

Treatment Strategies for the Management of Chronic Illness: Is Specialization Always Better?

Peter J. Veazie

University of Minnesota
321 19th Ave S
Minneapolis, MN 55455-0430
Phone: (612) 626-8903/ Fax: (208) 988-1797
veaz0005@umn.edu

Paul E. Johnson

University of Minnesota
321 19th Ave S
Minneapolis, MN 55455-0430
Phone: (612) 626-0320/ Fax: (612) 626-1316
pjohnson@csom.umn.edu

ABSTRACT: *Chronic illness cannot be cured, only controlled. In this paper, we describe an investigation of treatment strategies designed to control the natural progression of type 2 diabetes. We propose that treatment strategies are often specialized to types of patients, and their performance is sensitive to accurate categorization. We investigate the proposition that when incorrect categorization occurs, more specialized strategies may perform worse than less specialized strategies. Following analysis of necessary conditions based on an expected utility model, we present a dynamical systems model of patient care and define two measures of control based on the trajectory of patient health states. The first measure characterizes the accumulated level of control (the extent that health goals are maintained); the second measure characterizes dynamic structure (the time dependencies among health states). Computer simulation is used to analyze the effect of incorrect categorization on these measures.*

KEY WORDS: Chronic illness, Treatment strategies, Dynamical systems, Computer simulation.

1 Introduction

Much of medical care is designed to cure illness and mend injuries; the duration of care is typically short-term and success is defined by specific results. This type of patient-care endeavors to return deviant physiological systems to normal functioning. Chronic illness is an exception to this model because the underlying adverse physiological conditions are not curable; instead, the treatment of chronic illness focuses on controlling health effects throughout the life of the patient. Consequently, the duration of care is long-term and success is defined in terms of control.

The state of health for a chronically ill patient is dynamic and often unstable. When left unattended, the nature of disease governs the trajectories of patient health in the space of possible health states. The patient's state of health eventually deteriorates in normal functioning and,

depending on the illness, can lead to death. The problem faced by physicians who manage chronic illness is one of developing treatment strategies that control the disease process by directing the trajectory of patient health-states over time. More successful strategies typically maintain health-state trajectories in acceptable regions of the space of possible health states.

From the perspective of the healthcare provider, strategies of patient care (hereafter, treatment strategies) are often tailored to patient categories. Physicians and other healthcare providers use categorization to reduce the complexity of the information and knowledge they process. Decision strategies tailored to these categories further reduce cognitive effort (Gigerenzer et al., 1999). Useful categories embed information that provides associated treatment strategies an advantage over more general strategies (i.e. strategies that are not tailored to specific patient categories). Experience leads to more informative categories and better performing treatment strategies by refining the structure of knowledge and defining better decision heuristics.

For curable illnesses and injuries, patient categories are primarily based on medical diagnosis. For chronic illness, patient categories must also account for factors that influence a patient's future health state; consequently, factors that govern physician and patient behavior can be important contributors to patient categorization. Researchers have proposed that physicians use psychosocial dimensions (e.g. compliance behavior) as well as disease states to develop a patient category structure having finer granularity than medical diagnosis alone (Johnson et al., 2001; O'Connor et al., 1997). Treatment strategies based on these additional dimensions reflect relevant variation across categories and may outperform more general strategies.

Treatment strategies can be characterized along three dimensions: (1) structure, (2) performance, and (3) specificity. We hypothesize two types of structure for the treatment strategies employed to manage chronic illness. The first is one in which physicians make decisions and choose clinical moves based on predictions of the future patient states. We term this a feedforward treatment strategy (Brehmer, 1990). Feedforward treatment strategies depend on mental models that include dynamic (time dependent) information regarding (1) patient disease processes, (2) the consequences of past and present courses of action, including patient compliance, and (3) knowledge of the way patients move through the clinical care system (Freyd, 1987)). A feedforward treatment strategy is supported by systems of care in which physicians follow individual patients over time and by clinics in which patients are tracked and monitored so that information about patient state (including past compliance) is available at the time of the patient encounter (Brehmer, 1990; Brehmer, 1992). The second type of structure for treatment strategies is based on the concept of a feedback (as opposed to feedforward) process control. In this type of strategy, physicians make decisions and choose clinical moves using information about the patient's current state as evident in the immediate context of care. A feedback strategy would be expected in a system in which patients are not typically followed by specific physicians, but receive care based on whichever provider is available when the need for care arises. A feedback strategy presumes a mental model that is simpler and makes fewer cognitive and organizational resource demands than a feedforward strategy (Brehmer and Allard, 1991).

Performance and specificity characterize attributes particular to individual treatment strategies. Performance provides a normative measure of how well a strategy controls patient health trajectories (depicted as the height of the bars in Figure 1). Specificity reflects the difference in performance between the category to which a strategy is tailored and other categories. Figure 1 shows a measure of specificity as the differences denoted by S_A and S_B .

Dividing S_A and S_B by each respective strategy's highest performance gives a measure of relative specificity; in a following section we use a related concept, relative generality, equal to one minus relative specificity. Categorization generates a relationship between performance and specificity whereby better performance is achieved through greater specificity.

We compare two kinds of strategies: those that are applied without regard to patient categories (e.g. clinical guidelines) and those that are tailored to patient categories. We label the former, general strategies, and the latter, specialized strategies; both kinds may vary in performance and specificity. Specialized strategies would be expected to outperform general strategies; strategies with higher specificity should outperform those with lower specificity. However, levels of performance and specificity can exist such that treatment strategies with lower specificity outperform those with higher specificity.

