Competitive Analysis of Financial Games

R. El-Yaniv
University of Toronto

A. Fiat
Tel-Aviv University

R. Karp *
UC Berkeley
and
ICSI

G. Turpin
University of Toronto

TR-92-064

Abstract

In the *unidirectional conversion problem* an on-line player is given the task of converting dollars to yen over some period of time. Each day, a new exchange rate is announced, and the player must decide how many dollars to convert. His goal is to minimize the *competitive ratio*, defined as $\sup_E \frac{P_{OPT}(E)}{P_X(E)}$, where E ranges over exchange rate sequences, $P_{OPT}(E)$ is the number of yen obtained by an optimal off-line algorithm, and $P_X(E)$ is the number of yen obtained by the on-line algorithm X. We also consider a continuous version of the problem, in which the exchange rate varies over a continuous time interval. The on-line player's *a priori* information about the fluctuation of exchange rates distinguishes different variants of the problem. For three variants we show that a simple *threat-based strategy* is optimal for the on-line player and determine its competitive ratio. We also derive and analyze an optimal policy for the on-line player when he knows the probability distribution of the maximum value that the exchange rate will reach. Finally, we consider a *bidirectional conversion problem*, in which the player may trade dollars for yen or yen for dollars.

*Supported by NSF Grant No. CCR-9005448

1 Introduction

Many ongoing financial activities, such as currency exchange or mortgage financing, must be carried out in an *on-line* fashion, with no secure knowledge of future events. Faced with this lack of knowledge, players of these financial games often use models based on assumptions about the future distribution of relevant quantities such as exchange rates or mortgage rates, and aim for acceptable results *on the average*.

An alternate approach in such situations - and the one we explore here - is to use *competitive analysis* (first applied to on-line algorithms by Sleator and Tarjan in [ST85]). In this approach, the performance of an on-line strategy is measured against that of an optimal off-line strategy having full knowledge of future events. For some measure of cost or of profit, we try to minimize the worst-case ratio of on-line cost to optimal cost or of optimal profit to on-line profit; if this ratio is bounded for all event sequences, we deem the on-line strategy to be *competitive* and call the supremum of this ratio the *competitive ratio* of the strategy. One advantage of this performance measure over the traditional average-case measure is that for most non-trivial financial activities it is extremely difficult to come up with an accurate probabilistic model.

In devising a competitive strategy, the on-line player will not be able to escape the necessity of making *some* assumptions, or having *some* knowledge about future events, but these need not be probabilistic in nature. For example, instead of knowledge about the distribution of future exchange rates, an on-line strategy might be based only on knowledge of the bounds on possible exchange rates over the period in question, and should work well no matter how erratically (or unfortunately) the rates vary from day to day.

Despite the obvious significance and appeal of investment problems, we have found only two previous references concerned with the competitive analysis of on-line financial decision-making. The paper [Cover91] deals with the problem of portfolio selection for investment in the stock market. Here, the (on-line) player must determine the proportion of his wealth invested in each of the stocks in the market. Cover presents a simple on-line strategy that (based on the market history) dynamically changes the distribution of its current wealth among the stocks. It is shown that this on-line strategy performs asymptotically as well as any static off-line portfolio selection strategy. Such a strategy fixes an optimal investment portfolio at the start of the game (based on a full knowledge of future events), but it is not allowed to change it thereafter. Note that the on-line portfolio selection strategy is very robust; it does not assume anything about the behavior of future events. The paper [Ragh91] analyzes the competitive performance of an on-line investment algorithm against a "statistical adversary," whose request sequence is required to satisfy certain distributional requirements. Thus, both of these papers restrict the adversary: in the first case, by limiting him to a static policy, and in the second case by restricting the request sequence.

We suspect that other potential investigators have been deterred by the belief that, in problems such as stock market investment, knowledge of the future is such an overwhelming advantage that it would be impossible for an on-line algorithm to achieve a bounded competitive ratio. Nevertheless, we shall show that, if suitable constraints are placed on the future evolution of prices or exchange rates, then finite, and indeed rather small, competitive ratios can be achieved in fairly realistic situations.

