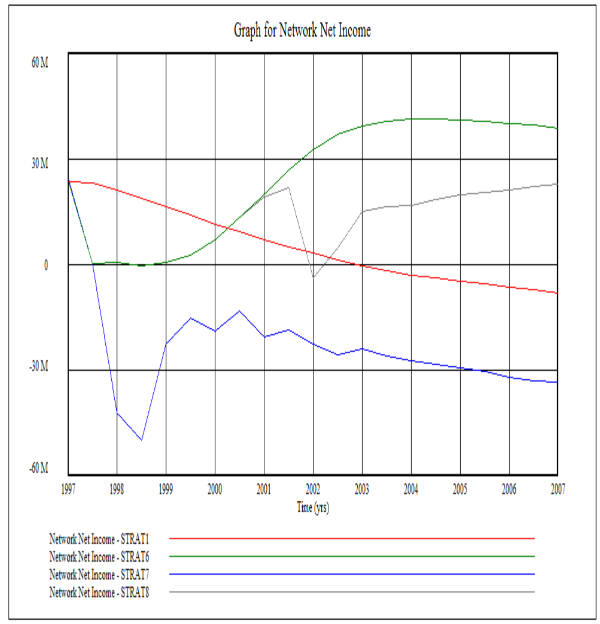

Network Net Income

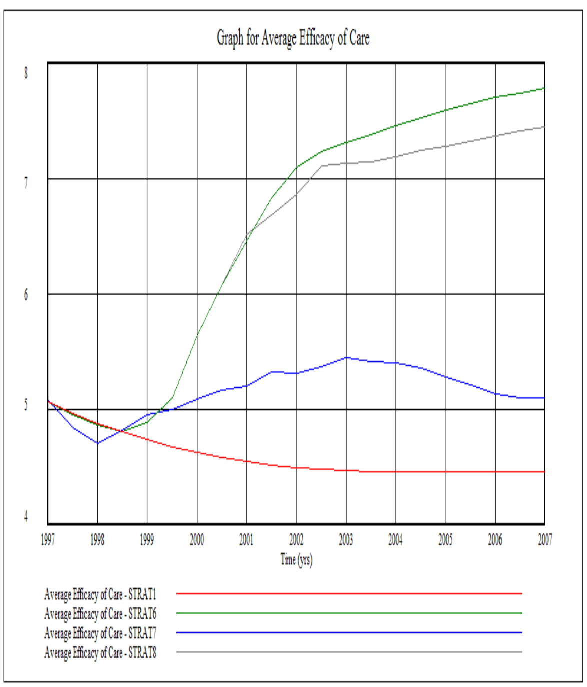

Efficacy of Care

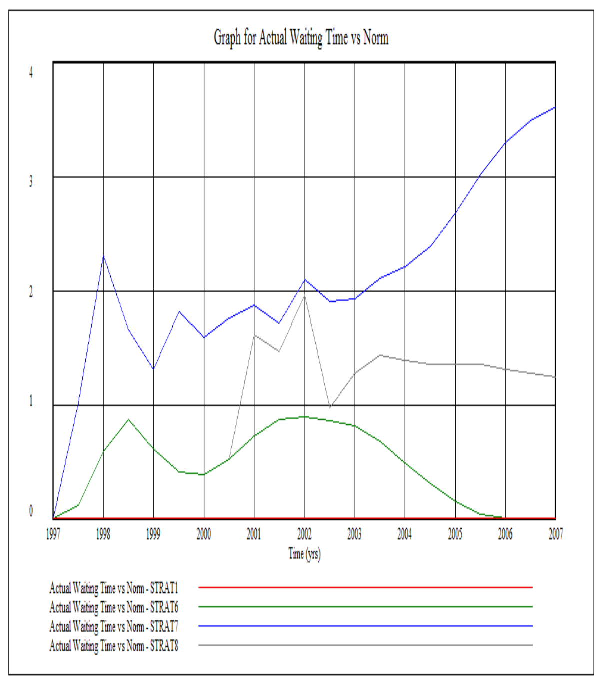

Waiting Time vs. Norm

Figure 12. Comparison of results of four simulations: Do nothing (red), Improve efficiency (green), Improve efficiency with CIM (blue), and Improve efficiency with delayed CIM (gray)