Higher specificity implies greater variation in performance across classes. This is depicted in Figure 1 as the greater difference in expected outcomes for high specificity strategies than low specificity strategies (i.e. S_A and S_B are larger in the high specificity graph of Figure 1). Consequently, applying a treatment strategy tailored for one category to patients in another category can decrease treatment effectiveness (e. g. in Figure 1, ΔO_A and ΔO_B are larger in the high specificity graph). This suggests strategies with high specificity may not be preferred if they are applied to patients in other categories.

2 Necessary Conditions for General Strategies

We use an expected utility model to identify necessary conditions for preferring general treatment strategies to specialized treatment strategies. We represent patient categories as a partition Σ on a patient state space S . Each element of a partition comprises the set of patient states that compose a patient category. Patient categories are related to treatment strategies by a function g that maps $S \cup \Sigma$ onto a subset A of all possible treatment strategies. The elements s of the partition Σ are indexed by $i \in \{1, \dots, k\}$ where k is the total number of element in the partition (i.e. number of patient categories). Treatment strategies $a \in A$ are indexed by $j \in \{0, 1, \dots, k\}$. The index $j=0$ denotes a general treatment strategy (i.e. $a_0 = g(S)$); the remaining indices match the corresponding elements of the partition and indicate specialized strategies (i.e. $a_j = g(s_i)$ for all $j = i$). We denote strategy performance by u_{ij} ; the utility associated with applying treatment strategy $a_j \in A$ to a patient in category $s_i \in \Sigma$. The following matrix represents the performance (utility) structure relative to s_i and a_j :

$$\begin{array}{c}
 \text{Patient Categories} \\
 \hline
 \begin{array}{c}
 \mathbf{s}_1 \\
 \mathbf{s}_2 \\
 \vdots \\
 \mathbf{s}_k
 \end{array}
 \end{array}
 \begin{array}{c}
 \text{Physician Strategies} \\
 \hline
 \begin{array}{cccccc}
 & a_0 & a_1 & a_2 & \cdots & a_k \\
 \hline
 & u_{10} & u_{11} & u_{12} & \cdots & u_{1k} \\
 & u_{20} & u_{21} & u_{22} & \cdots & u_{2k} \\
 & \vdots & \vdots & \vdots & \ddots & \vdots \\
 & u_{k0} & u_{k1} & u_{k2} & \cdots & u_{kk}
 \end{array}
 \end{array}
 \tag{1}$$

$\underbrace{\hspace{1.5cm}}_{U^0} \quad \underbrace{\hspace{3.5cm}}_{U^*}$

The matrix labeled \mathbf{U}^* , containing elements $\{u_{ij}: i, j \geq 1\}$, comprises the utilities for specialized strategies. The diagonal of \mathbf{U}^* contains utilities associated with the patient categories for which the specialized strategies are tailored (i.e., the diagonal of \mathbf{U}^* contains states $g^{-1}(a_j)$ for all $j \geq 1$). The column vector labeled \mathbf{U}^0 , containing elements $\{u_{ij}: j = 0\}$, comprises the utilities of the general strategy across patient categories.

We analyze three models, each constrained by two assumptions: First, we assume for each treatment strategy, a_j , the utility associated with the patient category $\mathbf{s}_i = g^{-1}(a_j)$ is greater than or equal to the utility of that strategy applied to other categories (i.e. $u_{mm} \geq u_{im}$ for each $m \geq 1$). This assumption embodies the proposition that treatment strategies are tailored to specific patient categories by virtue of improved utility. Second, we assume the distribution of patients across categories is independent of the treatment strategy decision (although the distribution may be a function of past decisions). This assumption reflects a common ontological commitment regarding the temporal order of causation: causes precede effects.

Expected utility theory implies the general treatment strategy a_0 is preferred to the set of specialized strategies $\{a_1, \dots, a_k\}$ if the expected utility associated with the use of a_0 exceeds the expected utility of using the set $\{a_1, \dots, a_k\}$:

$$E(u_{ij} | a_0) > E(u_{ij} | \{a_1, \dots, a_k\}). \quad (2)$$

The expected utility of the general strategy (the left side of equation 2) is

$$E(u_{ij} | a_0) = \sum_i (u_{i0} \cdot p(\mathbf{s}_i)) \quad (3)$$

where p denotes a probability mass function. We apply two additional assumptions to the first model: (1) for each patient category \mathbf{s}_i the utility of the general strategy is proportional by a constant factor \mathbf{b} to the optimal utility associated with the specialized strategy $a_j = g(\mathbf{s}_i)$; and, (2) the probability of correctly categorizing a patient is the same across categories. The first assumption implies

$$u_{i0} = \mathbf{b} \times u_{ii} \quad (4)$$

for all i , where \mathbf{b} is the proportionality factor representing the relative utility of the general strategy. This assumption requires that we restrict our analysis to specialized strategies with non-zero utilities on the diagonal of \mathbf{U}^* . From equation 3, the general strategy's expected utility is rewritten as

$$E(u_{ij} | a_0) = \mathbf{b} \cdot \sum_i (u_{ii} \cdot p(\mathbf{s}_i)). \quad (5)$$

Denoting the expected utility for $i = j$, with respect to the marginal distribution $p(\mathbf{s})$, as $\bar{u}_{i=j}$, equation 5 becomes

$$E(u_{ij} | a_0) = \mathbf{b} \cdot \bar{u}_{i=j}. \quad (6)$$

The expected utility of specialized strategies (the right side of equation 2) is

$$E(u_{ij} | \{a_1, \dots, a_k\}) = \sum_{i,j,m} u_{ij} \cdot p(a_j | \mathbf{s}_i, m) \cdot p(m | \mathbf{s}_i) \cdot p(\mathbf{s}_i). \quad (7)$$

