Proceedings of the 2004 Winter Simulation Conference
R .G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

MODELING THE PROGRESSION AND TREATMENT OF HIV

Steven M. Shechter
Andrew J. Schaefer

1048 Benedum Hall
Department of Industrial Engineering
University of Pittsburgh
Pittsburgh, PA 15261, U.S.A.

R. Scott Braithwaite
Mark S. Roberts

200 Lothrop Street, MUH E820
Division of General Internal Medicine
University of Pittsburgh
Pittsburgh, PA 15261, U.S.A.

ABSTRACT

Current treatment of HIV patients is based on various guidelines that have changed with the advent of newer antiretroviral therapies and the emergence of resistance to them. However, there remains uncertainty over the best time to initiate HIV therapy or when to switch. Observational cohort studies or clinical trials are limited in the number of scenarios they can examine, whereas simulation modeling is well suited for considering various treatment policies. We describe a Monte Carlo simulation of a cohort of HIV positive patients that explicitly models two key components of HIV progression: adherence and the acquisition of resistance. Simulation results closely match cohort statistics such as survival time and length of time on the first three treatment regimens. We also describe sensitivity analyses and experiments such as testing the effects of starting therapy at different levels.

1 INTRODUCTION

HIV treatment has evolved considerably in the last twenty years. Until the late 1990s, the standard of care was the use of a single antiretroviral drug, and because the HIV virus mutates rapidly, monotherapy was not very effective due to the emergence of resistant strains (Shernoff and Smith 2001). As a result, researchers developed triple-drug therapies (also known as cocktail therapy, or Highly Active Antiretroviral Therapy (HAART)) in the late 1990s that allowed for a multi-pronged attack against the HIV virus. There are four major classes of antiretroviral drugs (see Appendix for a glossary of terms): protease inhibitors (PIs), nucleoside reverse transcriptase inhibitors (NRTIs), non-nucleoside reverse transcriptase inhibitors (NNRTIs), and more recently, fusion inhibitors (Peiperl and Coffey 2004). HAART has significantly reduced the mortality rates for HIV patients (Palella et al. 1998).

With the advent of new HIV therapies, the optimal use of them has also become less clear. Initially, the medical community took a "hit hard, hit early" strategy, but more recently the focus has shifted more to a "hit hard, but only when necessary" strategy (Harrington and Carpenter 2000). This is reflected in the changing guidelines for HIV therapy from recommending starting patients on HAART whenever their CD4 count falls below 500 to waiting until it falls below 350 (Stine 2003). No clinical trial data support starting above this level (Harrington and Carpenter 2000).

Simulation is an invaluable tool for treatment modeling. While randomized controlled trials (RCTs) are a necessary and final test of a new treatment policy, they are not practical for testing a wide variety of alternatives because of costs, sample size considerations, and possible ethical concerns. Cohort observations are also limited in scope and, additionally, are subject to the biases inherent in patients falling into the different branches of the cohort. All of these concerns are mitigated with a valid computer model of the clinical process.

In this paper, we describe the development of a complex Monte Carlo simulation that allows for testing a wide variety of assumptions and treatment policies for HIV patients. It is based on data from a cohort of antiretroviral naive patients from several clinics and hospitals in the United States. Unlike previous HIV simulation models (Wein, Zenios, and Nowak 1997; D'Amato, D'Aquila, and Wein 1998; Freedberg et al. 2001; Richter et al. 2002) our model explicitly considers the development of resistance and the effect adherence to the prescribed drug regimen has on this development. As indicated above, resistance is a major problem in effective HIV therapy. Additionally, HAART is a difficult drug regimen to comply with because of regimen complexity, side effects, and psychosocial issues such as depression and stress (Chesney 2003). One study demonstrated a significant correlation between better adherence levels and lower mortality rates (Carmona, Knobel, and Guelar 2000).

In Section 2 we describe the model in more detail. Section 3 provides validation results while section 4 reports results from sensitivity analyses. In Section 5, we describe some of the results from different policy experiments. Finally, we provide conclusions and directions for future research in Section 6.

