Proceedings of the 2009 Winter Simulation Conference
M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin and R. G. Ingalls, eds.

A SIMULATION MODEL OF HIV TREATMENT UNDER DRUG SCARCITY CONSTRAINTS

Robert T. Koppenhaver
Andrew Schaefer

1048 Benedum Hall

University of Pittsburgh
Pittsburgh, PA 15261, USA

R. Scott Braithwaite

950 Campbell Ave, 11
ACSL-G Bldg 35A
Yale University
West Haven, CT 06516, USA

Mark Roberts

200 Meyran Ave

University of Pittsburgh
Pittsburgh, PA 15213, USA

ABSTRACT

We consider a population of HIV infected patients. In resource poor environments, decision makers must allocate antiretroviral drugs (ARVs) to patients in need of them the most. Further complicating matters is that once a patient is given ARVs, the decision maker must decide when to deny further access to ARVs. We compare various methods for determining which patients should receive ARVs and when to switch a patient off of ARVs. We examine the World Health Organization's (WHO) treatment recommendations and how the level of drug shortages can influence the performance of these recommendations. Instead of a single recommendation, the WHO offers three distinct treatment policies with no mention of when to use them. We find that the severity of drug shortages can greatly impact the performance of these policies and the performance gap can be as high as 1.4 years.

1 INTRODUCTION

The Human Immunodeficiency Virus (HIV) affects 33 million people worldwide, with an estimated 22 million living in areas where the supply for drugs does not meet demand. In such situations, decision makers are forced to prioritize among the patients to determine who should be treated. The only treatment option for chronic HIV remains to be antiretroviral therapy (ART). Antiretrovirals have been shown to prolong both lifetime and quality-adjusted life years. For patients, ART can result in a suppressed viral load and boosted CD4 count, until the virus adapts to the drugs and forms a resistant mutation. Once a mutation has developed, ART's effectiveness is diminished and it may be beneficial to begin treating another patient who is waiting. It is for these reasons that proper antiretroviral management in the areas most affected by the HIV epidemic is one of the most urgent problems in global health.

The World Health Organization (WHO) reports that at the end of 2005 there were 1.3 million people receiving ART in low and middle income countries (WHO HIV statistics 2005). However that figure represents only 20% of the demand for the drugs. Countries that face drug shortages also experience different levels of scarcity. Some of the hardest hit countries, in terms of drug scarcity, are in areas of the world where the prevalence of HIV is the highest. In particular, countries in Sub-Saharan Africa are forced to deal with drug scarcity when managing the HIV epidemic. Coverage levels in this area of the world are generally very low and can range from 6% in Niger to 72% in Rwanda with the majority of countries having coverage less than 40%. Increasing access to ART is important in reducing the spread of the virus as well as increasing life expectancy of the patients. Therefore, the efficient distribution and management of these medications, especially in resource-poor areas has substantial potential to alter the course of the epidemic. Consequently in the past few years, the WHO has significantly increased the number of drugs available to patients in the areas affected the most by the HIV epidemic. Given this increased supply of drugs, the need for proper management of them has become an important issue that has yet to be adequately addressed. Specifically, decision makers must be aware of the health benefits and consequences that specific treatment plans have on the overall population's health.

While eliminating scarcity and increasing access to ART is a potential long-term solution to the HIV epidemic, managing treatment decisions in a population where the supply of drugs does not meet the demand is a complex, important and urgent problem. Clinicians are faced with two basic treatment questions: "Who to treat" and "When to stop treatment". While the WHO has published treatment decision guidelines, there is ambiguity in their specific application to resource-poor environments. WHO guidelines state that a Rule of Rescue (RoR) policy is appropriate when deciding who to treat (WHO Guidelines 2006). However the guidelines offer multiple recommendations regarding how to define treatment failure. This is an important issue since upon experiencing treatment failure, it may be more beneficial to begin treating another patient. Specifically, the WHO offers 3 recommendations on when ART treatment has failed based on a patient's CD4 count progression

since initiating ART. Previous work using simulation models have accurately predicted the effects of various treatment policies in a cohort of patients (Braithwaite et al. 2005, Freedberg et al. 2002) but are not equipped to capture the effects of drug scarcity within the cohort. In this paper, we describe the use of a modified, physiologically-based simulation model of HIV to investigate the effects of various HIV treatment decisions in resource–poor environments.