Our main example is one-way currency conversion, which deals with converting assets from one form to another, for example, dollars to yen. For this problem, we show that a surprisingly small competitive ratio can be achieved under very moderate assumptions about the on-line player's knowledge concerning the exchange rates: only upper and lower bounds on possible exchange rates need be known. We investigate different versions of this problem, given by varying the on-line player's knowledge. We conclude with some preliminary results for two-way currency conversion, where the on-line player converts currency back and forth.

2 The Unidirectional Conversion Problem

2.1 Problem statement

In the *unidirectional conversion problem* an on-line player is given the job of converting dollars to yen over some period of time. This is unidirectional conversion since the on-line player may not convert yen already purchased back to dollars. At any given time there is an *exchange rate* giving the number of yen that can be purchased for one dollar. The problem has a *continuous version*, in which the exchange rate fluctuates during some continuous period of time and the on-line player may trade continuously, and a *discrete version*, in which a new exchange rate is announced on the morning of each trading day, and remains fixed throughout the day. At the end of the trading period the trader *must* convert the remaining dollars at the current rate. The problem is designated as *on-line* since the player must determine his transactions without knowing what the future rates will be.

Considering the discrete version first, let $E = e_1, e_2, \ldots$ be a sequence of exchange rates. Let $P_X(E)$ denote the number of yen obtained by an on-line conversion strategy X that starts with D_0 dollars (without loss of generality $D_0 = 1$) and converts the dollars to yen in accordance with E. Let OPT denote the optimal off-line strategy (i.e., the one that waits and converts all its dollars to yen when the highest exchange rate is reached). For any on-line algorithm X, the *competitive ratio* is defined as $\sup_E \frac{P_{OPT}(E)}{P_X(E)}$. A small competitive ratio implies that X can do well in comparison with OPT. The set-up for the continuous case is similar, except that the exchange rate sequence E is replaced by an exchange rate function $E(t)$ defined over some time interval $[0, T]$, and trading can take place continuously during this interval.

2.2 Assumptions about the fluctuation of the exchange rate

The on-line player's *a priori* information about the exchange rate sequence or exchange rate function defines particular variants of the game. Let

- M = upper bound on possible exchange rates,
- m = lower bound on possible exchange rates,
- a = initial exchange rate,
- $\Phi = M/m$ (call this the *fluctuation ratio*),
- n = the number of trading days (in the discrete case).

2.3 Results

In this paper we derive optimal on-line algorithms for three variants of the unidirectional conversion problem.
Variant 1 The continuous version, with M, m and a known to the on-line player;
Variant 2 The discrete version, with M, m and n known to the on-line player;
Variant 3 The discrete version, in which the on-line player knows n and Φ, but does not know either M or m.

2.4 Optimal unidirectional conversion strategy

In this section we describe a general conversion strategy that yields optimal algorithms for the three problem variants. The strategy consists of three rules. Execution of the strategy requires knowledge of r, the competitive ratio that the algorithm is trying to attain. Later we will show how the smallest attainable competitive ratio is calculated.
Rule (i) At the end of the game (on the last day in the discrete case, or at the last instant in the continuous case), spend all the remaining dollars;
Rule (ii) Except at the end of the game, purchase only when the current rate is the highest seen so far;
Rule(iii) Whenever the exchange rate reaches a new maximum convert just enough to ensure that a competitive ratio of r would be obtained if an adversary dropped the exchange rate to m and kept it there throughout the rest of the game.

For the three variants that we are considering, these rules completely determine the on-line algorithm. We refer to this algorithm as the *threat-based algorithm*, since it converts dollars to yen only when forced to by the threat that the exchange rate will drop permanently to m.

2.5 The continuous case, with m, M and a known

In this case we can formulate the threat-based algorithm in terms of two functions, $D(x)$ and $Y(x)$, giving the number of dollars and yen, respectively, that the on-line player possesses just after the exchange rate has assumed the value x for the first time. Assume that the on-line player is trying to attain the competitive ratio r. We may assume that the exchange rate is monotone increasing, since both the optimal off-line algorithm and the threat-based algorithm conduct transactions only when the exchange rate reaches a new maximum. We may also assume that the exchange rate is a continuous function, since discontinuous jumps merely enable the threat-based algorithm to exchange its dollars at more favorable rates. Suppose the exchange rate increases continuously from a to M. Then Rules (i),(ii) and (iii) imply that the functions $D(\cdot)$ and $Y(\cdot)$ satisfy the following conditions:
Condition (i) $D(x) = 1$ and $Y(x) = 0$ when $x \leq rm$, since Rule (iii) does not require any exchanges until the rate reaches rm.
Condition (ii) $mD(x) + Y(x) = x/r$ for all $x \in [rm, M]$; here x is the number of yen the off-line player realizes if the maximum exchange rate is x, and $mD(x) + Y(x)$ is the amount the threat-based algorithm would realize if the rate dropped from x to m and its $D(x)$ remaining dollars were converted at rate m.
Condition (iii) For all $x \in [rm, M]$, $xD'(x) = -Y'(x)$ since each dollar exchanged at rate x yields x yen.