2 MODEL DESCRIPTION

In the medical decision making literature, our model would be labeled a Monte Carlo Markov cohort microsimulation in which individual patients enter the model independently of each other, and events are updated at fixed time intervals (in our case, monthly). It is common to model disease progression at fixed time intervals, and because HIV progression in individual patients is our primary concern, our model does not take the form of more traditional discrete-event simulation models which consider competition for resources and random times until events.

2.1 Model Flow

Patients enter the model one at a time with a CD4 count, viral load, and age generated from distributions based on the actual cohort of patients. An entering patient is then updated monthly until he dies, at which point the next patient goes through the same routine, with a different stream of random numbers. In this way, we can run many replications of patients starting with the same CD4 count, viral load, and age to estimate the average outcomes for patients with those initial characteristics. The basic flow of the model is shown in Figure 1.

At the start of each monthly cycle, we determine if the patient should start HAART or not. The main factor that drives the start decision is the patient's CD4 count; if the patient's CD4 count falls below a pre-defined constant, the patient begins therapy. He then remains on therapy until he either runs out of available therapies or dies.

If the patient is on HAART at the start of the month, we determine the level of adherence exhibited by the patient over that month. We assumed that patients who fail to take one drug at a particular time are more likely to miss the other drugs at that same time, and we implemented this by increasing the probability of nonadherence to the other drugs if the patient misses any one of the drugs.

Once we know if the patient is on a HAART regimen and his level of adherence and resistance, we update the CD4 count, viral load, and age that change over the course of that month. CD4 and viral load changes were obtained from regression equations based on the patient cohort. The changes depend, among other things, on whether the patient is on HAART, the level of adherence and existing resistance, and which regimen number the patient is on. With these updated values, we determine if the patient dies of HIV or non-HIV-related causes that month. Because the success of HAART has turned HIV into a chronic disease (Selwyn and Rivard 2003), a significant number of HIV patients are dying from causes not related to HIV (Cohen et al. 2002). The probability of dying from HIV-related causes is a function of age, time on HAART, CD4 level, and viral load, and the risk of dying from non-HIV-related causes is a function of age.

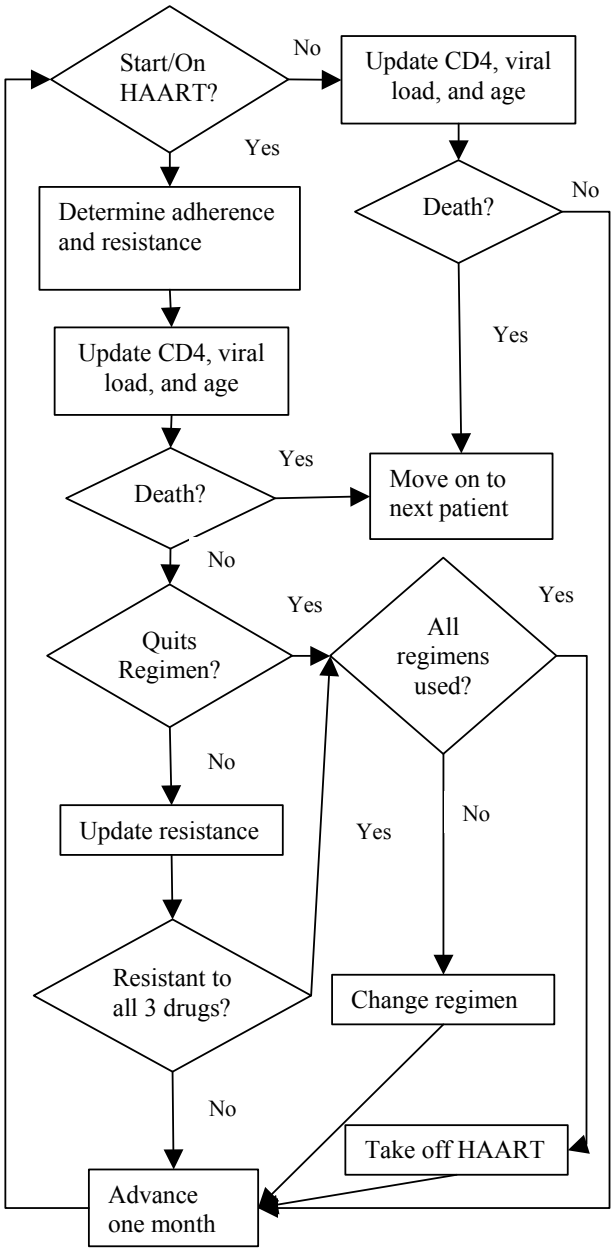

Figure 1: Overview of Model Flow

If the patient does not die, then we determine if the patient changes to a different drug regimen. This can happen in two ways: 1) the patient decides on his own or with his doctor to stop taking that particular regimen completely, or 2) the current regimen has become ineffective because there is resistance to all three drugs the patient is taking. If Case 1 occurs, then the patient gets assigned a new regimen and advances to the next month.

