2005s-34

Early Initiation of Highly Active Antiretroviral Therapies for Aids: Dynamic Choice with Endogenous and Exogenous Learning

Pierre Lasserre, Jean-Paul Moatti, Antoine Soubeyran

Série Scientifique
Scientific Series

Montréal
Octobre 2005

CIRANO
Centre interuniversitaire de recherche
en analyse des organisations

CIRANO

Le CIRANO est un organisme sans but lucratif constitué en vertu de la Loi des compagnies du Québec. Le financement de son infrastructure et de ses activités de recherche provient des cotisations de ses organisations-membres, d'une subvention d'infrastructure du Ministère du Développement économique et régional et de la Recherche, de même que des subventions et mandats obtenus par ses équipes de recherche.

CIRANO is a private non-profit organization incorporated under the Québec Companies Act. Its infrastructure and research activities are funded through fees paid by member organizations, an infrastructure grant from the Ministère du Développement économique et régional et de la Recherche, and grants and research mandates obtained by its research teams.

Les organisations-partenaires / The Partner Organizations

PARTENAIRE MAJEUR
. Ministère du Développement économique, de l'Innovation et de l'Exportation

PARTENAIRES
. Alcan inc.
. Banque du Canada
. Banque Laurentienne du Canada
. Banque Nationale du Canada
. Banque Royale du Canada
. Bell Canada
. BMO Groupe financier
. Bombardier
. Bourse de Montréal
. Caisse de dépôt et placement du Québec
. Fédération des caisses Desjardins du Québec
. GazMétro
. Hydro-Québec
. Industrie Canada
. Ministère des Finances du Québec
. Pratt & Whitney Canada
. Raymond Chabot Grant Thornton
. Ville de Montréal

. École Polytechnique de Montréal
. HEC Montréal
. Université Concordia
. Université de Montréal
. Université du Québec
. Université du Québec à Montréal
. Université Laval
. Université McGill
. Université de Sherbrooke

ASSOCIÉ À :
. Institut de Finance Mathématique de Montréal (IFM^2)
. Laboratoires universitaires Bell Canada
. Réseau de calcul et de modélisation mathématique [RCM^2]
. Réseau de centres d'excellence MITACS (Les mathématiques des technologies de l'information et des systèmes complexes)

ISSN 1198-8177

Early Initiation of Highly Active Antiretroviral Therapies for Aids: Dynamic Choice with Endogenous and Exogenous Learning*

Pierre Lasserre†, Jean-Paul Moatti‡, Antoine Soubeyran§

Résumé / *Abstract*

Le bien-fondé d'administrer précocement des traitements antirétroviraux à haute activité (HAART) aux personnes infectées par le VIH reste objet de débats dans le monde, car leurs bienfaits à court terme peuvent compromettre les traitements futurs si se développent des souches résistantes du virus. Par ailleurs retarder le recours aux HAART comporte un coût d'opportunité thérapeutique si la santé du patient se dégrade au point qu'il ne peut plus bénéficier par la suite des traitements encore en cours de développement. Nous formulons un modèle à deux périodes où l'adoption du traitement de première période est irréversible et engage le futur, alors que des informations et connaissances nouvelles, exogènes et endogènes, déterminent les conditions entourant la décision thérapeutique de deuxième période. Paradoxalement, sous des conditions reflétant bien les enjeux du recours aux HAART, il s'avère que l'effet résistance éventuel a d'autant moins de chance d'importer pour la décision optimale, que sa gravité est élevée.

Mots clés : décisions thérapeutiques, incertitude, information, irréversibilité, traitement, apprentissage endogène, apprentissage exogène

Criteria for initiation of highly active antiretroviral treatments (HAART) in HIV-infected patients remain a matter of debate world-wide because short-term benefits have to be balanced with costs of these therapies, and restrictions placed on future treatment options if resistant viral strains develop. On the other hand, postponing the introduction of HAART may involve a therapeutic opportunity cost if a patient's health is allowed to deteriorate to such an extent of becoming unable to benefit from new treatments currently under development when they become available. We introduce a two period model where period one treatment adoption is an irreversible act with future, but uncertain, consequences. New information, both endogenous and exogenous, arises over time and shapes the conditions surrounding the second period therapeutic decision. A surprising result is that, under conditions that appear close to those surrounding the HAART debate, the magnitude of the feared resistance effect has no effect on leaves the optimal treatment decision as far as it is high enough.

Keywords: *therapeutic decisions, uncertainty, information, irreversibility, treatment, endogenous learning, exogenous learning*

Codes JEL : I18, C61

---

* Financial support from the FQRSC funds, CHRS, IFM2 and the CNRS is gratefully acknowledged. We thank Stéphane Méchoulan and the participants to the 3rd Canadian Health Study Group Conference. We also thank two anonymous referees; one of them gave us comments too helpful and insightful to be given credit properly within the paper. Forthcoming, *Journal of Health Economics*.

† Département des sciences économiques, Université du Québec à Montréal, CIRANO, and GREQAM.

‡ Faculté d'économie et gestion, Université de la Méditerranée, GREQAM, and INSERM Research Unit 379.

§ Département des sciences économiques, Université de la Méditerranée, and GREQAM, Route des Milles, 13290 Les Milles, France, tel. : 33 (0) 4 42 93 59 80, e-mail: soubey@univ-aix.fr

1. Introduction

During the last seven years, the clinical care of HIV-infected people has been substantially influenced by the introduction of the Highly Active Antiretroviral Treatment (HAART). This treatment combines several therapies associating new HIV-specific protease inhibitors (PIs), and, more recently, non nucleoside reverse transcriptase inhibitors (NNRTIs), with previously existing antiretroviral drugs, the Nucleoside Reverse Transcriptase Inhibitors (NRTI). Short-term studies (Hogg et al., 1998a; Murphy et al., 2001) have clearly proved HAART therapies to be effective in decreasing viral replication and reducing morbidity and mortality among HIV-infected patients. However, major uncertainties remain about the optimal criteria for therapy initiation, as well as the dosages and specific HAART combinations that will ensure long-term efficacy (Gallant, 2000).

Current recommendations for initiation of HAART in patients infected with human immunodeficiency virus type 1 HIV are based on the combination of two biological markers, CD4 T-lymphocyte cell counts and plasma HIV RNA levels. The relative prognostic value of each marker following initiation of therapy has not been fully characterized. Earlier guidelines were heavily based on the principle of 'hit early, hit hard', although the long-term implications of this approach were unknown (Kyriakides and Guarino, 2001). Since then, the only clear international consensus is that HAART should be initiated before CD4 cell counts become lower than 200/microL because uniformly higher rates of disease progression to death and AIDS or death among patients starting ARV therapy have been observed in patients having access to HAART below this threshold (Hogg, 2001). On the other hand, there is no clear evidence whether delay in initiation of ARV therapy (ART) until this threshold of 200/microL may lead to a poorer viral load response for patients with human immunodeficiency virus (Phillips et al., 2001). Due to lack of evidence, country guidelines may significantly differ for recommendations about patients with CD41 lymphocytes between 200 and 350 cells/mL. Differences are even more pronounced for patients with CD41 lymphocytes > 350 cells/mL and when combining CD4 thresholds with those used for viral load (Rubio, 2002; Recommendations

of the Panel on Clinical Practices for Treatment of HIV, 2002; Idemyor, 2002).

Because of increased awareness of the activity and toxicity of current drugs, the threshold for initiation of therapy has shifted to a later time in the course of HIV disease. This shift is not so much the result of any better knowledge of the influence of patients' characteristics on future outcomes from treatment, as it reflects general knowledge or worries about the effect of the drugs irrespective of the initial condition of the patient. The optimal time to initiate therapy remains imprecisely defined (Yeni et al., 2002; Delfraissy, 2002). A major issue has to do with the development of resistance, and the subsequent loss of drug activity, which may be caused by a variety of pharmacological and biological (Descamps et al., 2000) as well as behavioral factors (Paterson et al., 2000). One potential long-term consequence is cross-resistance to alternative HIV treatments not yet prescribed (Deeks et al., 1997).

The prospect of resistance is a key parameter for clinical decision-making about initiation of HAART because it threatens to make individual patients unresponsive to their first line regimen and to reduce the effectiveness of switching to other available or future regimens in case of failure of previous HAART combinations. Indeed, virological treatment failure has been reported in circa 50% of unselected patients within one year initiation of a PI-containing regimen (van Heeswijk, 2001), and such failure is more frequent when patients had received previous ARV treatment (Le Moing et al., 2002).[1]

The potential resistance induced by treating too early must be balanced against the potentially equally harmful consequence of not applying any treatment. In that case, a patient may deteriorate to such an extent that he will not be able to derive the benefits

[1]As resistance to antibiotics, transmission of HIV resistant strains to others is essentially an externality. In this paper, we do not focus on that externality facing the healthy population although it could easily be adapted to study that particular aspect of resistance. We rather focus on how the prospect of resistance induced by medication into the viruses carried by a given patient should be taken into account. When dealing with this issue, a different externality arises from the fact that direct individual treatment costs are typically sensitive to the number of individuals under therapy, and, possibly, to the experience acquired by treating other patients earlier. We clarify conditions under which, paradoxically, identical patients may be given different treatments. We also emphasize that early treatment may allow clinicians to later exploit differences in susceptibility to treatment among initially undistinguishable patients.

from a new, more advanced, therapy when it becomes available. This will be called the therapeutic opportunity cost effect.