Mathematical models have been used extensively in medicine when clinical trials would be impractical (Shechter et al. 2005, Saka et al. 2007). In particular, simulation models have been used to address a variety of questions pertaining to HIV care. These models have been constructed using clinical data and have been shown to be a valid mechanism for simulating the progression of a patient's health during the chronic stage of HIV. Previous work has focused on simulating a cohort of HIV+ patients in resource-rich environments. These models implicitly assume that drugs will be available to patients whenever they are needed. Therefore to model a resource-poor population, we must alter the mechanics of these previous models while still maintaining their clinical validity. We adapted an existing *individual* simulation model of HIV to become a population model of HIV that incorporates the effects of drug scarcity.

1.1 Chronic HIV Infection

Chronic HIV infection is a complex biological process in which the HIV virus attacks a patient's immune system until, at some time, the onset of AIDS occurs and eventually death. For patients with HIV there are two measures of health that are widely used: CD4 count and viral load (VL). The CD4 count measures the concentration of CD4 cells in the patient's system and is essentially a measure of how strong the patient's immune system is (higher values are preferred). These cells are crucial to immune system functions. HIV attacks these cells and uses them to produce copies of the HIV virus. When patients visit their physician, the CD4 count is tested and many decisions related to treatment. For these reasons, CD4 count has been referred to as the most important prognostic variable in HIV treatment (Braithwaite et al. 2006). While the CD4 count measures the strength of the patient, VL measures the strength of the virus in the patient's system. Specifically viral load is the concentration of the virus in the patient's bloodstream (lower values are preferred). CD4 count and VL are closely related, as one would expect. There is almost a perfect negative correlation between the two. For example, if a patient's VL is very high there are more copies of HIV that will attack the CD4 cells leading to a low CD4 count.

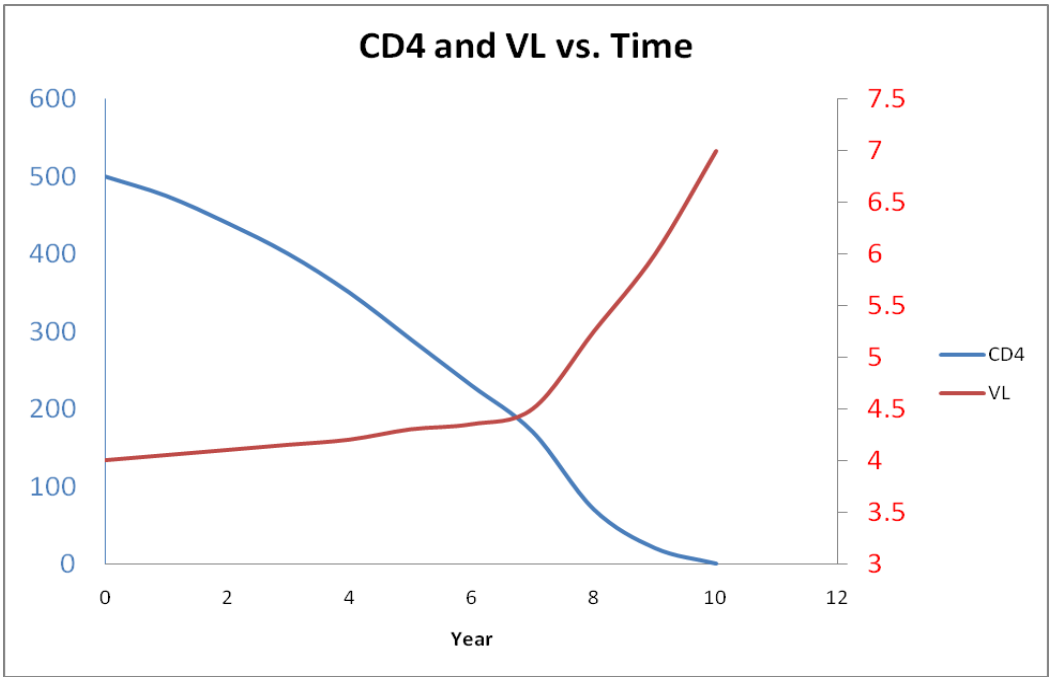

Figure 1: Natural history of chronic HIV infection

Further complicating matters is the development of resistant mutations. The HIV virus is remarkably fast in adapting to the antiretroviral drugs that patients take and at some time after initiating ART, a mutation will emerge in the patient's system that is resistant to the drugs that the patient is taking. The event of a resistant strain emerging is known as treatment fail-

ure. Once a patient has failed a regimen, they will receive significantly less benefit from taking it and generally the only course of treatment would be to begin taking a new drug regimen.