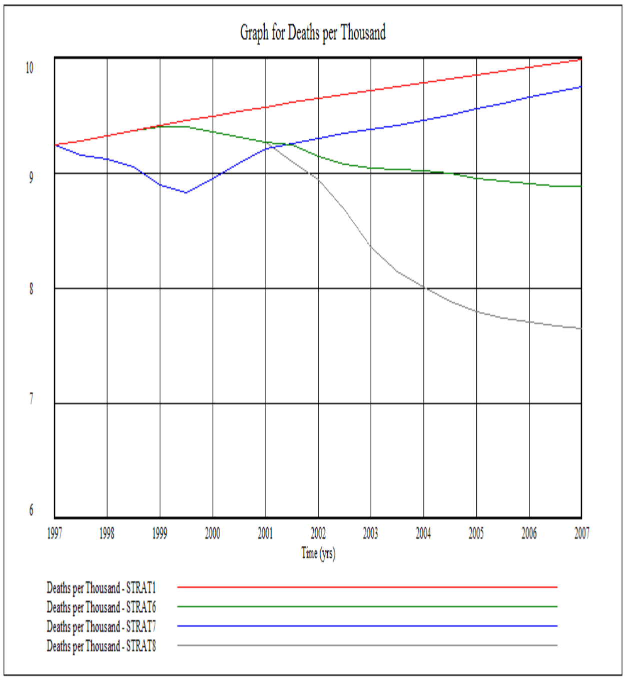

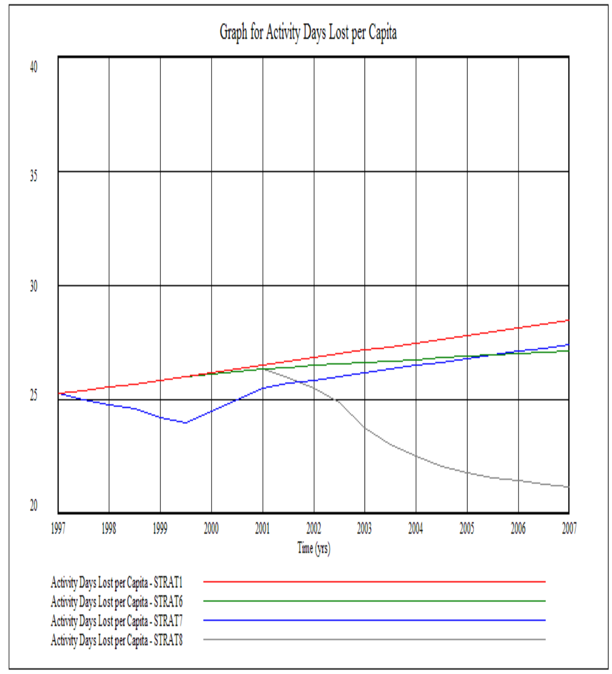

Deaths per Thousand | Activity Days Lost per Capita

Figure 13. Comparison of four simulations in terms of health status measures:
Do nothing (red), Improve efficiency (green), Improve efficiency with CIM (blue),
and Improve efficiency with delayed CIM (gray)

Conclusion

The two applications presented in this paper illustrate the valuable insights that can emerge when chronic illness management (CIM) programs are planned in the context of the delivery systems within which they will be implemented. Implementing CIM without understanding its effect on delivery systems can compromise the impact of CIM and, in extreme circumstances, threaten the delivery systems' viability. Reduced provider revenues that may ultimately result from successful CIM programs must also be planned for and offset by more efficient operations and by attracting additional patients to participating providers.