To determine if Case 2 occurs, we first need to determine if any new resistant mutations developed this month. We do this by transforming the mutation rate (based on the patient's viral load) into a probability of observing a muta-

tion that month. Then we consider selection pressure on the different drug classes mentioned earlier and determine if the mutation is resistant to the drug class that was selected for (we do not consider the fusion inhibitors in our model as these were just approved by the FDA in 2003, and therefore little data is available on them (Peiperl and Coffey 2004). Cross resistance is often a problem in HIV therapy, so if the model determines there is resistance to one of the drugs in the selected class, it checks for cross-resistance to the other drugs in the same class (if the patient is taking more than one drug from that class). Cross resistance is especially a problem for the PI and NNRTI drug classes (Centers for Disease Control and Prevention 2002). The probabilities of primary resistance and cross resistance are based on drug class. If, after updating the resistance levels, the virus is resistant to all three drugs the patient is taking, the patient get assigned a new regimen.

New regimens are assigned by considering the cumulative number of mutations to each of the three drug classes and determining which class is least likely to encounter a resistant mutation. The patient is assigned two drugs from that class and a third from one of the remaining two. Upon receiving a new regimen, we assign a new regimen-specific probability of adherence and continue on to the next month. Patients are limited to at most 8 regimens, and we assume they are always on HAART from the time they start until they either die or exhaust all available regimens.

2.2 Model Statistics

The model outputs a number of statistics based on averages of the total number of patient runs. Of primary importance are survival time, time until treatment failure of the first three regimens, average time on HAART, the percentage who die from HIV vs. non-HIV-related causes, the proportion of regimen changes resulting from patients stopping on their own vs. triple-drug resistance, and the proportion who die within 1, 3, and 5 years. The model also generates survival and time until treatment curves for easy comparison with the curves generated by the real cohort data (discussed in Section 3).

2.3 Variance Reduction

One of the primary reasons for building a medical simulation is to have a fairly inexpensive, quick, and risk-free method of evaluating alternative treatment policies. In doing so, we want to reduce the variance on the differences in outputs between policies so that we can have greater confidence that the observed differences are real. To that end, we implemented the variance reduction technique of common random numbers (CRN) (Law and Kelton 2000).

The key to the CRN technique is to use the same random numbers for similar reasons between simulations. For example, if the first patient in the simulation of Policy A has a probability p_{1A} of dying from HIV in month 1 and we generate a uniform[0,1] random number u_1 to determine if he dies, then we want to use that same u_1 to determine if the first patient in the simulation of policy B dies from HIV with probability p_{1B} in the first month.

To facilitate this method, we dedicated one random number stream to HIV deaths, one to non-HIV deaths, one to adherence, and one to resistance. Furthermore, because different policies may induce different survival times, we had to make sure that each new patient started at the same place in the dedicated random number stream. Since one does not know a priori how many random numbers a given policy will use, we also dedicated sections within each stream to each patient. For example, we use one random number per month from the HIV death stream until the patient dies. Therefore, we can safely assume that we need at most 100*12=1,200 random numbers from that stream for each patient. Then if patient 1 from Policy A lives 120 months and the same patient from Policy B lives 133 months, they each use 120 and 133 random numbers from the HIV death stream, respectively. Then patient 2 from each policy starts with the 1,201st random number from the HIV death stream in the first month.