In this formulation m is a dichotomous variable for which $m = 0$ represents correct categorization of a patient, and $m = 1$ represents incorrect categorization of a patient. Denoting the probability of correct categorization as \mathbf{p} (i.e. $\pi = p(m = 0 | \sigma_i)$ for all i categories), equation 7 can be written as

$$E(u_{ij} | \{a_1, \dots, a_k\}) = \mathbf{p} \cdot \sum_i (u_{ii} \cdot p(\mathbf{s}_i)) + (1 - \mathbf{p}) \cdot \sum_i \left(p(\mathbf{s}_i) \cdot \sum_{j \neq i} (u_{ij} \cdot p(a_j | \mathbf{s}_i, m = 1)) \right) \quad (8)$$

or simply

$$E(u_{ij} | \{a_1, \dots, a_k\}) = \mathbf{p} \cdot \bar{u}_{i=j} + (1 - \mathbf{p}) \cdot \bar{u}_{i \neq j}, \quad \forall (i,j) \geq 1. \quad (9)$$

Substituting equations 6 and 9 into equation 2 and solving for \mathbf{b} gives the necessary condition for preferring the general strategy to the specialized strategies in this model:

$$\mathbf{b} > \mathbf{p} + (1 - \mathbf{p}) \cdot \frac{\bar{u}_{i \neq j}}{\bar{u}_{i=j}}. \quad (10)$$

We define the ratio $\frac{\bar{u}_{i \neq j}}{\bar{u}_{i=j}}$ as the relative generality \mathbf{g} of the strategies $\{a_1, \dots, a_k\}$. Relative generality is one minus relative specificity. These concepts are mirror images of each other; hence, either can be used in this analysis without loss of clarity.

We assume $\bar{u}_{i \neq j} \leq \bar{u}_{i=j}$, hence relative generality is less than or equal to 1 (i.e., $\mathbf{g} \leq 1$). As \mathbf{g} approaches 1, the specialized strategies perform equally well across patient categories; consequently, there is no cost associated with incorrect categorization and \mathbf{b} must exceed 1 to prefer the general strategy. Similarly, as \mathbf{p} approaches 1, the chance of incurring a cost due to incorrect categorization is diminished, and again, \mathbf{b} must exceed 1 to prefer the general strategy. As \mathbf{g} approaches 0, the specialized strategies do not function with the patient categories where $i \neq j$, implying a greater cost associated with incorrect categorization. In the last case, the second term on the right side of equation 10 is 0, implying the relative performance of the general strategy must exceed the probability of misclassification.

Figure 2 shows a contour plot of \mathbf{b} on the parameter space defined by \mathbf{p} and \mathbf{g} . The maximum of \mathbf{p} and \mathbf{g} defines the lower bound of \mathbf{b} for preferring the general strategy; the lower bound is achieved if minimum of \mathbf{p} and \mathbf{g} is equal to 0. For example, a physician using

specialized strategies with \mathbf{g} near 0 and a 0.8 probability of correctly categorizing patients must have a general strategy that performs better than 80% of the specialized strategies' performance if the general strategy is to be preferred. As \mathbf{g} increases, the lower bound of the general strategy's performance is higher. This result conforms to intuition: as generality increases, the cost of incorrect categorization is diminished and a competing general strategy must increase performance to remain the preferred strategy. Similarly, for a given level of generality, as the probability of correctly categorizing patients increases, the probability of incurring a loss due to incorrect categorization decreases; again, a general strategy would require better performance to compete with the increased accuracy with which the specialized strategies are applied.

The preceding model assumes \mathbf{b} is constant across patient categories. A model in which \mathbf{b} varies across categories can be analyzed by expressing the utilities of the general strategy as a proportion of the mean utility across u_{ij} for $i = j$ of the specialized strategies¹. Substituting

$$u_{i0} = \mathbf{b}_i \cdot \bar{u}_{i=j} \quad (11)$$

into equation 3 gives

$$E(u_{ij} | a_0) = \bar{u}_{i=j} \cdot \sum_i (\mathbf{b}_i \cdot p(\mathbf{s}_i)). \quad (12)$$

Following the reasoning used in the preceding analysis, the necessary condition for preferring the general strategy in this case is

$$\bar{\mathbf{b}} > p + (1-p) \cdot \frac{\bar{u}_{i \neq j}}{\bar{u}_{i=j}}. \quad (13)$$

Here $\bar{\mathbf{b}}$ denotes the mean relative performance of the general strategy representing the summation in equation 12. Figure 2 and the conclusions of the preceding analysis apply to this model as well, but we use the mean relative performance $\bar{\mathbf{b}}$ of the general strategy in place of \mathbf{b} .