2.1 Overview of Individual HIV Model

The *individual* model is a previously developed probabilistic simulation of the natural history of HIV disease that was specifically designed to model disease progression, treatment failure and long term outcomes in a biologically realistic manner.(Braithwaite et al. 2005, Braithwaite et al. 2006) The individual HIV model simulates a cohort of patients one at a time and follows them until death. At each time interval (one month) the model examines the health characteristics of each patient (CD4 count, viral load, accumulated mutations, treatment status, etc.) and determines how each patient's health will progress in the subsequent time period. In addition to updating the health states of each patient, the model calculates HIV-related mortality as a function of health.

A unique characteristic of this model is that it models the development of antiretroviral resistance acquisition biologically. This model has been calibrated using data from the Veterans Aging Cohort Study (VACS) and has closely reproduced Kaplan-Meier curves for time to treatment failure and survival, (Braithwaite et al. 2005) provided accurate 3-year mortality estimates for various initial values of age, CD4 count, and viral load for a distinct cohort from the calibration set, and was able to predict the rate of accumulation of resistant mutations (Braithwaite et al. 2005). The model has been sued to predict the influence of alternative starting thresholds (Braithwaite et al. 2008) and the effect of variable adherence on survival and quality of life.

2.2 Overview of Population HIV Model

We modified the individual model to consider a cohort of patients simultaneously instead of sequentially. This parallelization allows us to consider the entire population of patients and explicitly model the effects of drug scarcity within a closed population of patients and determine its effect on the population's overall survival. Because the purpose of this work is to investigate the effect of resource constraints on population treatment outcomes, we have chosen to maintain the HIV population and the proportion of the population that there are sufficient resources to treat as constant. Therefore, once a patient dies, another patient is immediately generated to replace him/her. When a patient dies or the model determines that a patient has failed his final ART regimen, a dose becomes available. At this point the model determines which patient will receive the newly freed dose according to a specified treatment policy.

2.3 Management of ART Decisions

In resource-poor environments the main issue facing decision makers is which patients should be treated because there are often more patients who would benefit from treatment than can be treated given available resources. For this reason, the WHO recommends a rule of rescue (RoR) type treatment plan for distributing available doses of ART. Under the RoR policy, when drugs become available they are given to the sickest (in terms of CD4 count) patients in the population. Typically, patients remain on a medication until death, although some recommendations suggest discontinuing a drug after resistance has developed and the regimen has failed. The WHO currently has guidelines for determining when virological failure and immunological failure. Because VL measurement are typically prohibitively expensive in resource-poor environments, we use the WHO guidelines for immunological failure based on CD4 count alone as the criteria for switching between drug regimens. The WHO provides 3 measurable events that may be used to determine that an ART regimen has failed:

- CD4 count below 100 after 6 months of therapy (WHO-1)
- A return to, or a fall below, the pre-therapy CD4 baseline after six months of therapy (WHO-2)
- A 50% decline from the on-treatment peak CD4 value (WHO-3)

Each of these measures defines different times of treatment failure for individual patients. We examine how each of the above metrics for defining treatment failure affects the overall life expectancy of the population. Once patients are started on ART they are not switched off of it until death or treatment failure has occurred and another patient requires the dose. We also assume that patients with CD4 count above 350 are not eligible for starting ART.

To model the long-term effect of drug scarcity we assume the following: 1) the population size is fixed, and 2) new patients enter the model only upon the death of a previous patient. These assumptions lead to a population of a fixed size at all times. We consider this situation because it allows for the most accurate estimate of a particular coverage level. By allowing the population to grow in size, we introduce two possible methods to address coverage. First we could fix the coverage level

regardless of the population size. This ensures the same percentage of patients will be treated at all times, however, this provides the same estimates of survival as a fixed population. Second, we can fix the raw number of drugs available and allow new patients to enter the population. This is desirable because it can capture the effect of the epidemic growing in a particular area. However this results in survival estimates gradually decreasing due to the coverage levels declining as the population increases.