3 VALIDATION

Resistance and adherence—critical to HIV progression—are also components for which it is difficult to get good measures. We therefore altered the mutation rate and probability of adherence within reasonable bounds until the model outputs matched closely with the cohort outputs with respect to overall survival time and time until treatment failure of the first three regimens (see Figures 2 and 3 for the survival curve and the curve for time until treatment failure of the first regimen. The curves for the second and third regimen are not shown. Figures are reprinted with permission from Braithwaite et al. 2004). The probability of adherence is with respect to a single drug of the three-drug regimen. For these runs, we ensured simulated patients started on HAART right from the beginning to correspond with the cohort patients. The close proximity of each of the curves gave us

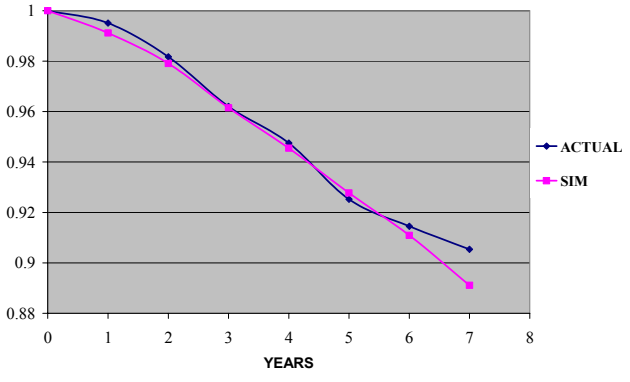

Figure 2: Overall Survival Time

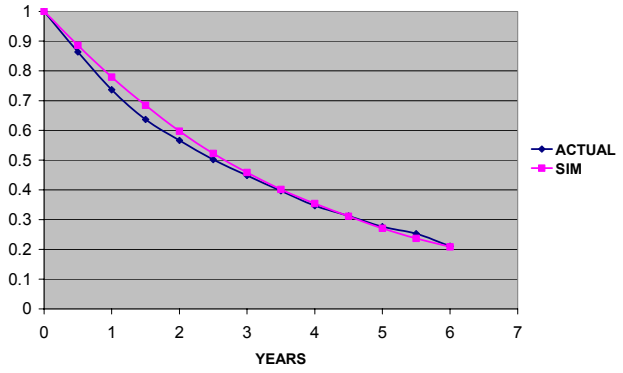

Figure 3: Time until Failure of Regimen 1

confidence that the simulation was a reasonable representation of the patients in our cohort.

4 SENSITIVITY ANALYSES

Calibrating a model to two parameters does not guarantee that the model is accurate with respect to those or other parameters in the model. Therefore, we performed one-way sensitivity analyses on some of these factors to examine their effects on output. Table 1 gives the overall mean survival time (and the half-length of the 95% confidence interval) for the baseline run and the baseline values of certain parameters. The mut_rate represents the average number of mutations per year. The prob_comp parameter gives the probability of a patient complying to any one drug in a given month, and assoc_comp indicates the degree to which not complying to one drug affects the probability of not complying with the other drugs in the regimen (if assoc_comp = 0, then there is no correlation, and if assoc_comp = 1, then the patient will certainly not take the other drugs). Prob_mut_res gives the probability that a new mutation is actually resistant to at least one of the drugs still susceptible to resistance. The table then shows how the survival time changes as these parameters change while holding the other parameters at their baseline values. Again, for these runs, patients start on HAART upon entering the model.

The assoc_comp parameter represents the clustering of adherence between drugs at the same time—i.e., the degree to which not complying with one drug at a certain time causes the person to not comply with the other drugs at the same time. A value of 0 indicates the there is no association whereas a value of 1 indicates 100% nonadherence to other drugs if there is nonadherence to one drug. Prob_mut_res gives the probability that a given mutation is actually resistant to one of the drug classes (currently we assign the same probability regardless of the drug class).

Keeping the chosen parameters within reasonable bounds, none of them changed the overall survival by more than 1.5 years. Furthermore, because the model calibrated well with the baseline values for mut_rate and prob_comp, the relative lack of sensitivity to these parameters gave us further confidence in using the derived values.