Combining models of the dynamics of chronic illness in populations with models of the dynamics of delivery systems can create powerful tools for planning chronic illness care. Improvements are possible in both of the models we have presented as examples. The models of diabetes and heart failure created in the Whatcom County work are good starting points, but need to be expanded to other chronic illnesses and so that they may deal with interactions among these illnesses. For example, efforts to reduce hypertension in a population may affect several illnesses at once. Depression is a chronic illness by itself and may also be a consequence of other chronic illnesses. Untreated, depression may be an impediment to getting patients to take the active role in their treatment necessary for effective CIM. Modeling interactions among chronic illnesses will provide a more realistic picture of their impact on delivery systems.

Similarly, delivery systems can be modeled in greater detail than was done in the Healthcare Microworld, with the inclusion of characteristics specific to selected real-world cases. Multiple specialties and different primary care settings can also be represented. A model could also represent the various facets of IDCOP and other programs designed to improve the performance of delivery systems, and the effects of clinical information systems and other technologies for improving both efficiency and patient outcomes.

A model combining chronic illness dynamics with the details of delivery systems performance could raise health program planning to a new level, paving the way for more economical services, better care for patients, and better working environments for providers.

References

Hirsch GB, Homer J. 2004. Integrating Chronic Illness Management, Improved Access to Care, and Idealized Clinical Practice Design in Health Care Organizations: A Systems Thinking Approach. *Third International Conference on Systems Thinking in Management*, Philadelphia, Pennsylvania.

Hirsch GB, Immediato CS. 1998. Design of Simulators to Enhance Learning: Examples from a Health Care Microworld. *Proceedings of the International Conference of the System Dynamics Society*, Quebec City, Quebec.

Hirsch GB, Immediato CS. 1999. Microworlds and Generic Structures as Resources for Integrating Care and Improving Health. *System Dynamics Review* 15(3):315-330.

Hirsch GB, Killingsworth WR. 1975. A New Framework for Projecting Dental Manpower Requirements. *Inquiry*, June 1975.

Hirsch GB, Wils W. 1984. Cardiovascular Disease in the Dutch Population: A Model-Based Approach to Scenarios. *Ministry of Health: Conference on Health Care Scenarios*. The Hague, Netherlands, August, 1984.

Homer J, Hirsch G, Minniti M, Pierson M. 2003. Models for Collaboration: How System Dynamics Helped a Community Organize Cost-Effective Care for Chronic Illness. *21st International Conference of the System Dynamics Society*, New York, New York.

Institute for Healthcare Improvement (IHI). 2003. Information on the IDCOP approach is available at http://www.ihi.org/idealized/idcop/index.asp .

Institute of Medicine Committee on Quality of Health Care in America (IoM). 2001. *Crossing the Quality Chasm: A New Health System for the 21st Century*. National Academy Press: Washington, DC.

Levin G, Roberts EB, Hirsch GB, Kligler DS, Wilder JF, Roberts N. 1976. *The Dynamics of Human Service Delivery*. Ballinger: Cambridge, Massachusetts.

Luginbuhl W, Forsyth B, Hirsch G, Goodman M. 1981. Prevention and Rehabilitation as a Means of Cost-Containment: The Example of Myocardial Infarction. *Journal of Public Health Policy* 2(2): 103-115.

Rich M, Nease RF. 1999. Cost–Effectiveness in Clinical Practice: The Case of Heart Failure. *Archives of Internal Medicine*, 159, August 9, 1999, 1690-1700

Wagner EH, Sandhu N, Newton KM, *et al*. 2001. Effect of Improved Glycemic Control on Health Care Costs and Utilization. *JAMA* 285(2): 182-189.

West JA *et al.* 1997. A Comprehensive System for Heart Failure Improves Clinical Outcomes and Reduces Medical Resource Utilization. *American Journal of Cardiology*, 79, January 1, 1997, 58-63.