The therapeutic decision must be made on the basis of data that are not fully matured, under major uncertainties about their consequences on HIV-infected persons. Learning takes place over time, and may take two forms: it may be endogenous, and it may be exogenous. Endogenous learning may take many forms for the decision makers responsible for administering HAART therapies; hovever an essential part of it is learning by treating, the particular form of learning by doing applying in this context. When they treat a patient, clinicians relate particular patient characteristics with particular outcomes, and the more patients they treat, the better they learn whether a particular response in a patient is random or is the sign of a systematic outcome. By treating patients, the clinicians thus learn about each individual responsiveness to treatment and the probable responsiveness of the same individual patients to further, future, treatment.

Exogenous learning is information generated independently from the clinician decision maker. For example fundamental and development research about HAART is carried out by various teams and generates information published in scientific journals, belayed by the pharmaceutical industry, etc.. That information is observed by the clinicians who make decisions on HAART treatment, who update their knowledge over time. But they do not influence the speed at which the information is generated nor its content. This information generation process is exogenous to the decision maker.

The model presented in this paper formalizes and solves the therapeutic dilemma whether or not, and when, to initiate a therapy; it also deals with the economic dilemma of weighting financial cost considerations against health benefits. It incorporates endogenous and exogenous arrival of information. The model analyses and helps understand the optimal decision process. It establishes and rationalizes some surprising results.

Although expected total costs and benefits are linear in the second period, learning by treating, by allowing the selection of patients according to expected responsiveness, makes the objective function concave. This calls for interior treatment shares in the

second period for previously treated patients, while the optimal decision for naïve patients is a corner solution. [2] Furthermore, the first period decision about the proportion of patients receiving treatment affects the second period choice. This happens via two channels. First learning by treating implies that period 2 treatment costs and benefits depends on the number of patients treated in period 1; second, irrespective of acquired knowledge by the clinician decision maker, the therapeutic benefits of treatment in period 2 depend on whether or not a patient has been treated before.

The combination of interior and corner solutions to the second period treatment decisions is important in analysing sequential health decisions involving uncertainty and learning for the following reason. A corner solution is not affected by a small parameter change; an interior solution normally is. As a result, some parameters affect the optimal decision in period 2 while some do not, depending on whether they bear on the corner solution or the interior one. Furthermore this parameter sensitivity depends on the first period decision, because the period 2 decision is conditional on it. In a sensible configuration of the relevant scientific parameters, viral resistance turns out to be irrelevant to the period 1 treatment decision, despite the fact that both the first period share and one of the second period shares are strictly interior.

Perhaps more surprising is that resistance to future therapy is more likely to be irrelevant, the higher the magnitude of the feared effect. This happens because the optimum decision rule implies complete protection against an outcome involving resistance when the magnitude of that effect is high, while it may involve only partial protection when the prospect of resistance is less dreadful.

The rest of the paper is organized as follows. In the next section we introduce a

[2] In a model that left the decision maker unable to select patients according to expected responsiveness, but had rising per patient costs of treatment in each period due to decreasing returns, a similar outcome would obtain. Although previously treated and naïve patients would be bundled within the same, convex, cost function, since they received the same treatment in the same period, expected benefits would still differ for the two groups so that a simple rule would arise for period 2: if viral resistance has been induced by Therapy 1 in the group that received it, treat naïve patients in priority, then (possibly) also treat some previously treated patients; if Therapy 1 turns out not to create viral resistance, treat previously treated patients in priority, then (possibly) also treat some naïve patients.

model that exhibits the main features just discussed: irreversible therapeutic decisions are made under uncertainty in a dynamic setup; general information about therapies unfolds exogenously during the period under scrutiny so that decisions may be regretted *ex post* even if they were not mistakes *ex ante*; simultanously information on individual responsiveness to treatment is generated endogenously by treating patients. The problem is solved by stochastic dynamic programming, starting, in Section 3, with the last period, and continuing, in Section 4, with the first period and the overall solution. Section 5 gives the intuition underlying the decision rule, explains and rationalizes why and when the resistance effect may, or may not, affect the optimal period 1 decision, and further discusses the link of the model with real options and the role of information. In the Conclusion, we go back to the general medical literature and raise issues that could not be tackled without further work.

2. Sequential therapeutic choice with intermediary revelation of information

2.1 General framework

We use a simple two-period model to mimic the choices faced by health decision maker confronted with a successive flow of new therapies with major *ex ante* uncertainties about their effectiveness. This decision maker is a clinician concerned both with patients well being and with therapeutic costs at the aggregate level.

All costs and utility levels are expressed in comparable present value units. At the beginning of the first period, a new therapy, Therapy 1, becomes available. Although its effects are not yet fully established, the clinician decides whether or not to prescribe it to some or all patients. The therapy has two effects: a known expected current utility improvement ε for the patient; and an unknown future effect on the efficacy of future treatments. Future treatments may consist in administering a second therapy, Therapy 2, which will become available in the second period, in administering Therapy 1, or in

using no therapy at all. We assume that no relevant specific information on individual patients is known to the decision maker at the beginning of period 1. Consequently, while individual responses to Therapy 1 may differ across patients *ex post*, the expected period 1 utility improvement is the same for all treated patients *ex ante*.

Unlike the current period effect, whose expected value is known, the expected effect of using Therapy 1, on future treatments will be known only at the beginning of period 2. At that time, the treatment options are: no treatment; administering old Therapy 1; administering new Therapy 2. Before a decision is made, the results of further research and experiment reveal whether Therapy 1 induces viral resistance thus reducing the benefits from both Therapy 1 and Therapy 2, or, on the contrary, leaves the patient fitter and in a better position to benefit from either therapy. If Therapy 1 turns out to create resistant viral strains, it will be labelled 'bad' (b); in the opposite case, it will be labelled 'good' (g). To repeat, the difference between a good, and a bad, realization, for Therapy 1 does not lie in its current utility effect but in its future impact through the mechanism of viral resistance and therapeutic opportunity cost. If Therapy 1 proves effective (g), that beneficial effect will be felt during the next period in the form of a better response to medication administered in period 2.

Whether Therapy 1 turns out g or b is a characteristic of the therapy, not a matter of individual patient responsivenes. At the beginning of period 1 it is assumed that Therapy 1 will turn out g with probability γ, and b otherwise. A fraction a of the total number of HIV-infected patients are prescribed Therapy 1; the remaining fraction $(1-a)$ are not treated.

The duration of the first period corresponds to the time necessary for property b or g of Therapy 1 to reveal itself, and for Therapy 2 to reach the prescription stage; these two processes are the result of exogenous scientific research and experimentation. Then period 2 begins; the therapeutic decision made at the beginning of period 2 is the last relevant decision and its effects cover the rest of the decision making horizon.

2.2 Therapy 1 affects the impact of Therapy 2

Thus, at the beginning of the second period, the information about the long-run effects of Therapy 1 is revealed and Therapy 2 becomes available.[3] The expected effect of Therapy 2 on a patient's well-being depends on whether the patient has received Therapy 1 or not and, if the patient has received Therapy 1, it depends on whether Therapy 1 has revealed itself b or g. We note E_b and E_g the corresponding expected utility gain on the average patient that has undergone Therapy 1; we note E_n the expected utility gain on an untreated, or naïve, patient. The foregoing discussion implies:

$$E_b < E_n < E_g \tag{1}$$

If Therapy 1 reveals itself g a patient left untreated in period 1 may expect less benefit from Therapy 2 than a previously treated patient, because the former has been allowed to deteriorate while the latter has been protected from various symptoms of the desease by Therapy 1 without experiencing the development of any viral resistance. This is the therapeutic opportunity cost, measured by the difference

$$E_g - E_n. \tag{2}$$

In contrast, if Therapy 1 turns out b, naïve patients benefit more from the new treatment than previously treatment patients because the former do not harbor resistant strains of the virus selected or promoted by Treatment 1. This is the resistance effect, measured by the difference

$$E_n - E_b. \tag{3}$$

[3] As pointed out by a referee, patients may not in reality receive or perceive the same information as the decision maker. This, and the fact that costs may not be perfectly internalized, may make it difficult for the decision maker to obtain patient agreement for the desired treatment.

The therapeutic opportunity cost expresses regrets experienced *ex post* for not having prescribed Therapy 1 when the latter turns out g; the resistance effect expresses regrets experienced *ex post* for having prescribed Therapy 1 when the latter turns out b.