Both analyses presented thus far assume the probability of correct categorization is independent of patient categories. Without this assumption, we can state a more general criterion for the preference of the general strategy. From equations 2 and 12 the necessary condition is

$$\bar{\mathbf{b}} > \frac{E(u_{ij} | \{a_1, \dots, a_k\})}{\bar{u}_{i=j}}. \quad (14)$$

The numerator in the right side term of the inequality is equal to or less than the denominator, the equality holding only when the probability of correct classification is 1. A general strategy must therefore perform better on average across patient categories than specialized strategies relative to the optimal performance of the specialized strategies (i.e.

relative to the expected utility of correctly applied specialized strategies). Factoring equation 14 and rewriting gives

$$\bar{b} > \sum_i p(\mathbf{s}_i) \cdot \left(\frac{(\mathbf{p}_i \cdot \mathbf{u}_{ii})}{\bar{u}_{i=j}} + \frac{\left((1 - \mathbf{p}_i) \cdot \sum_{j \neq i} (u_{ij} \cdot p(a_j | \mathbf{s}_i, m=1)) \right)}{\bar{u}_{i=j}} \right). \quad (15)$$

This formulation reveals two category-specific components: (1) the decrease in performance due to the probability of not applying specialized strategies to the targeted category (represented by the first term in the parenthesis on the right side of the inequality); and, (2) the modifying affect of a non-zero utility associated with misapplying treatment strategies (represented by the second term on the right side of the inequality). The import of these components is based on the probability of correct categorization \mathbf{p}_i , and the utilities u_{ij} . If patients are always correctly categorized (i.e. $\mathbf{p}_i = 1$ for all i), or if each specialized strategy performs equally well across categories (i.e. $u_{ii} = u_{ij}$ for all $j \geq 1$), then the first term equals 1 and the second term equals 0. In this case, the general strategy must outperform the optimal application of the specialized strategies, which would contradict the assumption that specialization is driven by improved utility. If patients are always incorrectly categorized ($\mathbf{p}_i = 0$ for all i), then the first term is equal to 0 and the general strategy must outperform the consistent misapplication of the specialized strategies. If the utilities u_{ij} equal 0 for all $i \neq j$ (i.e. the off-diagonal elements of \mathbf{U}^*), the second term is 0 and \bar{b} is bound solely by the relative performance of the diagonal elements of \mathbf{U}^* .

In this section we have identified necessary conditions for the preference of general strategies based on the model presented in matrix 1. We compared the use of general strategies versus a number of specialized strategies considered as a set. Alternatively, general strategies can be compared with each specialized strategy individually. Results are the same as those presented here, only they apply separately to each column of \mathbf{U}^* .

We next use computer simulation to investigate the relationship between general and specialized strategies in the context of dynamic interactions between treatment strategies and the health states of patients with type 2 diabetes.

3 Dynamical Systems Model of Patient Care

We represent patients as a map $G_{P_t, \mathbf{q}}$ from the space of possible treatments \mathbf{S}_{R_x} to the space of possible patient health states \mathbf{S}_{P_t} : patients respond, via changes in state, to the moves generated by treatment strategies. We represent treatment strategies as a map $G_{R_x, \mathbf{f}}$ from the patient state space \mathbf{S}_{P_t} to the treatment space \mathbf{S}_{R_x} : treatment strategies generate moves in response to patient state information.

The symbols \mathbf{q} and \mathbf{f} denote parameters specifying patients and treatment strategies respectively. For \mathbf{q} and \mathbf{f} ranging over a set of patient categories, the map $G_{P_t, \mathbf{q}}$ represents patient types and $G_{R_x, \mathbf{f}}$ represents treatment strategies of varying specificity (we denote general strategies as $\mathbf{f} = \cdot$). There are two general cases: (1) \mathbf{f} is the same as \mathbf{q} , and (2) \mathbf{f} is not the same as \mathbf{q} . The first case corresponds to the application of a category-specific treatment strategy to a

patient in the same category; the second case corresponds to the application of a category-specific treatment strategy to a patient of another category.

The interaction between patients and physicians (i.e. the compositions of the patient and treatment maps) generate trajectories of patient health states and treatment moves:

$$F_{P_t(f,q)} \circ G_{P_t,q} \circ G_{R_x,f}: \mathbf{S}_{P_t} \rightarrow \mathbf{S}_{P_t} \quad (16)$$

and

$$F_{R_x(q,f)} \circ G_{R_x,f} \circ G_{P_t,q}: \mathbf{S}_{R_x} \rightarrow \mathbf{S}_{R_x}. \quad (17)$$

The pairs $\langle \mathbf{S}_{P_t}, F_{P_t(f,q)} \rangle$ and $\langle \mathbf{S}_{R_x}, F_{R_x(q,f)} \rangle$ are discrete dynamical systems. We focus on the dynamical system $\langle \mathbf{S}_{P_t}, F_{P_t(f,q)} \rangle$ (i.e. the effects of treatment strategies). Specifically, we analyze patient state trajectories generated by F_{P_t} :

$$\mathbf{x}_n = F_{P_t(f,q)}^n(\mathbf{x}_0) \text{ for } n \in (0, 1, 2, \dots) \quad (18)$$

where $F_{P_t(f,q)}^n(\mathbf{x}_0)$ is the n^{th} iteration of the dynamic function on initial patient state $\mathbf{x}_0 \in \mathbf{S}_{P_t}$.

We use computer simulation to operationalize the preceding dynamical system and investigate the relative effects of general and specialized treatment strategies. Both the patient and treatment strategies, $G_{P_t,q}$ and $G_{R_x,f}$, can be encoded as computer programs; their interaction (i.e. the composition $F_{P_t(f,q)}$ in equation 16) generates trajectories in the patient state space.

3.1 Patient model

Figure 3 depicts the structure of the patient model that encodes the map $G_{P_t,q}$. The dashed box encloses the patient model; the components labeled medication effort and psychosocial effort are inputs from the treatment strategy based on patients' glycosylated hemoglobin levels (HbA_{1c}). The model updates patient HbA_{1c} in response to inputs generated by the treatment strategy.