Table 1: Baseline and Sensitivity Analyses Output

Baseline Values	
mut_rate	0.01
prob_comp	0.755
assoc_comp	0.9
prob_mut_res	0.5
survival	**22.86 (.24)**

Parameter	**Value**	**Survival**
mut_rate	0.005	24.09 (.25)
mut_rate	0.015	21.97 (.24)
prob_comp	0.7	21.75 (.24)
prob_comp	0.8	23.97 (.25)
assoc_comp	0.5	22.60 (.24)
assoc_comp	1	23.09 (.24)
prob_mut_res	0.3	23.88 (.25)
prob_mut_res	0.7	22.12 (.24)

5 EXPERIMENTS

5.1 Effects of Starting Criteria on Total Lifetime

As mentioned in the introduction, one of the most important questions in HIV therapy is when to begin therapy. This is usually framed in terms of what CD4 threshold to wait until to start therapy. Commonly analyzed thresholds are 200, 350, and 500.

Figure 4 shows the median survival times from waiting until these three thresholds to begin therapy for various hypothetical patients. We assumed each patient started with a CD4 of 500 and varied the starting age to be 30, 40, or 50 and the starting viral load to be 4, 5, or 6. For each of these strata, we simulated 10,000 patients to estimate the average response for those types of patients. As can be

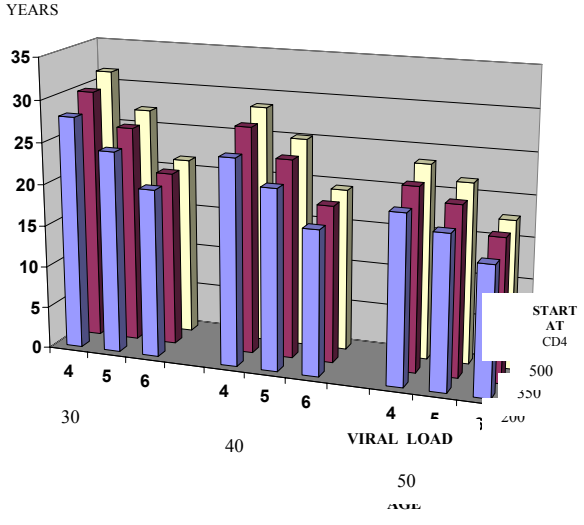

Figure 4: Median Survival Times by Treatment Initiation Threshold

seen, for each type of patient, starting therapy earlier results in a higher median survival time. For example, for a 50 year old patient with a viral load of 4 the estimated median survival time is 20.25, 22, or 23.42 years, for starting therapy at CD4 counts of 200, 350, and 500, respectively.

5.2 Effects of Starting Criteria on Quality Adjusted Life Years

Adverse side effects of HAART, which both decrease quality of life and can increase the risk of death, may cause physicians and patients to consider delaying the start of therapy. Medical decisions often involve a tradeoff between quality of life and quantity of life which is why a commonly used outcome measure in medical decision making is the total quality adjusted life years (QALYs) (Drummond et al. 1997). Our baseline model considered total life years, which may not be the best measure upon which to base decisions. We therefore wanted to incorporate quality-adjustments into our model to see how this affects starting decisions. It seems plausible that if the quality of life from being on HAART is low enough, it may make sense to delay the start of therapy.

To generate QALYs, we included utility weights for two different states in the model: HIV off of HAART, and HIV on HAART. Future work will consider the quality of life at finer levels of detail such as considering what CD4 category the patient is in. We assumed an off-HAART utility of .8 and tested two on-HAART utilities of .5 and .7. We also considered the increased risk of non-HIV-related death when patients are on HAART by multiplying that death rate by 3.

Figure 5 shows the gain or loss in median QALYs (z-axis) by starting therapy at a CD4 count of 500 vs. starting at 200 for each combination of age and viral load (x-axis) and each on-HAART utility of .5 and .7 (y-axis). We assume all patients start with a CD4 of 500. When the patient's utility for being on HAART is .7, then it is still better to start therapy at 500 for all categories except for the age = 50, viral load = 4 category (a gain of .317 QALYs by starting at 200). For the same age and viral load, if the patient's utility for being on HAART is .5, then there is an even bigger gain in QALYs (1.44) by waiting until the CD4 falls below 200. Furthermore, there are other categories for the .5 utility for which it appears better to delay therapy. A natural question is to ask which starting threshold results in the optimal QALYs, which we leave for future research. Finally, we note that if we just considered total life years, it was always better to start therapy earlier even with the toxicity multiplier of 3 (data not shown). Therefore, the variability of patient utilities can play an important role in the decision making process.