Therapy 1 may also be applied or reapplied in period 2, with expected utility gains $e_b < e_n < e_g$.[4] In fact HAART therapies (Therapy 2) combine new HIV specific PI's and NNRTI's with previously existing antiretroviral drugs (Therapy 1). We assume that Therapy 2 incorporates the new scientific knowledge acquired during period 1, so that it does at least as well as Therapy 1 in each possible case: $e_b \leq E_b$; $e_n \leq E_n$; $e_g \leq E_g$. We further assume that cost considerations do not interfere with the dominance of Therapy 2, so that Therapy 1 will simply not be used in period 2.[5]

2.3 Costs and patient specific responsiveness

Total costs depend on the total number of patients undergoing Therapy 1, in period 1, or Therapy 2, in period 2. Discounted unit costs differ between periods because the treatments differ, the length of the periods differ[6], and discount factors differ. If costs per patient were constant, identical patients would receive identical treatment, leading to corner solutions (zero or one) for the proportions of patients administered a treatment. Decreasing returns may explain the fact that proportions strictly between zero and one are typically observed.

However another, more profound, justification for intermediate proportions of patients treated (strictly between zero and one), would be if patients differed in treatment

[4] Although it applies to patients that are naïve at the beginning of period 2, e_n is different from ε, which applies to naive patients at the beginning of period 1, because we are not assuming that periods have the same duration, and because all utility and cost magnitudes are expressed in present value terms.

[5] In fact the technological progress occuring between the two periods may not concern only therapeutic efficiency but also production technology and costs, allowing access to more sophisticated medications. In any case, this hypothesis is not crucial. Relaxing it complicates the model and its resolution without bringing any further insights.

[6] We think of period 2 as representing the whole remaining time once Therapy 2 is available, while period 1 represents a shorter period over which information about the resistance-inducing characteristic of Therapy 1 becomes available and Therapy 2 becomes operationnal.

responsiveness, and could be distinguished accordingly.[7] In fact, while the decision maker does not hold any information that would allow her to discriminate between patients at the beginning of period 1, administering Therapy 1 in period 1 may reveal specific information about individual patients. At the beginning of period 2, the clinician may then be in a position to discriminate between patients according to expected responsiveness rather than make decisions based only on the average expected effects identified in the previous subsection.

Differences in patient responsiveness, if they arise, may be described in terms of differences in therapeutic effect and/or in costs of treatment. Although this makes a difference to the patient, it does not to the decision maker, who cares about the net effect: utility minus cost. We merge these two possible channels of differentiation under a single umbrella, called cost for simplicity. Thus what we will call cost includes both the conventional supply side cost of administering the treatment, and, on the demand side, any possible utility loss or gain to a particular patient due to patient specific responsiveness to treatment.

Period 1:

Since period 1 is relatively short, we assume decreasing returns due to some fixed factors. This is the short-run situation of scare resources which clinicians usually face: increasing the number of patients undergoing therapy usually imposes a strain on fixed factors and may require the acquisition of new fixed equipments that would otherwise not be required. Thus we hypothesize an increasing convex cost function of the number of patients undergoing therapy in period 1. Although there may be differences in patient responsiveness that will translate into different *ex post* realizations, the clinician has no information that she could use to select patients accordingly. Consequently, the total expected cost of treatment *ex ante* depends only on the total number of patients: $c(z) = \frac{1}{2}cz^2$ where c is a positive parameters that reflects technology and the constraints

[7] We are grateful to an anymous referee for suggesting this possibility, as well as many other improvements to the paper.

of health services, while z is the number of patients undergoing Therapy 1. Let n and a be respectively the total population of HIV infected patients and the proportion of that population that receives Therapy 1; then $z = na$; if population is normalized to 1, and units are chosen so that $c = 1$, then $z = a$ and the total expected cost of Therapy 1 is

$$tc(a) = \frac{1}{2}a^2 \tag{4}$$

Period 2:

At the beginning of period 2, the clinician has a wider set of decision possibilities to choose from. For each possible type of patient, she can abstain from any treatment; she can prescribe Therapy 1; or she can prescribe Therapy 2. We have already made assumptions implying that using Therapy 1 in period 2 is never optimal so we will not discuss the cost of administering Therapy 1 in period 2.

Whether it is administered to previously treated patients or to naïve patients, Therapy 2 involves a common protocole for each patient. Again, in the absence of any predictable differences in responsiveness between patients, the expected total cost depends on the total number (previously treated plus naïve patients) of patients treated. However, if period 2 has a longer duration, the assumption of decreasing returns may not be appropriate.[8] Consequently, when possible differences in responsiveness are not taken into account, expected total costs of Therapy 2 are assumed to be linearly increasing in the total number of patients treated: take any patient; his expected cost of Treatment is $\bar{C}$.

Let us introduce differences in individual responsiveness. If treating some patients in period 1 allows the clinician to rank them according to expected responsiveness, then more responsive patients should be given priority over less responsive ones for treatment in period 2. If all previously treated patients are given Therapy 2, the average

[8] Convex costs in the total number of patients are compatible with the results that we are going to present, but not necessary.

cost per patient will be unchanged at $\bar{C}$; but if only the more responsive ones are treated, the average expected cost per patient will be lower than $\bar{C}$. Consequently, among the group of previously treated patients, the expected marginal cost of treating a new patient increases with A, the proportion of patients treated in period 2. The ability to select patients according to treatment responsiveness is measured as the ability to rank them in order of decreasing expected responsiveness (increasing expected cost) on a continuous, finite, cost interval. When this ability does not exist, all patients have the same expected cost of treatment. An improvement in the ability to select means that fewer and/or smaller mistakes occur in the ranking process so that the difference in expected cost between the most responsive (least expected cost) patient and the least responsive (highest expected cost) patient rises, while the average over all patients remains $\bar{C}$. We assume that the ability to rank patients improves with the number of patients treated in period 1. Precisely, as shown in Appendix 1:

Lemma 1 *if the ability to select patients according to expected individual responsiveness is non-existent when $a = 0$ and increases as a rises, and;*

if the expected cost per previously treated patient is distributed uniformly according to the number of previously treated patients given Therapy 2, and the distribution average is $\bar{C}$, [9]

then the total expected cost of treating a proportion A of na previously treated patients is (taking $n = 1$)

$$TC(A, a) = aC_0(a)A + \frac{1}{2}aA^2(C_1(a) - C_0(a))$$

where $C_0(a)$ and $C_1(a)$, respectively the expected individual cost of treating the patient considered the most responsive, and the expected individual cost of treating the patient

[9] This distribution is conditional on a. It would be interesting to specify under which conditions on the unconditional distribution of costs per patient and on the learning process, such a uniform conditional distribution of expected cost per patient would arise. In Appendix 2 we provide an example of a sufficient condition as an illustration.

considered the least responsive, are such that $\frac{C_0(a)+C_1(a)}{2} = \bar{C}$.

This lemma illustrates that the ability to discriminate between patients according to treatment responsiveness confers convexity in A to the total cost of Treatment 2. Precisely, in the case of an otherwise linear total cost function and a uniform conditional distribution of expected individual-patient responsiveness, it makes the total cost function quadratic in A.

The average of $C_0(a)$ and $C_1(a)$ is $\bar{C}$ for any a and the difference $C_1(a) - C_0(a)$ increases from zero to some positive number as a increases. For example, taking $C_0(a) = \bar{C} - \frac{1}{2}a$ and $C_1(a) = \bar{C} + \frac{1}{2}a$, gives

$$TC(A, a) = \left(\bar{C} - \frac{1}{2}a\right) aA + \frac{1}{2}(aA)^2 \tag{5}$$

This is the total expected cost function assumed to hold for previously treated patients in the rest of the paper.

A major issue in decisions on HAART therapies has to do with balancing costs and benefits: if Therapy 1 turns out b, the decision maker *ex post* wishes she had not administered that therapy; the average previouly treated patient is no longer in a position to benefit from Therapy 2, while, on the contrary, a naïve patient will benefit; this implies that a natural assumption characterizing the problem at hand is:

$$E_b < \bar{C} < E_n \tag{6}$$

The problem faced by the clinician is a collective optimization problem, namely the maximization of expected incremental utility to the whole population of HIV-infected patients, minus total expected cost of treatments.[10] Because the second-period decision is conditional on the choice made in period 1, the solution is best obtained by backward

[10] We abstract from considering the consequences of the HIV treatment decisions on other categories of patients.

induction. This means maximizing the net expected incremental utility at the beginning of the second period for each possible state of the world, considering, in each state, the decision made in period 1. Once these various second-period programs have been solved, it becomes possible to optimize the first-period optimal *ex-ante* choice.

3. Second-period optimization

At the beginning of period 2, there are two possible states of nature: Therapy 1 has turned out b, or it has turned out g; and the decision maker faces two groups of patients: previously treated patients, and naïve patients. Furthermore, the decision maker has acquired some capability to select patients according to individual expected responsiveness; this ability exists only with respect to the group of previously treated patients. In state b, the variables to choose are the proportions A_b and Λ_b of respectively na previously treated and $n(1-a)$ untreated patients; in state g, the corresponding choice variables are A_g and Λ_g.