We are able to estimate average survival in the population under different treatment policies outlined by the WHO for different levels of coverage (the percentage of the population that can be treated). Each WHO policy outlines a clinical event that signals to the decision maker that a resistant mutation has occurred and that the patient has failed that particular regimen. We can use these events within the simulation to determine the effect each policy will have on the overall survival of the population.

2.4 Equitable Access to ART

In addition to survival, we can examine a somewhat coarse measure of the equity associated with a given treatment policy using this model. In settings where access to treatment is very low, it may be the case that the treatment policy that achieves the best survival may seem unfair to patients who are waiting to begin ART. For example, if our coverage level is 10%, we may want to examine policies that perform reasonably well, but allow for more patients to be treated. This could lead to an improved quality of life for a greater number of patients. To approach this we can calculate the actual coverage (total number of patients receiving therapy/total number of patients) for various levels of initial coverage for each of the treatment policies.

3 RESULTS

Our results show that for different coverage levels, different policies achieve the highest survival. Low coverage levels (0%-30%) lead to the WHO-3 (50% CD4 decline from on-treatment peak) policy performing the best. With moderate coverage (30%-60%) the WHO-2 (CD4 count falls below pretreatment level after 6 months of treatment) policy outperforms the others. These ranges are important for two reasons. First, we can see from the plot, 78% of resource-constrained countries, have coverage levels less than 60%. Second and more importantly, the largest gap in performance among the WHO policies occurs within this range. This can be explained by the large difference in the time spent on therapy across the WHO policies for different coverage levels.

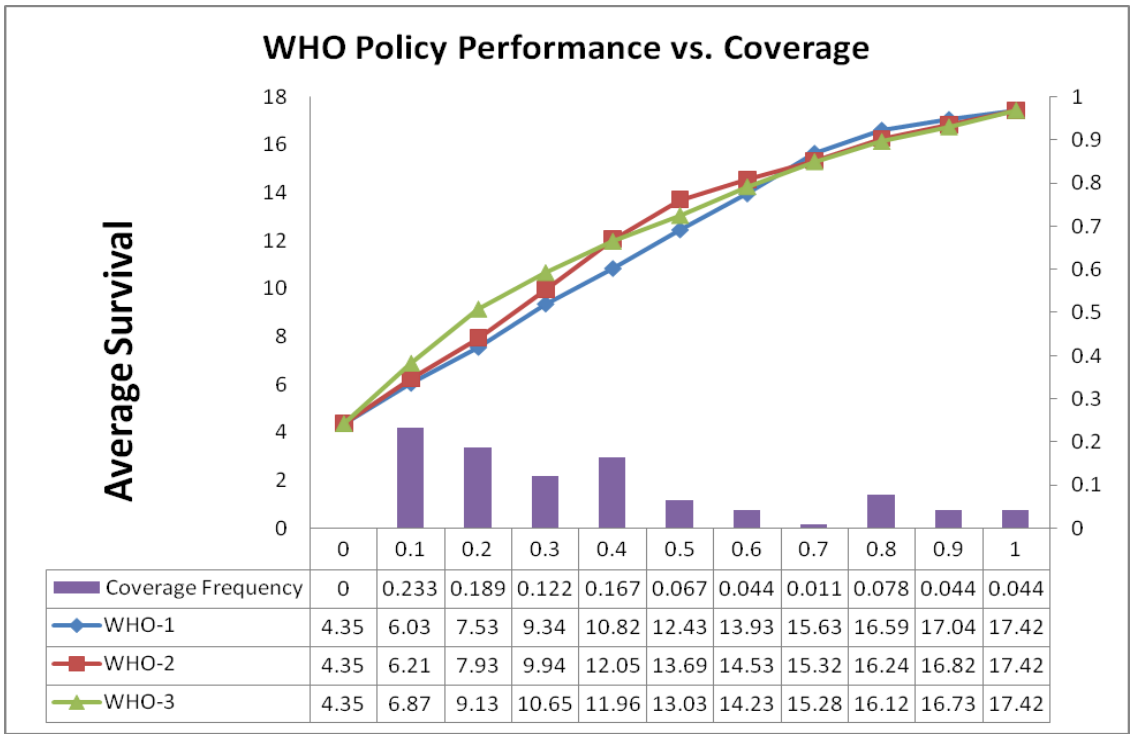

	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Coverage Frequency	0	0.233	0.189	0.122	0.167	0.067	0.044	0.011	0.078	0.044	0.044
WHO-1	4.35	6.03	7.53	9.34	10.82	12.43	13.93	15.63	16.59	17.04	17.42
WHO-2	4.35	6.21	7.93	9.94	12.05	13.69	14.53	15.32	16.24	16.82	17.42
WHO-3	4.35	6.87	9.13	10.65	11.96	13.03	14.23	15.28	16.12	16.73	17.42