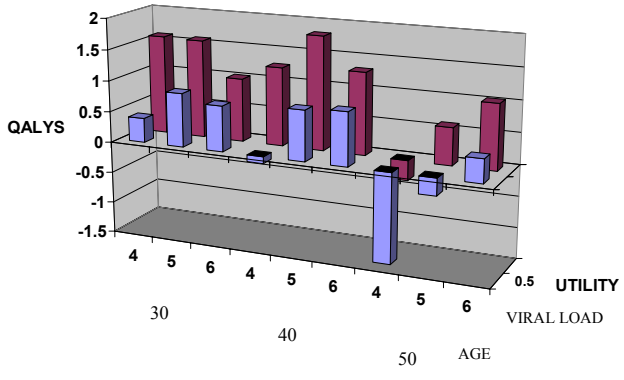

Figure 5: Difference in QALYs from Starting at 500 – Starting at 200

6 CONCLUSIONS

We have described our simulation of HIV patients undergoing HAART until they either die (from HIV or non-HIV-related causes) or they have exhausted all reasonable regimens. We model the monthly changes in critical components of HIV progression (CD4 count, viral load, adherence, and resistance) and track numerous output such as survival time, time until treatment failures, proportion of deaths attributable to HIV vs. non-HIV causes, and average time on HAART. The model and output can be used to compare different treatment policies such as when to start or switch therapy.

Our model can easily be changed to include future enhancements. Some of these include a more detailed modeling of the available regimens, resistance and adherence. For example, we currently consider only three major classes of drugs and do not consider specific drugs within each class. Along those lines, we want to model specific mutations to specific drugs. Also, we would like to model better the dependency of adherence on the specific drugs being taken.

This model can provide insights into a variety of therapeutic decisions regarding HIV care. The value in our simulation model is that we can test alternatives with little cost, in a short amount of time, and without risk to real patients. After incorporating more enhancements to our model, we hope that it can be used as a tool to guide clinical trials.

ACKNOWLEDGMENTS

The authors would like to thank Victoria Gettinger for her help in creating some figures and obtaining some references for this paper.

APPENDIX: GLOSSARY OF TERMS

Antiretroviral drugs: Drugs designed to stop or suppress retroviruses, one of which is HIV.

CD4 cells: White blood cells that help the body fight off infection. These are the cells that the HIV virus attacks.

CD4 count: The number of CD4 cells per microliter of blood.

Viral load: The amount of HIV RNA per milliliter of blood

REFERENCES

Braithwaite, R.S., A.C. Justice, C.H. Change, J.S. Fusco, S.R. Saffanti, J.B. Wong, and M.S. Roberts. 2004. Estimating the future burden of non-HIV related mortality in HIV patients. Submitted to *Annals of Internal Medicine*.

Carmona, A., H. Knobel, and A. Guelar. 2000. Factors influencing survival in HIV infected patients treated with HAART [Abstract TuOrb417]. Presented at 13th International AIDS Conference, Durban, South Africa, July 9-14.

Centers for Disease Control and Prevention. 2002. Guidelines for using antiretroviral agents among HIV-infected adults and adolescents: recommendations of the Panel on Clinical Practices. *MMWR* 51 (RR-7): 1-26.

Chesney, M. 2003. Adherence to HAART regimens. *AIDS Patient Care and STDs* 17 (4): 169-177.

Cohen, M.H., A.L. French, L. Benning, A. Kovacs, K. Anastos, M. Young, H. Minkoff, and N.A. Hessol. 2002. Causes of death among women with human immunodeficiency virus infection in the era of combination antiretroviral therapy. *The American Journal of Medicine* 113 (2): 91-98.

D'Amato, R.M., R.T. D'Aquila, and L.M. Wein. 1998. Management of antiretroviral therapy for HIV infection: modelling when to change therapy. *Journal of Antiviral Therapy* 3: 147-158.

Drummond, M.E., B. O'Brien, G.L. Stoddart, and G.W. Torrance. 1997. *Methods for the Evaluation of Health Care Programmes*. 2nd ed. Oxford: Oxford Medical Publications.