Conditional on Therapy 1 having turned out b, and given period 1 treatment decision a, the optimum net expected value of administering Therapy 2 to proportions A_b and Λ_b of respectively na previously treated and and $n(1-a)$ untreated patients is (taking $n=1$ without loss of generality and using (5)):

$$V(a|b) = \max_{A_b,\Lambda_b} E_b a A_b - \left[\left(\bar{C} - \frac{1}{2}a\right) aA_b + \frac{1}{2}(aA_b)^2\right] + E_n(1-a)\Lambda_b - \bar{C}(1-a)\Lambda_b$$

where the first two terms give total expected utility gains common to all aA_b previously treated patients that receive Therapy 2, net of expected treatment costs including patient specific costs or utility changes; and the last two terms give total expected utility gains to naïve patients, net of expected treatment costs; as discussed earlier, naïve patients cannot be distinguished according to responsiveness so that the expected cost per patient is $\bar{C}$ for any of the $(1-a)\Lambda_b$ patients from that group given Therapy 2.

Defining $X_b = aA_b$ and $Y_b = (1-a)\,\Lambda_b$ as decision variables, the problem becomes

$$\max_{X_b, Y_b} X_b \left(E_b - \bar{C} + \frac{1}{2}a \right) - \frac{1}{2}X_b^2 + Y_b \left(E_n - \bar{C} \right), \quad 0 \leq X_b \leq a;\ 0 \leq Y_b \leq 1-a. \tag{7}$$

The problem is linear in Y_b and concave in X_b. The solution for X_b is:

$$X_b^* = \begin{cases} 0 \ , \ a \leq a_b \\ E_b - C_0(a) \ , \ a > a_b \end{cases} \tag{8}$$

where, if it exists, the critical value a_b is given by the condition:

$$a_b = -2\left(E_b - \bar{C} \right) \ . \tag{9}$$

Since $C_0(a) = \bar{C} - \frac{1}{2}a$ is the lowest possible expected cost per previously treated patient, the rule calls for abstaining from administering Therapy 2 to previously treated patients except, when a is higher than a_b, for patients expected to respond best. The condition $a \leq a_b$ means that $E_b - C_0(a) \leq 0$, i.e. even the best responding patient is not expected to benefit from Therapy 2.

The solution for Y_b is:

$$Y_b^* = 1 - a \ . \tag{10}$$

It calls for administering Therapy 2 to all naïve patients.

The analysis is almost identical if Therapy 1 turns out g. Adapting the notation in obvious fashion, the optimum net expected value of administering Therapy 2 to X_g

previously treated and Y_g naïve patients is:

$$V(a|g) = \max_{X_g, Y_g} X_g \left(E_g - \bar{C} + \frac{1}{2}a \right) - \frac{1}{2}X_g^2 + Y_g \left(E_n - \bar{C} \right), \quad 0 \leq X_g \leq a;\ 0 \leq Y_g \leq 1 - a \tag{11}$$

The solution for X_g is:

$$X_g^* = \begin{cases} a \ , \ a \leq a_g \\ E_g - C_0(a) \ , \ a > a_g \end{cases} \tag{12}$$

where a_g is the critical value of a at which the solution for X_g shifts from a corner solution where all previously treated patients are administered Therapy 2, to an interior solution where some patients that are expected to respond poorly are not administered the therapy. If it exists, a_g is defined by the condition:

$$a_g = 2\left(E_g - \bar{C} \right) \ . \tag{13}$$

The solution for Y_g is:

$$Y_g^* = 1 - a \ . \tag{14}$$

It calls for administering Therapy 2 to all naïve patients as when Therapy 1 turns out b.

While naïve patients receive Therapy 2 whether Therapy 1 turns out g or b, the key difference between what happens in state g and what happens in state b occurs at the level of previously treated patients. If a corner solution occurs in state g, it involves treating all previously treated patients; while if a corner solution occurs in state b, it involves treating none of the previously treated patients. In case of interior solutions, in state b, being able to discriminate according to patient responsiveness enables the decision maker to select the patients expected to respond best and administer them Therapy 2. In contrast, in state g, the natural decision would be to treat all previously

treated patients; being able to discriminate according to patient responsiveness enables the decision maker to improve that rule by eliminating the worst responding patients from that group. It remains true, and can be shown that, whatever the level of patient discrimination made possible by prior learning through Therapy 1, more previously treated patients will be given Therapy 2 if Therapy 1 turns out g than if it turns out b.

Substituting (8) and (10) into (7) yields $V^*(a|b)$ the value function in state b; substituting (12) and (14) into (11) yields $V^*(a|g)$ the value function in state g. Weighting these two functions by their respective probabilities of $1-\gamma$ and γ and adding them up yields the second period value function, giving the net expected period 2 total welfare gain as a function of the proportion

of patients having received Therapy 1 in period 1:[11]

$$
\begin{aligned}
U(a) &= \gamma V^*(a|g) + (1-\gamma) V^*(a|b) \\
&= \begin{cases}
\gamma\left(E_n - \bar{C} + a\left(E_g - E_n\right)\right) + (1-\gamma)(1-a)\left(E_n - \bar{C}\right), & 0 \le a < a_g \\
\gamma\left(\frac{1}{2}\left(E_g - \bar{C} + \frac{1}{2}a\right)^2 + (1-a)\left(E_n - \bar{C}\right)\right) & \\
\qquad + (1-\gamma)(1-a)\left(E_n - \bar{C}\right), & a_g \le a < a_b \\
\gamma\left(\frac{1}{2}\left(E_g - \bar{C} + \frac{1}{2}a\right)^2 + (1-a)\left(E_n - \bar{C}\right)\right) & \\
\qquad + (1-\gamma)\left(\frac{1}{2}\left(E_b - \bar{C} + \frac{1}{2}a\right)^2 + (1-a)\left(E_n - \bar{C}\right)\right), & a_b \le a \le 1
\end{cases}
\end{aligned}
\tag{15}
$$

4. First period optimization

Period 1 treatment decision is based on the summation of period 1 payoffs and period 2 payoffs.

$$W(a) = u(a) + U(a) \tag{16}$$

[11] Function $U(a)$ is stated here for the case $a_g < a_b$, i.e. $[E_g + E_b]/2 \le \overline{C}$. The reader can adjust the formulation in the alternative case $a_g \ge a_b$.

where $U(a)$ is given by (15) and $u(a) = \varepsilon a - tc(a)$ gives the utility gains expected in period 1 for the proportion a of patients administered Therapy 1, net of the total expected cost of treatment in period 1, given by (4).[12] The total payoff function, an example of which is presented graphically in Figure 1, inherits from (15) its three-segment structure. Each segment is defined over one of the intervals delimited by 0, a_g, a_b, and 1, which correspond to the various possible combinations of period 2 treatments. The slope of W is a continuous function of a, a smoothness due to the fact that some optimum period 2 shares are interior over some ranges. For the parameters used to draw Figure 1, the function is concave, a property conferred to it by the assumed convexity of $tc(a)$;[13] it also reaches its maximum at an interior value of a. While other parameters may produce value functions that reach their maximum at $a = 0$ or $a = 1$, the configuration of Figure 1 corresponds to the observed stylized facts and best underlines the trade offs between resistance and therapeutic opportunity cost, and between acting now, learning actively, or learning passively, involved in decisions about HAART therapies. From both the medical and the economic point of views, the first-period treatment decision is then most difficult and interesting because these trade offs are of commensurate magnitudes, resulting in an interior choice for a. We will focus the analysis on this configuration, leaving to the reader the onus of making the minor adjustments required for alternative configurations that may better reflect other decision making situations.[14]

[12] This assumes that patients left untreated do not experience any utility change nor change in medical costs during period 1 as a result of the treatment of other patients.

[13] Concavity also results if one assumes decreasing returns to learning by treating.

[14] The following parameters were used: $\bar{C} = 1$; $E_b = 0.7$; $E_n = 1.1$; $E_g = 1.2$; they imply $a_g = 0.4$ from (13); and $a_b = .6$ from (9). Other parameters were: $\varepsilon = .5$, and $\gamma = .5$.