We use HbA_{1c} as the outcome variable based on evidence that it corresponds with the physiological health status of patients with Type 2 diabetes mellitus (United Kingdom Prospective Diabetes Group, 1998), plus the fact that HbA_{1c} is an important health indicator used in the management of patients with type 2 diabetes (American Diabetes Association, 1999). Higher HbA_{1c} levels indicate worse patient health.

The patient model in Figure 3 comprises six concepts:

- 1) HbA_{1c} level—the patient health state variable.
- 2) Adherence—the extent to which a patient complies with the prescribed treatment regimens.
- 3) Side effects—the adverse physical manifestations associated with medications.
- 4) Stress—the adverse psychophysical response to the psychosocial environment.

- 5) Medication effort—a representation of the amount and number of medications prescribed by a given treatment strategy.
- 6) Psychosocial effort—the psychological pressure, education, and motivation applied by the health care system to a patient regarding self-care behavior.

Ten relationships integrate these concepts:

- 1) Increased medication effort decreases HbA_{1c} levels, thereby improving patient health.
- 2) Increased medication effort increases side effects.
- 3) Increased side effects decreases patient adherence with prescribed medication regimens.
- 4) Increased adherence increases side effects by virtue of increasing the effective medication dose (e.g., a patient that does not take his medicine does not experience medication induced side effects).
- 5) Increased psychosocial effort either increases or decreases adherence depending on the level of psychosocial effort and specific patient parameterization used in the definition of patient categories.
- 6) Increased adherence decreases HbA_{1c} by increasing the effective medication effort.
- 7) Increased psychosocial effort increases stress.
- 8) Increased side effects increases stress.
- 9) Increased stress increases HbA_{1c} (Daniel et al., 1999).
- 10) Increased Hb_{1c} increases adherence. This relation is derived from the link between HbA_{1c}, comorbidities, and motivation: high HbA_{1c} corresponds to more comorbidities, and more comorbidities imply more manifest health consequences, thereby motivating greater compliance with treatment regimens.

We operationalize the model by (1) defining a discrete-time update function for HbA_{1c} (labeled h in the following equations) and (2) expanding each term to integrate model relationships. For a given patient i at time t , we assume change in HbA_{1c} levels is effected by two independent factors: (1) a disease effect $\mathbf{d}_{i,t}$ that increases h and, (2) a medication effect $DE_{i,t}$ that decreases h . Patient HbA_{1c} is updated at time t according to the function²

$$h_{i,t+1} = h_{i,t} + \mathbf{d}_{i,t} - DE_{i,t} . \quad (19)$$

For each patient i at a given time t we assume the disease effect $\mathbf{d}_{i,t}$ has a truncated normal distribution

$$\mathbf{d}_{i,t} \sim \theta(\mathbf{m}_{i,t}, \mathbf{s}_{i,t} / \mathbf{d}_{i,t} \geq 0) . \quad (20)$$

The parameter $\mathbf{m}_{i,t}$ is the mean of the corresponding non-truncated distribution and comprises a patient specific time invariant characteristic m_i and a positive perturbation $s_{i,t}$ generated by the patient's current level of stress

$$\mathbf{m}_{i,t} = m_i + s_{i,t} . \quad (21)$$

The time invariant characteristic m_i is set for each patient by a draw from a log-normal distribution

$$m_i \sim \text{LogNormal}(-1.830664, 0.8671747). \quad (22)$$

This specification represents the distribution of the average 6-week HbA_{1c} positive change among 6,768 patients with type 2 diabetes. The data were obtained from encounter records of a Minnesota staff-model HMO; all patients with type 2 diabetes that had HbA_{1c} tests during the years 1994 to 1998 are included.

The stress effect $s_{i,t}$ in equation 21 is modeled as

$$s_{i,t} = 1 - (1 - 0.73)^{\frac{1}{2} \cdot (a_{i,t} \cdot d_{i,t} + F_{i,t})}. \quad (23)$$

In this formulation d is the current medication effort, a is the current patient adherence, F is the current psychosocial effort, and the constant 0.73 is the standard deviation of m_i across the 6768 patient's in the empirical data. Equation 23 is an arbitrary functional specification selected to satisfy the assumption that increased stress increases the mean disease effect by an amount in some interval $(0, t)$. We select an upper threshold $t = 0.73$ to bound the effect of stress within one expected deviation. The effect of stress is maximal (equal to 0.73) when adherence, medication and psychosocial efforts are each equal to 1. The effect of stress is minimal (equal to zero) when either (1) adherence or medication effort is equal to 0, or (2) psychosocial effort is equal to 0.

The $s_{i,t}$ parameter in equation 20 (i.e. the standard deviation of the non-truncated normal distribution) is determined as a function of the mean m_i for each patient

$$s_i = .0104857 + .5189586m_i + .5775702m_i^2 - .0532329m_i^3. \quad (24)$$

The constants in this equation are estimated by a regression of patient's standard deviation in 6-week HbA_{1c} positive changes on a third-order polynomial of patient's average 6-week HbA_{1c} positive changes.