Figure 2: Performance of each WHO policy across coverage levels with coverage histogram

For low coverage levels (0%-30%), the WHO-3 policy had the least amount of time spent on therapy among the treated patients. This led to more patients having access to drugs and accounted for higher average survival among the three policies. With moderate coverage (30%-60%), we find that the WHO-2 policy leads to a slightly higher time on therapy than the WHO-3 policy and produced the best survival in that coverage range. In high coverage levels (70%-100%) the WHO-1 (CD4 count falls below 100 after 6 months on treatment) policy performed the best and also led to the longest time spent on treatment. This is intuitive since in high coverage levels, we would desire a longer amount of time spent on therapy since there is more of an opportunity for drugs to become available. Thus even with policies that yield very large amounts of time spent on therapy, we would expect to see relatively high estimates of expected survival. Conversely, if we apply a policy with a low amount of time spent on therapy, we would be removing people from treatment relatively early which will result in a reduced estimate of expected survival.

We can also see that the largest gap in policy performance occurs when coverage is less than 50%. This can be explained by the differences in time spent on treatment. In these coverage levels, the Nonetheless, this motivates the need for proper management of antiretrovirals in resource-constrained environments since most countries that experience drug shortages have coverage levels in this range.

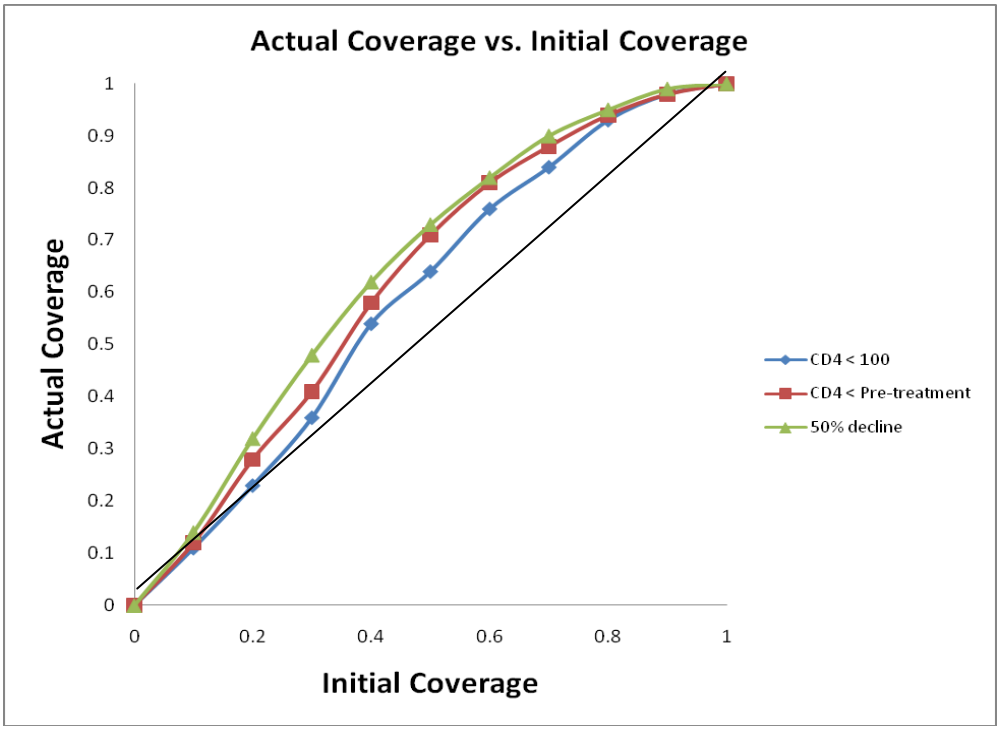

Figure 3: Actual percentage of patients treated for each policy

These results seem to suggest that there are significant variations in the performance of the WHO's recommendations for HIV treatment in resource-poor environments. Our findings suggest that the best policy is strongly dependent upon the coverage level of a particular population. The performance of a policy also seems to be strongly linked to the amount of time on treatment that a particular policy induces.