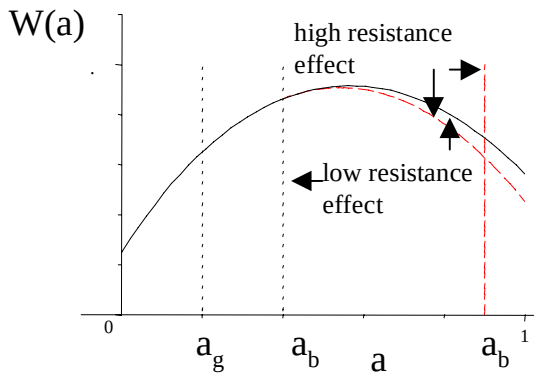

Figure 1: Net expected two period value function

When the solution is interior as in Figure 1, its properties can be analyzed from the first-order condition $\frac{\partial W}{\partial a} = 0$. This condition may be satisfied on any of the three a intervals on which $W(a)$ is defined. These intervals each corresponds to a particular therapeutic rule to be used in period 2:[15] the properties of the period 1 solution may differ accordingly. Setting $\frac{\partial W}{\partial a} = 0$, the first-order condition for an interior solution yields one possible value of a^* for each of the intervals. Since the intervals are not overlapping and cover all possible proportions, and since the W function is concave so that its maximum is unique, there is at most one of the three possible candidate values of a^* that also lies in the corresponding interval of definition of $W(a)$ (see (15) and (16)) : [16]

$$a^* = \begin{cases} \gamma\left(E_g - E_n\right) - (1-\gamma)\left(E_n - \bar{C}\right) + \varepsilon \text{ if } 0 \leq a^* \leq 2\left(E_g - \bar{C}\right) \\ \frac{2}{4-\gamma}\gamma\left(E_g - \bar{C}\right) - \frac{4}{4-\gamma}\left(E_n - \bar{C} - \varepsilon\right) \text{ if } 2\left(E_g - \bar{C}\right) \leq a^* < 2\left(\bar{C} - E_b\right) \\ \frac{2}{3}\left(\gamma\left(E_g - \bar{C}\right) - (1-\gamma)\left(\bar{C} - E_b\right)\right) - \frac{4}{3}\left(E_n - \bar{C} - \varepsilon\right) \text{ if } 2\left(\bar{C} - E_b\right) \leq a^* \leq 1 \end{cases} \quad (17)$$

[15] As described in the previous section, for a given parameter combination, the rule depends on a and consists in: either (line 1 of (17)) treating all previously treated and all naïve patients in state g while treating all naïve patients and no previously treated ones in state b; or (line 2 of (17)) treating all naive patients, and some previously treated ones in state g while treating all naïve patients and no previously treated ones in state b; or else (line 3 of (17)) treat all naïve patients and all but some previously treated ones in state g, while treating all naïve ones and none but some of the previously treated ones in state b.

[16] We continue with the case case $a_g < a_b$; as before it is immediate to adjust the formula for the alternative case case $a_g \geq a_b$.

For the parameter combination used in Figure 1,[17] the solution is $a^* = 0.57$: this value of a^*, obtained from the third line of (17), also lies in the interval $\left[2\left(\bar{C} - E_b\right),\ 1\right]$, i.e. $[.4,\ 1]$ as required.

Formula (17) indicates that the optimum proportion of patients submitted to Therapy 1 in general depends on ε the direct expected patient utility gain from that therapy in period 1, on the expected period 2 utility gains E_g or E_b that are made possible by the use of Therapy 1 in period 1, on the utility gain E_n made possible in period 2 by abstaining from Therapy 1 in period 1, and on the parameters c (normalized to one) and $\bar{C}$ that characterize the expected cost of applying Therapy 1 in period 1 and Therapy 2 in period 2. Since they are not observed *ex ante*, period 2 costs and utility gains are weighted according to the probability γ that Therapy 1 turns out g.

5. Discussion

5.1 The optimal decision

The optimal decision strategy can be interpreted as applying the following principle. In period 1 choose the proportion a of patients treated in the current period, but do not commit to any decision for period 2. Instead, postpone the period 2 decision, simply choosing the rule that will be used in period 2, based on the new information that will be acquired in the meantime, and on the initial choice of a.

The period 2 rule defines the proportions of patients from the group of previously treated patients, and the group of naïve patients, that will receive Therapy 2. The proportions differ according to the group considered, according to the exogenously gathered information (the state of nature), and according to the endogenously gathered information (learning by treating). Consequently, the period 1 decision on a affects the period

[17] See footnote 14. It is simple, although tedious, to spell out the conditions on the model parameters that ensure the maximizing value a^* to lie in each of the three possible intervals. For example, in order a^* to lie in $[a_b, 1]$, the parameters must satisfy $a_b \leq a^* \leq 1$ or $2\left(\bar{C} - E_b\right) \leq \frac{2}{3}\left(\gamma\left(E_g - \bar{C}\right) - (1-\gamma)\left(\bar{C} - E_b\right)\right) - \frac{4}{3}\left(E_n - \bar{C} - \varepsilon\right) \leq 1$. See the discussion below for an interpretation.

2 decision through two channels. First it determines the size of the two groups of patients existing at the beginning of period 2. Second it determines what is known about individual patients responsiveness to therapy among the group of previously treated patients; this in turn allows the clinician to discriminate between previously treated patients according to their susceptibility to treatment.

Under such circumstances, the period 1 choice of a is not based only on period 1 costs and benefits; it also seeks to promote future exposure to good outcomes and reduce exposure to bad outcomes. To illustrate by way of the solution corresponding to Figure 1 again, where a^* is given by the third line in (17), consider the optimal second-period contingent shares in that instance. In case Therapy 1 turns out to induce viral resistance (outcome b), the expected net-of-cost utility gain from Therapy 2 would be negative on average for previously treated patients: zero exposure to that bad outcome could be achieved by not administering any therapy; however, among these patients, the best responding ones would benefit from Therapy 2, provided the clinician had enough information to select them with sufficient precision. This can be done only if enough learning by treating has occured in period 1. The first period share of patients undergoing Therapy 1 is therefore set higher than it would be otherwise. As far as naïve patients are concerned, the net expected gain of administering Therapy 2 is positive on average for any a; while, *ex post*, some of them may turn out not to respond well, there is no way to select them and leave them untreated because no information has been gathered about them: maximum exposure is achieved by administering the therapy to all naive patients.

In case of a good outcome g for Therapy 1, the net expected gain of administering Therapy 2 to previously treated patients is positive on average: a high exposure can be achieved by administering the therapy to all such patients; however, *ex post*, some of them may still turn out to respond poorly. A further gain can be achieved if the least susceptible patients can be selected with enough accuracy and left untreated. Again, the decision maker acquires that ability by treating more patients in period 1 than she

would otherwise, thereby increasing exposure to the good outcome. No such possibility exists with respect to naïve patients: they all receive Therapy 2.

5.2 Viral resistance and therapeutic opportunity cost

As mentioned at the beginning of this paper, a major issue in HAART therapy is the development of resistance to treatments. The magnitude of the expected viral resistance effect is $E_n - E_b > 0$. On the other hand there is an opportunity cost of not administering a HAART therapy before the possible adverse resistance effect is known: the patient may deteriorate to the point of not being able to benefit from new therapeutic developments. The magnitude of the expected therapeutic opportunity cost is $E_g - E_n$. The dilemma faced by the clinician decision maker in period 1 results from the tension between these two magnitudes, and takes all its dramatic sense when the optimal value of a is interior: in such case, one would expect the optimal value of a to be sensitive to both the resistance effect and the therapeutic opportunity cost, making these scientific data crucial to the therapeutic decision.

This intuition turns out not to be necessarily true. More precisely, under parameter restrictions corresponding best to the example of HAART therapies, the magnitude of the resistance effect is irrelevant provided it is big enough. Figure 1 has been designed to illustrate this paradoxical result. The continuous curve illustrate a parameter combination where E_b, while it satisfies restrictions (1) and (6), is relatively high. This means, other things equal, that the expected viral resistance effect $E_n - E_b$ is relatively small. In that case, the optimal value a^* occurs to the right of a_b. This corresponds to a value of a within the third possible interval, $[a_b,\ 1]$, so that the solution is given by the third line of (17) : the formula indicates that a^* in that case depends on E_b; precisely, as intuition suggests, a^* is a decreasing function of the expected viral resistance effect. However, at higher levels of the resistance effect, the situation is different as illustrated by the dashed curve. Then the optimum occurs in the second interval $[a_g,\ a_b]$, both because the $W(a)$ curve has a different shape and because the value of a_b is now higher.

On that interval, as can be verified by inspecting the second line of (17), a^* does not depend on E_b. Moreover, the resistance effect may increase by any amount (i.e. E_b may diminish by any amount), there will be no effect on the optimum. In fact the higher the expected viral resistance effect, the shorter the interval $[a_b,\ 1]$ over which it is relevant; that interval may even disappear.

What is the economic and therapeutic intuition explaining this unexpected result? It resides with the decision rule to minimize exposure to bad outcomes and maximize exposure to good outcomes: when the resistance effect is high, the optimal rule is to choose zero exposure to the bad outcome. However zero exposure is compatible with some patients receiving Therapy 1 in period 1: it is achieved by not administering Therapy 2 to these patients in case Therapy 1 turns out b. However, if Therapy 1 turns out g, having a pool of previously treated patients allows a better exposure to the good outcome.

This result occurs in a framework that may be more of an example than a general case. However this example is based on key features of HAART therapies and on important stylized facts. In particular the proportions of patients receiving therapy in both periods are strictly interior.[18] Learning creates heterogeneity among patients. The model shows that the possible irrelevance of the resistance effect is compatible with differences in patient responsiveness, although these differences can be shown not to be necessary for the result to occur. They are simply one of the main justifications for observed differences in individual treatments.