The effect of medication effort DE in equation 19 is modeled as a proportion of the current A1c level

$$DE_{i,t} = P_{i,t} \mathcal{A}_{i,t}. \quad (25)$$

The proportionality factor $P_{i,t}$ is calculated as

$$P_{i,t} = \int_{D_t}^{D_{t-1}} g(D) dD. \quad (26)$$

Here $D_{i,t} = a_{i,t}d_{i,t}$ is the effective medication effort (reflecting the attenuating effect of adherence on medication effort) and $g(D)$ is the marginal effect of D on P (i.e. dP/dD). Integrating g over its domain (i.e. over the interval $[0,1]$) gives the maximum effect of medication effort ($MaxEffect$). We assume g is monotonically decreasing. In our analysis $g(D)$ is a linear function of D such that $g(0) = 2 \cdot MaxEffect$ and $g(1) = 0$:

$$g(D) = 2 \cdot MaxEffect \cdot (1 - D). \quad (27)$$

Hence, as shown in Figure 4, the effect of changing medication effort from $D_{i,t-1}$ to $D_{i,t}$ can be calculated as the area of a trapezoid:

$$P_{i,t} = \frac{1}{2} (2 \cdot MaxEffect \cdot (1 - D_{i,t}) + 2 \cdot MaxEffect \cdot (1 - D_{i,t-1})) \cdot (D_{i,t} - D_{i,t-1}). \quad (28)$$

Adherence $a_{i,t}$ (embedded in the definition of D) is calculated as the product of two effects. The First effect, S , captures (1) the negative influence of medication side effects ($SideEffect$), (2) the positive influence of A1c level as an indicator of potential comorbidities (higher HbA_{1c} levels correspond to more comorbidities and worse manifest health conditions), and (3) the positive influence of the HbA_{1c} gradient with respect to the previous change in adherence representing improved self-efficacy (Bandura, 1977; Bandura, 1997; Bandura, 2001; Schwarzer and Fuchs, 1995). The second effect, B , captures the effect of psychosocial effort ($PsychosocialEffort$) as a unimodal function representing increased adherence for levels of psychosocial effort below a threshold and decreased adherence above the threshold.

$$a_{i,t+1} = S(SideEffects, h_{i,t}, Dh_{i,t}/Da_{i,t}) \cdot B(PsychosocialEffort, \mathbf{a}, \mathbf{b}). \quad (29)$$

The function S , is arbitrarily specified to meet two conditions: (1) adherence must remain in the interval $[0,1]$, and (2), as a function of the variables $\mathbf{x} = (Side\ Effects, h_{i,t}, Dh_{i,t}/Da_{i,t})^T$, S must be able to generate both positive and negative deviations from a patient specific base level of adherence a_i . We use the specification

$$S_{i,t} = \max \left(\min \left(a_i + \left(\frac{2}{1 + e^{-(\boldsymbol{\gamma} \cdot \mathbf{x})}} - 1 \right), 1 \right), 0 \right) \quad (30)$$

where $\boldsymbol{\gamma}$ is a vector of weight parameters specifying the relative importance and direction of effect for each component in \mathbf{x} . In equation 29, B denotes a beta function with parameters \mathbf{a} and \mathbf{b} . B is scaled to equal 1 at its mode. Numerous linear and nonlinear functions can be achieved via the parameterization (\mathbf{a}, \mathbf{b}) of B .

As shown in Figure 5, adherence can either increase or decrease in response to psychosocial effort depending on the level of psychosocial effort and the patient parameterization of the Beta function. The functions denoted as $(S = 0.5, \mathbf{a} = 1, \mathbf{b} = 5)$ and $(S = 0.8, \mathbf{a} = 5, \mathbf{b} = 1)$ in Figure 6 represent patients that monotonically respond to increased psychosocial effort; the first reacts by decreasing adherence, the second by increasing adherence.

The functions denoted as ($S = 1$, $\mathbf{a} = 2$, $\mathbf{b} = 5$) and ($S = 0.6$, $\mathbf{a} = 5$, $\mathbf{b} = 2$) represent patients that respond well to increased psychosocial up to a threshold point and then decrease adherence after psychosocial effort exceeds the threshold. The scale S of the Beta function is set by equation 30; higher values of S imply greater possible adherence.

Medication side effects (*SideEffects* in equation 29) enter the model as a proportion \mathbf{u} of the effective medication effort

$$SideEffects_{i,t} = \mathbf{u} \mathcal{D}_{i,t}. \quad (31)$$

Two patient categories based on adherence behavior are defined for the purpose of this analysis: Category 1, composed of patients with Beta parameters $\mathbf{a} < \mathbf{b}$, and category 2, composed of patients with Beta parameters $\mathbf{a} \geq \mathbf{b}$. Category 1 represents patients that respond well to low levels of adherence but react negatively to high levels of adherence. Category 2 represents patients that respond well to higher levels of adherence and react negatively only to the highest levels of adherence. Specialized treatment strategies will be defined for each category of patient, and a general treatment strategy will be defined for all patients without regard to categorization.

3.2 Treatment Strategies

We use goal-directed machine learning to capture the knowledge structure of treatment strategies. More specifically, we use artificial neural networks to represent the treatment strategies $G_{R_x, \mathbf{f}}$. Networks representing specialized strategies of each patient category are trained using populations of category-specific simulated patients.

The parameter \mathbf{f} of the treatment strategy is an index of patient category. Levels of specificity are achieved by varying training experience: for each category, three networks are trained at different experience levels. Networks representing general strategies are trained using a population composed of simulated patients from a sample of two categories. A total of 7 networks are trained (i.e. [2 categories \times 3 experience levels] + 1 general strategy).

An online reinforcement-learning algorithm is used to train the neural networks. The reinforcement function embeds (1) physician goals regarding patient health states and (2) constraints on treatment. For example, goals derived from the clinical practice literature are the reduction of HbA_{1c} levels, the reduction of LDL levels, and the reduction of Bp level. Treatment constraints include minimizing the number of drugs given at any one time, and minimizing drug dosage early in treatment.