4 CONCLUSIONS

This study has shown that there may be significant differences in the WHO treatment recommedations. Our results are particularly interesting due to the large performance gap observed among the policies when coverage is less than 40%. Nearly 80% of countries experiencing drug shortages can treat less than 40% of their patients. Our work represents a step forward in modeling the complex dynamics associated with the HIV infection in populations where antiretrovirals are a scarce resource. These interactions among patients are important factors that must be taken into account when designing treatment policies. These results may not be generalizable to sub-Saharan populations since the simulation was calibrated using data from the

United States. African populations would experience much different mortality rates and thus the estimates of survival may be reduced in such situations.

REFERENCES

Braithwaite RS, Justice AC, Chang CC, Fusco JS, Raffanti SR, Wong JB, Roberts MS. Estimating the proportion of individuals infected with human immunodeficiencyF virus who will die of comorbid diseases. *American Journal of Medicine.* 2005;118(8):890-98.

Braithwaite RS, Shechter S, Roberts MS, Schaefer A, Bangsberg DR, Harrigan PR, Justice AC. Explaining variability in the relationship between adherence to antiretroviral medications and HIV mutation accumulation. *Journal of Antimicrobial Chemotherapy*. 2006;58 (5): 1036-43

Braithwaite RS, Shechter S, Chang CCH, Schaefer A, Roberts MS. Estimating the rate of accumulating drug resistance mutations in the HIV genome. *Value in Health*. 2007; 10 (3): 201-213

Braithwaite RS, Roberts MS, Chang CCH, Goetz MB, Gibert CL, Rodriguez-Barradas MC, Shechter S, Schaefer AJ, Nucifora K, Koppenhaver R, Justice AC. The influence of alternative thresholds for initiating HIV treatment on life expectancy and quality-adjusted life expectancy: A decision model. *Annals of Internal Medicine.* Feb 5;148(3):178-85, 2008

Braithwaite RS, Roberts MS, Goetz MB, Gibert CL, Rodriguez-Barradas MC, Nucifora K, Justice AC. Do benefits of earlier antiretroviral treatment initiation outweigh harms for individuals at risk for poor adherence? *Clinical Infectious Diseases*. In Press

Saka, G., J. E. Kreke, A. J. Schaefer, D. C. Angus and M. S. Roberts, "Predicting Disease Progression using Dynamic Microsimulation in Pneumonia-related Sepsis," 2007. *Critical Care*11:R65.

Schackman BR, Freedberg KA, Weinstein MC, Sax PE, Losina E, Zhang H, et al. Cost-effectiveness implications of the timing of antiretroviral therapy in HIV-infected adults. Arch Intern Med. 2002;162:2478-86. [PMID: 12437408]

Shechter, S. M., C. L. Bryce, O. Alagoz, J. E. Kreke, J. E. Stahl, A. J. Schaefer, D. C. Angus, M. S. Roberts, "A Clinically Based Discrete Event Simulation of End-Stage Liver Disease and the Organ Allocation Process," 2005. *Medical Decision Making,* volume 25, number 2, pp. 199-209.

Antiretroviral therapy for HIV infection in adults and adolescents, World Health Organization. http://www.who.int/hiv/pub/arv/adult/en/index.html. Accessed May 2009

HIV/AIDS Data and Statistics, World Health Organization . http://www.who.int/hiv/data/en/. Accessed May 2009.

AUTHOR BIOGRAPHIES

ROBERT T. KOPPENHAVER is a PhD student in the Industrial Engineering Department at the University of Pittsburgh. His email is <rzk9@pitt.edu>.

ANDREW SCHAEFER is an Associate Professor of Industrial Engineering within the Department of Industrial Engineering at the University of Pittsburgh. His email is <schaefer@pitt.edu>

R. SCOTT BRAITHWAITE is an Assistant Professor of General Internal Medicine at Yale University. His email address is <ronald.braithwaite@med.va.gov>.

MARK ROBERTS is Professor of Medicine, Health Policy and Management and Industrial Engineering and Chief of the Section of Decision Sciences and Clinical Systems Modeling in the Division of General Medicine at the University of Pittsburgh. His email address is <robertsm@upmc.edu>