Two features are important to the result. The first one is the presence of at least two distinct groups of patients in period 2, resulting in a mixture of corner solutions and interior solutions. In a model where all solutions were interior, the result would not hold, because there would not exist any situation where the decision maker could avoid exposure to the bad outcome completely. Even though, the intuition brought up by our

[18]There are several reasons why proportions of patients treated might be interior, some of which probably compatible with some relevance for the resistance effect. However, the simplest way to generate interior proportions of patients is to assume increasing expected per patient costs, and the most natural way to do so in the absence of patient specific information is to assume that expected per patient costs depend only on the number of patients receiving the therapy. Such a model, although less sophisticated, generates similar results.

result would help evaluate the situation: if the decision maker can reduce exposition to the bad outcome by choice of period 2 decision rule, then the magnitude of the bad outcome does not matter as much. In the present model, changes in expected viral resistance, when high enough, do not affect period 2 decision concerning previously treated patients: none is given Therapy 2 in state b. The second feature is the irreversibility of the actions taken by the decision maker, combined with uncertainty taking the form of new information arriving over time. Without this second feature, there would be no exposure to the bad outcome since the decision could be undone.

5.3 Patient homogeneity, learning, and real options

To a good clinician, each patient is unique. Yet, unless specific information is known about individual patients, they must be considered homogenous. Learning about individual health characteristics is a therapeutic act. Our model stresses this feature by assuming that no specific information is known at first about individuals in the wide group of HIV infected patients, but that learning about individuals occurs by treating. Consequently, at the beginning of period 2, there is a group of heterogenous individuals, consisting of patients who have received Therapy 1; and a group of homogenous individuals, consisting of naïve patients. In the first group, the clinician is able to select patients according to expected responsiveness to Therapy 2; in the second group, the clinician is unable to make such a distinction. Thus learning by treating produces patient differentiation, and patient differentiation can be used for a better allocation of resources devoted to Therapy 2. (see Moscaniri and Smith, 2001, for a view on experimentation)

Learning by treating is endogenous in that it depends on a. There is also exogenous learning in the decision problem addressed in this paper; it takes the form of information becoming available at the beginning of period 2 no matter the decision on a. This information on induced viral resistance results from the general R&D and experimentation carried out independently of, but observed by, the decision maker.

Exogenous learning over time is typical of real options. It is uncertainty on future viral resistance possibly induced by Therapy 1, combined with the irreversibility of any

action chosen in period 1, that gives rise to the particular decision rule described above: promote future exposure to good outcomes; avoid future exposure to bad ones. Yet the problem differs from most typical real option problems [19](e.g. Dixit and Pindyck, 1994) in at least three respects. First the focus on timing is limited here as there are only two periods. In a typical real-option problem, choosing an exercize date optimally is central to achieving the highest possible value for the option. The value of flexibility lies in the ability to make that choice. Here the timing flexibility consists in deciding what to do at two points in time; although limited, this flexibility generates an optimizing behavior which is typical of real options. The second difference is that both the state and the nature of the decision change over time in our HAART model, while, in a typical real option model, the state changes but the decision remains the same (e.g. invest or not). Indeed, in period 2 of our model, the clinician decision maker faces both a different state of nature (the effect of Therapy 1 on Therapy 2 has been revealed) and a new set of decisions (administer Therapy 2 or not, rather than Therapy 1). We think that this is an important feature of many decision problems, including health care decisions, that needs further analysis going beyond the current paper. The third difference is endogenous learning; endogenous learning is not incompatible with real options; it is simply not often introduced in that literature.

6. Conclusion and limitations

The additional gains in life expectancy associated with HAART combinations are not yet precisely known. Studies that attempted to model the impact of antiretroviral treatments have provided estimations in the range of 6 to 24 years for the increase in an individual's life expectancy from HIV infection to death for patients treated in North America (Blower, Gershengorn & Grant, 2000 ;Wood et al., 2000a).

However, uncertainties remain about the long-term efficacy of HAART therapies.

[19]There are few instances of real options being used in formulating health decisions. The ones that we are aware of have to do with investments in equipments or technology (Moretto & Levaggi, 2004; Palmer and Smith 2000). A related literature addresses food safety, biodiversity, genetically modified crops, and resistance to antibiotics or pesticides (Morel et al., 2003; Laxminarayan, 2003; Kassar and Lasserre, 2004; Salin, 2000).

For example, recent evidence has shown the persistence of viral replication even in successfully treated patients (Zhang et al., 1999; Finzi et al., 1999). In such context it is not surprising that complete consensus has not been reached world-wide, among clinicians and health authorities, about the best standards of practice for HAART delivery. In particular, clinical guidelines still differ between countries, and sometimes inside each country, about the eligibility criteria for HAART initiation. Empirical studies draw conflicting conclusions on the matter: some of them strongly question aggressive early use of HAART (Tebas et al., 2001), whereas others advocate very early initiation (as soon as HIV-infected patients have less 500 CD4 cells/microL) on cost-effectiveness grounds (Schackman, 2001). Interestingly, the former try to account for the resistance effect while the latter ignore it. Moreover, whatever the guidelines, surveys among prescribing physicians show a great variability of attitudes toward initiation of HAART treatment (Obadia et al., 1999 ; Reedjik et al., 1999 ; Kitahata, Van Rompaey & Shields, 2000 ; Landman et al., 2000).

In developing countries, where the vast majority of HIV infected people currently live, access to antiretroviral treatment was not considered a feasible technical and economic option until recently (Van Praag et al., 1997; Ainsworth and Teokul, 2000). Following the United Nations General Assembly Special Session on AIDS in 2001, a multi-lateral Global Fund to Fight AIDS, Tuberculosis and Malaria has been established at the beginning of 2002, and the goal of scaling up access to HAART in developing countries is increasingly shared by governments and international donor organisations. Between 1996 and 2000, expensive drug costs were the major barrier for diffusion of HAART in these countries. In the last three years, significant reductions in the prices of antiretroviral and other HIV-related drugs have been brought about in developing countries with the greatest need for access to HAART.

Our simple stochastic dynamic model of a sequential therapeutic choice with intermediary revelation and acquisition of information underlines the importance of expectations about effectiveness and costs in current and future therapies, as well as the importance of induced viral resistance and therapeutic opportunity cost. The fear that diffusion

of HAART may spread viral resistance tends to become the most powerful argument in favour of limiting or delaying access to antiretroviral treatment. At the empirical level, there is evidence from the Brazilian programme of universal coverage for HAART and from pilot experiments in African countries such as Senegal and Uganda that viral resistance and non-adherence are not a greater problem in cohorts of patients treated in developing countries when compared to data from developed countries (Tanuri et al., 2002; Silveira et al., 2002; Weidle et al., 2002; Laurent et al., 2002). Unilateral attitudes and arguments, such as the ones recommending to withhold or delay access to HAART in certain groups of patients or countries (Senak, 1997 ; Stewart, 1997 ; CDC, 1998) for fear of possible diffusion of drug-resistant HIV strains, express very questionable implicit trade-offs.

Moreover, a main conclusion of our model is that, when there is a significant risk of resistance due to therapeutic failure of initial existing treatments, differences in the estimation of this risk should not influence the optimal decision about the size of the HIV-infected population eligible for initiation of HAART. Paradoxically, it is in the case where expectations about resistance are rather optimistic (the phenomenon will be limited) that differences in estimations of this phenomenon may be a factor of variability in optimal treatment initiation. Because these conclusions are quite counter-intuitive, they may help clarify current inconsistencies between recommendations and practical behaviors of HIV/AIDS clinicians and public health experts on the one hand, and the expressed set of preferences and expectations of these same decision-makers, on the other hand (Gerbert et al., 2000).

A major limitation of our model is that we focus on the impact of the decision to initiate treatment on a population which is already HIV-infected. From a public health perspective, negative externalities associated with the diffusion of resistant HIV-strains into the rest of the population, are important factors (Geoffard, Philipson, 1996). Our model can be adapted to take these externalities into account.

REFERENCES

AINSWORTH, M. and TEOKUL, W. (2000) . Breaking the silence: setting realistic priorities for AIDS control in less developed countries. Lancet 356, 55-60.

British HIV association (BHIVA) Guidelines Coordinating Committee. (1997). Guidelines for antiretroviral treatment of HIV seropositive individuals. Lancet, 349, 1046-1092.

BLOWER, S.M., GERSHENGORN, H.B., GRANT, R.M. (2000). A tale of two futures : HIV and antiretroviral therapy in San Francisco. Science, 287, 650-654.

CARPENTER, C.C.J., COOPER, D.A., FISCHL, M.A. et al. (1997) Antiretroviral therapy for HIV infection in 1997. Updated recommendations of the International AIDS Society-USA panel. JAMA, 277, 1962-1969.

Center for Diseases Control (CDC). (1998). Report of the NIH Panel to define principles of therapy of HIV infection. Guidelines for the use of antiretroviral agents in HIV-infected adults and adolescents. MMWR , 47 (RR-5), 43-82.