Freedberg, K.A., E. Losina, M.C. Weinstein, A.D. Paltiel, C.J. Cohen, G.R. Seage, D.E. Craven, H. Zhang, A.D. Kimmel, and S.J. Goldie. 2001. The cost effectiveness of combination antiretroviral therapy for HIV disease. *The New England Journal of Medicine* 344 (11): 824-831.

Harrington, M. and C.J. Carpenter. 2000. Hit HIV hard, but only when necessary. *Lancet* 355: 2147-2152.

Law, A.M. and W.D. Kelton. 2000. *Simulation Modeling and Analysis*. 3rd ed. Mcgraw-Hill, Inc.

Palella, F.J., K.M. Delaney, A.C. Moorman, M.O. Loveless, J. Fuhrer, G.A. Satten, D.J. Aschman, and S.D. Holmberg. 1998. Declining morbidity and mortality among patients with advanced human immunodeficiency virus infection. HIV Outpatient Study Investigators. *New England Journal of Medicine* 338(13): 853-60.

Peiperl, L. and S. Coffey. 2004. Overview of Antiretroviral Drugs [online]. Available online via <http://hivinsite.ucsf.edu/InSite?page=ar-drugs> [accessed April 2, 2004].

Richter, A., B. Hauber, K. Simpson, J.A. Mauskpof, and D. Yin. 2002. A Monte Carlo simulation of AIDS treatment regimens. *Pharmacoeconomics* 20 (4): 215-224.

Selwyn, P.A. and M. Rivard. 2003. Palliative care for AIDS: Challenges and opportunities in the era of Highly Active Anti-Retroviral Therapy. *Journal of Palliative Medicine* 6: 475-487.

Shernoff, M. and R.A. Smith. 2001. HIV Treatments: A History of Scientific Advance [online]. Available online via <http://www.thebody.com/bp/jul01/treatments.html> [accessed April 2, 2004].

Stine, G.J. 2003. *AIDS Update*. Upper Saddle River, New Jersey: Prentice Hall.

Wein, L.M., S.A. Zenios, and M.A. Nowak. 1997. Dynamic multidrug therapies for HIV: A control theoretic approach. *Journal of Theoretical Biology* 185: 15-29.

AUTHOR BIOGRAPHIES

STEVEN M. SHECHTER M.S. is a PhD student in Industrial Engineering at the University of Pittsburgh. He received his Masters degree in Operations Research from Georgia Tech in 1999 and went to work for United Airlines for two years. There, he worked on a simulation model of their revenue management system. His current research interests are operations research, Markov decision processes, and simulation and their application to medical decision making problems. His email address is <sms13@pitt.edu> and his web address is <www.pitt.edu/~sms13>.

R. SCOTT BRAITHWAITE M.D. is an Assistant Professor of Medicine in the department of General Internal Medicine and Section of Decision Sciences at the University of Pittsburgh. He received his medical degree at the State University of New York at Stony Brook and completed a fellowship in medical decision making at Tufts University School of Medicine. His research interests are in decision analysis, simulation modeling and cost-effectiveness analyses applied to a variety of health care issues. His e-mail address is <braithewaiters@msx.upmc.edu>.

ANDREW J. SCHAEFER Ph.D. is an Assistant Professor in the Department of Industrial Engineering and Medicine at the University of Pittsburgh. He received an undergraduate degree and a master's degree from Rice University, and he received his PhD from Georgia Tech in 2000. His research interests include the application of operations research techniques to medical problems, stochastic optimization and integer programming. His email address is <schaefer@ie.pitt.edu> and his web address is <http://www.ie.pitt.edu/~schaefer/>.

MARK S. ROBERTS M.D., M.P.P. is Associate Professor of Medicine, Health Policy and Management and In-

dustrial Engineering and Chief of the Section of Decision Sciences and Clinical Systems Modeling in the Division of General Medicine at the University of Pittsburgh. He obtained his undergraduate degree in Economics from Harvard College, completed Medical School at Tufts University School of Medicine, and earned a Masters in Public Policy (MPP) from the Kennedy School at Harvard. His research interests include the mathematical modeling of disease, cost effectiveness and decision analytic modeling, health services research and clinical information systems/quality improvement. The modeling of complex biologic processes and evaluating the tension between realistic models of clinical process and their analytic tractability is a major area of interest and expertise. His e-mail address is <roberts@msx.upmc.edu>.