Interaction between patient and treatment strategy is achieved by inputting patient state information to the neural network and passing the resulting treatment moves from the network to a simulated patients. The simulated patient then generates a new patient state in response to the treatment move. Recursive operation on resultant patient states produces a trajectory in the patient state space.

4 Analysis

Strategies at different levels of specificity are compared using functions of patient health-state trajectories (denoted as $O_{(\mathbf{f}, \mathbf{q})}$). We consider two functions of HbA_{1c} trajectories: one summarizes the overall control achieved by treatment strategies; the other summarizes the

constraint imposed by treatment strategies on patient variation across time. The first function is the sum of state values along the trajectory of HbA_{1c} states:

$$O_{(f,q)}(HbA_1c_o) = \sum_n F_{Pt,(f,q)}^n(HbA_1c_o). \quad (32)$$

This outcome measure associates higher values with trajectories containing poor health states and lower values with those containing good health states. The second function, which characterizes the dynamic structure imposed on the trajectory by the treatment strategy, is defined as the determinant of the autocorrelation matrix of the trajectory. This function returns smaller values for strategies that better control variation across time in the patient's health state.

The expected outcome of either function for a set of strategies and a probability of applying a category-specific treatment strategy to a patient of another category $p(M) = m$ is

$$E(O | p(M) = m) = \int \sum_q \sum_f O_{(f,q)}(x) p(\mathbf{f} | \mathbf{q}, x, m) p(\mathbf{q} | x) f(x) dx. \quad (33)$$

The function $f(x)$ is the probability density of the patient state space, $p(\mathbf{q} | x)$ is the conditional probability mass function of category membership (the patients actual category), and $p(\mathbf{f} | \mathbf{q}, x, m)$ is the conditional probability mass function of categorization (the category index of the treatment strategy applied to the patient). Neither $f(x)$ nor $p(\mathbf{q} | x)$ are conditioned on m because both are independent of categorization.

We specify the conditional probability of categorization $p(\mathbf{f} | \mathbf{q}, x, m)$ as

$$p(\mathbf{f} | \mathbf{q}, m, x) = \begin{cases} (1 - m) & \text{if } \mathbf{f} = \mathbf{q} \\ p(\mathbf{q} | \mathbf{f} \neq \mathbf{q}, x) \cdot m & \text{if } \mathbf{f} \neq \mathbf{q} \end{cases}. \quad (34)$$

The probability of correct categorization (i.e. $\mathbf{f} = \mathbf{q}$) is one minus the probability of incorrect categorization (i.e. $\mathbf{f} \neq \mathbf{q}$). The probability of incorrect categorization is distributed across the remaining categories according to the distribution of those categories for a given health state.

When $p(M) = 0$ (i.e. category-specific treatment strategies are always applied to appropriate patients), the conditional probability of categorization becomes

$$p(\mathbf{f} | \mathbf{q}, m, x) = \begin{cases} 1 & \text{if } \mathbf{f} = \mathbf{q} \\ 0 & \text{if } \mathbf{f} \neq \mathbf{q} \end{cases}. \quad (35)$$

The expected outcome then reduces to

$$E(O | p(M) = 0) = \int \sum_q O_{(q,q)}(x) p(\mathbf{q} | x) f(x) dx. \quad (36)$$

When the specificity is zero, as is the case of a general strategy ($\phi = \cdot$), the expected outcome is

$$E(O | \mathbf{f} = \cdot) = \int \sum_q O_{(c,q)}(x) p(\mathbf{q} | x) f(x) dx. \quad (37)$$

The difference between equations 36 and 37 is generated solely by the effect of specificity on the outcome function.

Results presented at the conference will be in three parts: First, we characterize expected outcomes as a function of the probability of applying category-specific treatment strategies to other categories of patients for various levels of specificity (see Figure 6). Second, we identify the probability of incorrect categorization for which equivalent expected outcomes are achieved, such as the points *a* and *b* in Figure 6. Finally, we determine the probability of incorrect categorization (such as point *d* in Figure 6) at which a general strategy achieves the expected outcome of specialized strategies.

5 Conclusions

Chronic illness requires a distinct mode of care, one for which the goal is system control rather than system repair. Successful control entails continual management of a patient's health state, even when physical maladies are not evident. Management is achieved through treatment strategies often tailored to specific patient categories. However, as we have shown, the necessary conditions for strategy selection based on performance under optimal conditions may not be preferred due to the effects of specificity and faulty categorization.

In the context of type 2 diabetes, when category-specific treatment strategies are appropriately applied, the expected outcome using more specific strategies should exceed the expected outcome using less specific strategies. This result implies tailoring treatment strategies to patient categories based on psychosocial attributes, such as compliance with drug regimens, will be successful if patient categorization is accurate. However, as the probability of applying a category-specific strategy to patients of another category increases, the advantage of treatment specificity is diminished.

When the probability of appropriate patient categorization is less than 1, the use of treatment strategies with lower specificity may be preferred. In circumstances with this result, it is not necessary that the strategy with lower specificity have a high level of performance; the only requirement is that the lower specificity strategy performs *better* than the *misapplication* of the strategy with high specificity. This can result in a low level of performance.

If treatment strategies are too specific and there is a possibility of incorrect categorization, policies that strive to capitalize on tailoring strategies to manage the dynamics of chronic illness may fail. The benefits of tailoring strategies to patient categories should be considered in light of the expected accuracy of the categorization process.