DELFRAISSY, J.F. (Ed.). (1999). Guidelines for the use of antiretroviral therapies in HIV infection. Report to the Ministry of Health & Social Affairs. Flammarion Eds, Paris.

DESCAMPS, D., FLANDRE, P., CALVEZ, V. (2000). Mechanisms of virologic failure in previously untreated HIV-infected patients from a trial of induction-maintenance therapy. Trilège (Agence Nationale de Recherches sur le SIDA 072) Study Team. JAMA, 283, 205-11.

DEEKS, S.G., SMITH, M., HOLODNIY, M. & KAHN, J.O. (1997). HIV-1 protease inhibitors. A review for clinicians. JAMA, 277, 145-153.

DELFRAISSY J. F. (éd.) (2002) Prise en charge des personnes infectées par le VIH. Rapport 2002. Recommandation du groupe d'experts au Ministre de la Santé, de la Famille et des personnes handicapées. Médecine Sciences, Flammarion: Paris.

DIXIT, A. and PINDYCK, R. S. (1994). Investment under Uncertainty. Princeton University Press: Princeton.

FINZI D., BLANKSON J., SILLCIANO JD et al. (1999). Latent infection of CD4+T cells provides a mechanism for lifelong persistence of HIV-1 even in patients on effective combination therapy. Nat Med, 5, 512-517.

GALLANT, J.E. (2000). Strategies for long-term success in the treatment of HIV infection. JAMA, 283, 1329-1334.

GEOFFARD, P.Y., PHILIPSON, T. (1996). Rational epidemics and their public control. Intl Eco Rev, 37, 603-623.

GERBERT, B., BRONSTONE, A., CLANON, K., ABERCROMBIE, P., BANGSBERG, D. (2000). Combination antiretroviral therapy : health care providers confront emerging dilemmas. AIDS CARE, 12, 409-421.

HAMMER, S.M. et al. (2002) . Scaling up antiretroviral therapy in resource-limited settings: guidelines for a public health approach. World Health Organization, Geneva.

HECHT, F. M., GRANT, R. M., PETROPOULOS, C. J., et al. (1998). Sexual transmission of an HIV-1 variant resistant to multiple reverse-transcriptase and protease inhibitors. New England Journal of Medicine, 339, 307-11.

HOGG, R.S., HEATH, K.V., YIP, B. et al. (1998a). Improved survival among HIV-infected individuals following initiation of antiretroviral therapy. JAMA, 279, 450-454.

HOGG, R. S.; YIP, B.; CHAN, K. J.; WOOD, E.; CRAIB, K. J.; O'SHAUGHNESSY, M. V.; MONTANER, J. S. (2001). Rates of disease progression by baseline CD4 cell count and viral load after initiating triple-drug therapy, JAMA: The Journal Of The American Medical Association, 286(20), 2568-2577.

IDEMYOR, V. (2002). Continuing debate over HIV therapy initiation, HIV Clinical Trials, 3(2), 173-176.

KASSAR, I. and LASSERRE, P. (2004). Species Preservation and Biodiversity Value: A Real Options Approach, Journal of Environmental Economics and Management, 48(2), 857-79.

KITAHATA, M.M., VAN ROMPAEY, S.E., SHIELDS, A.W. (2000). Physician experience in the care of HIV-infected persons is related with earlier adoption of new antiretroviral therapy. J Acquir Immune Defic Syndr , 24, 106-114.

KYRIAKIDES, T. C.; GUARINO, P. (2001). Timing of antiretroviral treatment initiation, JAMA: The Journal Of The American Medical Association, 285(13), 1702-1703.

LANDMAN, R., MOATTI, J.P., PERRIN, V., HUARD, P. & The PAMPA Study Group. (2000). Variability of attitudes toward ealy initiation of HAAT for HIV infection; a study of French prescribing physicians. AIDS CARE, 12, in press.

LAURENT, C. et al. (2002). The Senegalese government's highly active antiretroviral therapy initiative: an 18-month follow-up study. AIDS 16, 1363-1370.

LAXMINARAYAN, R. (2003). Battling Resistance to Antibiotics and Pesticides *Resources for the Future,* Washington, D.C..

LE MOING, V.; CHÈNE, G., CARRIERI, M. P; ALIOUM, A.; BRUN-VÉZINET, F.; PIROTH, L.; CASSUTO, J. P.; MOATTI, J.-P.; RAFFI, F.; LEPORT et al. (2002). Predictors of virological rebound in HIV-1-infected patients initiating a protease inhibitor-containing regimen, AIDS (London, England), 16(1), 21-29.

LITTLE, S. J.; HOLTE, S.; ROUTY, J.-P.; DAAR, E. S.; MARKOWITZ, M.; COLLIER, A. C.; KOUP, R.A.; MELLORS, J. W.; CONNICK, E.; CONWAY ET AL. (2002). Antiretroviral-drug resistance among patients recently infected with HIV, The New England Journal Of Medicine, 347(6), 385-394.

MOREL, B., FARR0W, S., WU, F., et al. (2003). Pesticide resistance, the Precautionary Principle and the regulation of Bt corn: real and rational option approaches to decision making. In: LAXMINARAYAN, R. ed. Battling resistance to antibiotics and pesticides: an economic approach. Resources for the Future, Washington, 184-213.

MORETTO, M & LEVAGGI, R (2004). Investment in Hospital Care Technology under Different Purchasing Rules: A Real Option Approach, Working Papers 2004.75, Fondazione Eni Enrico Mattei.

MOSCARINI, G. and SMITH, L. (2001). The Optimal Level of Experimentation, Econometrica 69, 1629-44.

MURPHY, E. L.; COLLIER, A C.; KALISH, L A; ASSMANN, S F; PARA, M F; FLANIGAN, T P; KUMAR, P N; MINTZ, L; WALLACH, F R; NEMO ET AL (2001).

Highly active antiretroviral therapy decreases mortality and morbidity in patients with advanced HIV disease, Annals Of Internal Medicine, 135(1), 17-26.

OBADIA,Y., SOUVILLE,M., MORIN, M., MOATTI, J.P. (1999). French general practicioners'attitudes toward therapeutic advances in HIV care : results of a national survey. Intl J STD & AIDS, 10, 243-249.

PALMER, S. and SMITH, P.C. (2000). Incorporating option values into the economic evaluation of health care technologies, Journal of Health Economics, 19(5), 755-66.

PATERSON, D. L., SWINDELLS, S., MOHR, J. et al. (2000). Adherence to protease inhibitor therapy and outcomes in patients with HIV infection. Ann Intern Med, 133, 21-30.

PEDRAZA, M. A., DEL ROMERO, J., ROLDAN, F. (1999). Heterosexual transmission of HIV-1 is associated with high plasma viral load levels and a positive viral isolation in the infected partner, J Acquir Immune Defic Syndr, 21, 120-125.

PHILLIPS, A. N.; STASZEWSKI, S.; WEBER, R.; KIRK, O.; FRANCIOLI, P.; MILLER, V.; VERNAZZA, P.; LUNDGREN, J. D.; LEDERGERBER, B.; EUROSIDA STUDY GROUP (2001) HIV viral load response to antiretroviral therapy according to the baseline CD4 cell count and viral load, JAMA: The Journal Of The American Medical Association, 286(20), 2560-2567.

QUINN, T.C., WAWER, M.J., SEWANKAMBO, N., et al. (2000). Viral load and heterosexual transmission of human immunodeficiency virus type 1. Rakai Project Study Group, N Engl J Med, 342, 921-929.

Recommendations of the Panel on Clinical Practices for Treatment of HIV (2002). Guidelines for using antiretroviral agents among HIV-infected adults and adolescents. Recommendations And Reports: Morbidity And Mortality Weekly Report. Centers For Disease Control, 51(RR-7), 1-55.

REEDJIK, M., BINDELS, P.J.E., MOHRS, J., WIGERSMA, L. (1999). Changing attitudes towards antiretroviral treatment of HIV infection : a prospective study in a sample of Dutch general practicioners. AIDS CARE, 11, 141-145.

RISTIG, M B; ARENS, M Q; KENNEDY, M; POWDERLY, W; TEBAS, P. (2002) . Increasing prevalence of resistance mutations in antiretroviral-naïve individuals with established HIV-1 infection from 1996-2001 in St. Louis, HIV Clinical Trials, 3(2), 155-160.

RUBIO, R.; BERENGUER, J.; MIRO, J. M.; ANTELA, A.; IRIBARREN, J. A.; GONZALEZ, J.; GUERRA, L.; MORENO, S.; ARRIZABALAGA, J.; CLOTET ET AL. (2002) . Recommendations of the Spanish AIDS Study Group (GESIDA) and the National Aids Plan (PNS) for antiretroviral treatment in adult patients with human immunodeficiency virus infection in 2002, Enfermedades Infecciosas y Microbiologia Clinica, 20(6), 244-303.

SALIN, V. (2000). A Real Option Approach to Valuing Food Safety Risks, chapter 11 in Economics of HACCP: Costs and Benefits, Laurian J. Unnevehr, ed., Eagan Press: St. Paul, 225-240.

SCHACKMAN, B. R.; GOLDIE, S. J.; WEINSTEIN, M. C.; LOSINA, E.; ZHANG, H.; FREEDBERG, K. A. (2001) . Cost-effectiveness of earlier initiation of antiretroviral

therapy for uninsured HIV-infected adults, American Journal Of Public Health, 91(9), 1456-1463.

SENAK, M. (1997) Predicting antiviral compliance: physician's responsibilities vs. Patients' rights. J Intl Assn Phys in AIDS Care, 3, 45-8.

SILVEIRA, M. P. T.; DRASCHLER, M. DE L.; LEITE, J. C. DE C.; PINHEIRO, C. A. T.; DA SILVEIRA, V. (2002) Predictors of undetectable plasma viral load in HIV-positive adults receiving antiretroviral therapy in Southern Brazil, The Brazilian Journal Of Infectious Diseases: An Official Publication Of The Brazilian Society Of Infectious Diseases, 6(4), 164-171.

STEWART, G. (1997). Adherence to antiretroviral therapies. In Van PRAAG, E., FERNYAK, S. & KATZ, A.M. The implications of antiretroviral treatments. Informal consultation. Geneva, WHO/UNAIDS, 35-50.

TANURI, A.; CARIDEA, E.; DANTAS, M. C.; MORGADO, M. G.; MELLO, D. L. C.; BORGES, S.; TAVARES, M.; FERREIRA, S. B.; SANTORO-LOPES, G.; MARTINS et al. (2002) . Prevalence of mutations related to HIV-1 antiretroviral resistance in Brazilian patients failing HAART, Journal Of Clinical Virology, 25(1), 39-46.

TEBAS, P.; HENRY, K.; NEASE, R.; MURPHY, R.; PHAIR, J.; POWDERLY, W. G. (2001) . Timing of antiretroviral therapy: Use of Markov modeling and decision analysis to evaluate the long-term implications of therapy, AIDS (London, England), 15(5), 591-599.

VAN HEESWIJK, R. P.; VELDKAMP, A.; MULDER, J. W.; MEENHORST, P. L.; LANGE, J. M.; BEIJNEN, J. H.; HOETELMANS, R. M. (2001) . Combination of protease inhibitors for the treatment of HIV-1-infected patients: a review of pharmacokinetics and clinical experience, Antiviral Therapy, 6(4), 201-229.

VAN PRAAG, E., FERNYAK, S., KATZ, A.M. (1997) . Impact of antiretroviral treatments. Informal consultation. WHO/ASD/97.2, Geneva.

VOELKER, R. (1997) Debating dual AIDS guidelines. JAMA, 278, 613.

WAINBERG, M.A. & FRIEDLAND, G. (1998). Public health implications of antiretroviral therapy and HIV drug resistance. JAMA, 279, 1977-1983.

WEIDLE, P.J. et al. (2002). Assessment of a pilot antiretroviral drug therapy programme in Uganda: patients' response, survival, and drug resistance. Lancet 360, 34-40.

WOOD, E., BRAITSTEIN, P., MONTANER, J.S.G. et al. (2000). Extent to which low-level use of antiretroviral treatment could curb the AIDS epidemic in sub-Saharan Africa. Lancet, 355, 2095-2100.

YENI, P. G; HAMMER, S. M.; CARPENTER, C. C. J.; COOPER, D. A; FISCHL, M. A.; GATELL, J. M; GAZZARD, B. G; HIRSCH, M. S; JACOBSEN, D. M; KATZENSTEIN et al.. (2002) . Antiretroviral treatment for adult HIV infection in 2002: updated recommendations of the International AIDS Society-USA Panel, JAMA: The Journal Of The American Medical Association, 288(2), 222-235.

ZHANG L., RAMRATNAM B., TENNER-RAQ K et al. (1999). Quantifying residual HIV-1 replication in patients receiving combination antiretroviral therapy. N Engl J Med., ,340, 1605-1613.

Appendix 1: Proof of Lemma 1

We assume that the distribution of patients according to individual treatment cost (including differences in patient response to treatment) is known but that the location of any particular individual on the distribution is not observable unless further information is acquired. We assume that, by treating patients in period 1, the decision maker learns something about the cost of treating these patients in period 2: previously treated patients can be ranked according to expected individual cost of treatment and response, although not perfectly. Consequently, in period 2, low expected cost patients are treated first; high expected cost patients last. If a small proportion of previously treated patients are selected for Therapy 2 (A small), they will be low expected cost patients. At $A = 0$, the expected cost of treating the marginal patient in period 2 is lowest; it is noted $C(0,a) \equiv C_0(a)$; at $A = 1$, the marginal expected cost, noted $C(1,a) \equiv C_1(a)$, with $C_0(a) \leq C_1(a)$, is highest. If all previously treated patients are treated again in period 2, the average cost per patient is $\bar{C}$, and is not affected by learning; this is because learning does not affect the distribution of patients but only the ability to select best responding patients. However, if a lower proportion is treated, that is to say if selection is taking place, the average cost for that sub-group is lower than if the whole group was treated. Assume,[20] as in Lemma 1, that the expected individual cost of treating previously treated patients is uniformly distributed according to the number of patients treated in period 2, say between $C_0(a)$ and $C_1(a)$. Then the expected individual cost is linear in the number treated in period 2 and rises, between $C_0(a)$ when $A = 0$, and $C_1(a)$ when $A = 1$.

If $A < 1$, patients selected for treatment have a lower expected individual cost, while patients to be left untreated have a higher expected individual cost. The assumption of a uniform distribution implies that $C(A,a) = C_0(a) + A[C_1(a) - C_0(a)]$. As a result, the expected total cost of treating a proportion A of na previously treated patients is $TC(A,a) = \int_0^A naC(\alpha,a)\,d\alpha = naC_0(a)A + na\frac{1}{2}A^2[C_1(a) - C_0(a)]$. Thus the total period 2 expected cost/responsiveness is quadratic in A.

Since patients' ranking is more accurate, the higher the proportion a of patients treated in the first period, treating in period 1 creates a gap between $C_0(a)$ and $C_1(a)$: $C_1(a) - C_0(a) \geq 0$ and increasing in a. By assumption, learning affects the ability to rank individual patients, but not their actual cost of treatment; this implies that the expected per patient cost of treating all previously treated patients ($A = 1$) must be $\bar{C}$: $\frac{C_0(a)+C_1(a)}{2} = \bar{C}$. ■

Appendix 2: A sufficient condition for uniformly distributed expected costs per patient.

Suppose that individual cost of treatment are distributed uniformely on the interval $[0, n]$, where patient $z \in [0, n]$ has the individual cost of treatment for Therapy 2,

[20] In Appendix 2, we give an example that provides a theoretical foundation for that assumption.

$C(z) = z$. Suppose that there is one and only one patient z having individual cost $C(z) = z$.

Suppose that a proportion an of patients are given Therapy 1, which, as a by product, reveals their cost of treatment (or responsiveness) for T2.

Suppose that a proportion Aa of these patients are treated by Therapy 2.

Let us calculate the expected cost $C(A, a)$ of treating a proportion A of na previously treated patients.

Consider $n = 5$, a population of five patients $z \in \{0, 1, 2, 3, 4\}$ whose individual cost for Therapy T2 is $C(z) = z$. Assume that a sample of size 3 is selected for period 1 treatment, *i.e.* $a = 3/5$. There are $C_5^3 = 5!/(3!2!) = 10$ different possible samples, where individual costs are arranged by increasing order in each sample: $\{2, 3, 4\}$, $\{1, 2, 3\}$, $\{1, 2, 4\}$, $\{1, 3, 4\}$, $\{0, 1, 2\}$, $\{0, 1, 3\}$, $\{0, 1, 4\}$, $\{0, 2, 3\}$, $\{0, 2, 4\}$, and $\{0, 3, 4\}$.

The expected sample is, taking the expected value of each three sample term, $\{(5/10), (20/10), (35/10)\}$. Then, the lowest possible expected cost is $C_0(a) = 5/10$, and the highest possible expected cost is $C_1(a) = 35/10$.

Assume, as a proposition to be verified, that the expected cost per patient cost is $C(A, a) = (5/10) + A((35/10) - (5/10)) = (5/10) + A((30/10))$ and let A be defined as $A = k/(n-1) = k/2$, for $k \in \{0, 1, 2\}$.

Then, if $k = 1$, *i.e.* $A = 1/2$, $C(A, a) = 5/10 + (1/2)(30/10) = 20/10$. This is indeed the expected cost for the second entry, corresponding to $A = 1/2$, in the above 3-patient sample.

For $k = 2$, *i.e.*$A = 1$, $C(A, a) = 5/10 + 30/10 = 35/10$. This is the expected cost for the third entry in our 3-patient sample.

This construction can be generalized for any n for costs distributions $C(z) = z$, $z \in [0, n]$.