FIGURES




Figure 1. Treatment strategy performance and specificity




Figure 2: Contour map of a general strategy's relative performance b on the space defined by relative generality g and the probability of correct patient categorization p .




Figure 3: Patient Model

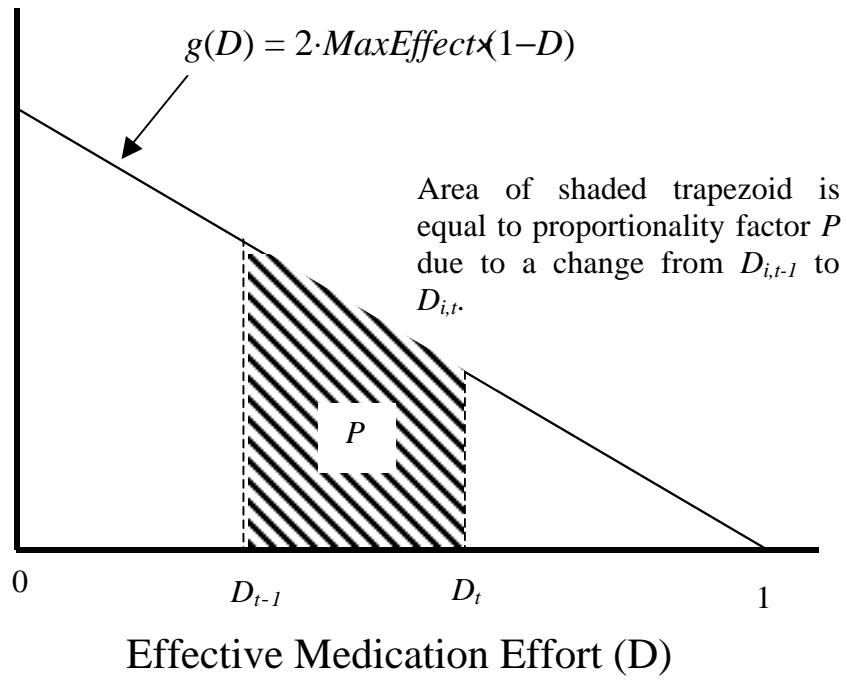


Figure 4: The effect of changing medication effort from D_{t-1} to D_t .




Figure 5: Adherence Functions

REFERENCES

- American Diabetes Association, 1999. Standards of medical care for patients with diabetes mellitus. *Diabetes Care*, 22(Supp. 1).
- Bandura, A., 1977. Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84: 191-215.
- Bandura, A., 1997. *Self-efficacy: The exercise of control*. W. H. Freeman and Company, New York.
- Bandura, A., 2001. SOCIAL COGNITIVE THEORY: An Agentic Perspective. *Annual Review of Psychology*, 52: 1-26.
- Brehmer, B., 1990. Strategies in real-time, dynamic decision making. In: R.M. Hogarth (Editor), *Insights in Decision Making*. Chicago University Press, Chicago, pp. 262-279.
- Brehmer, B., 1992. Dynamic decision making: human control of complex systems. *Acta Psychologica*, 81: 211-241.
- Brehmer, B. and Allard, R., 1991. Dynamic decision making: The effects of task complexity and feedback delay. In: J. Rasmussen, B. Brehmer and J. Leplat (Editors), *Distributed Decision Making: Cognitive Models for Cooperative Work*. John Wiley and Sons, New York, pp. 319-334.
- Daniel, M., O'Dea, K., Rowley, K.G., McDermott, R. and Kelly, S., 1999. Glycated Hemoglobin as an Indicator of Social Environmental Stress among Indigenous versus Westernized Populations. *Preventive Medicine*, 29: 405-413.
- Freyd, J.J., 1987. Dynamic mental representations. *Psychological Review*, 94(4): 427-438.
- Gigerenzer, G., Todd, P.M. and Gerd Gigerenzer, A.R., 1999. *Simple Heuristics That Make Us Smart*. Evolution and Cognition. Oxford University Press, New York.
- Johnson, P.E. et al., 2001. Understanding Variation in Chronic Disease Outcomes. *Health Care Management Science*(submitted for publication).
- O'Connor, P.J., Crabtree, B.F. and Yanoshik, M.K., 1997. Differences between diabetic patients who do not respond to diabetes care interventions: A qualitative analysis. *Family Medicine*, 29: 424-428.
- Schwarzer, R. and Fuchs, R., 1995. Self-efficacy and health behavior. In: M. Conner and P. Norman (Editors), *Predicting Health Behavior: Research and Practice with Social Cognition Models*. Open University Press, Buckingham.
- Sontag, E.D., 1990. *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Texts in Applied Mathematics, 6. Springer Verlag, New York.
- United Kingdom Prospective Diabetes Group, 1998. Intensive blood-glucose control with sulfonylureas or insulin compared with conventional treatments and risk of complications in patients with type 2 diabetes (UKPDS 33). *Lancet*, 352: 837-853.

NOTES

¹ The use of $\bar{u}_{i=j}$ is a convenience for analysis but it is not an assumption: The model is based on a number system with an algebraic structure such that $\forall(a,b \in \mathfrak{R}) \exists(c \in \mathfrak{R})(a \cdot c = b)$; therefore, the ratio of any two non-zero numbers exists and need not be assumed.

² At this level of description, equation 19 appears as a linear discrete-time system with control feedback *DE* (Sontag, 1990, p. 36); however, the definitions of ***d*** and *DE* embed the non-linearity of the system.