These conditions determine the functions $D(\cdot)$ and $Y(\cdot)$ uniquely, and thus the policy achieves its target competitive ratio r if and only if $D(M) \geq 0$; i.e., if and only if the above conditions can be satisfied without spending more than the initially available number of dollars, even when the adversary increases the exchange rate continuously from a to M.

We shall calculate the conditions under which the competitive ratio r is achievable.

Case 1: $a \leq rm$. In this case $D(x) = 1$ and $Y(x) = 0$ for $x \in [a, rm]$.

For $x \in [rm, M]$ we may differentiate the equation $mD(x) + Y(x) = x/r$ to obtain $mD'(x) + Y'(x) = 1/r$. Combining this with the equation $xD'(x) = -Y'(x)$ we obtain $D'(x) = -\frac{1}{r(x-m)}$. Thus $D(M) = D(rm) + \int_{rm}^{M} D'(x)dx = 1 - \int_{rm}^{M} \frac{1}{r(x-m)}dx = 1 - \frac{1}{r}\ln\frac{M-m}{rm-m}$. It follows that the competitive ratio r is achievable if and only if $\frac{1}{r}\ln\frac{M-m}{rm-m} \leq 1$. Of particular interest is the case where the initial exchange rate is unknown to the on-line algorithm, so that the pessimistic assumption $a = m$ must be made. In this case, the optimal attainable competitive ratio $r*$ is the unique root of the equation $r = \ln \frac{M-m}{rm-m}$.

Case 2: $a > rm$. In this case $mD(a) + Y(a) = a/r$ and $Y(a) = a(1 - D(a))$. This gives $D(a) = \frac{a-a/r}{a-m}$ and $D(M) = \frac{a-a/r}{a-m} - \int_a^M \frac{1}{r(x-m)}dx = \frac{a-a/r}{a-m} - \frac{1}{r}\ln\frac{M-m}{a-m}$. The competitive ratio r is attainable if and only if this quantity is nonnegative.

From these results it is easy to determine the optimal attainable competitive ratio. The first step is to determine whether the competitive ratio a/m is attainable; this will be the case if and only if $1 - \frac{m}{a}\ln\frac{M-m}{a-m} \geq 0$. If this rate is not attainable then we are in Case 1 and the optimal competitive ratio is $r*$. If the rate is attainable then we are in Case 2 and the optimal competitive ratio is the unique root of $\frac{a-a/r}{a-m} - \frac{1}{r}\ln\frac{M-m}{a-m} = 0$.

We omit the proof that the threat-based algorithm achieves the best competitive ratio of any on-line algorithm. A similar theorem will be proved below, in connection with Variant 2.

Some numerical examples will indicate the slow growth of the optimal competitive ratio as a function of M/m. We consider the case $a = m$, in which the optimal competitive ratio is $r*$; larger values of a serve to decrease the competitive ratio. When $M/m = 2$, $r* = 1.28$ (rounded to two decimal places); when $M/m = 4$, $r* = 1.60$; when $M/m = 8$, $r* = 1.97$; when $M/m = 16$, $r* = 2.38$.

2.6 Variant 2: Discrete case with n, M, m known

This case is similar to Variant 1 (with $a = m$), but with the following added complication. In analyzing Variant 1 it sufficed to consider an adversary that increased the exchange rate continuously up to its maximum value and then dropped it permanently to m. Here we have to consider a more complex class of adversaries, corresponding to all the different choices of the successive exchange rate maxima $a_1 < a_2 < \ldots < a_k$ for all $k \leq n$.

Without loss of generality, assume that the lower bound m on the exchange rates is 1. The upper bound M then equals the fluctuation ratio Φ. Let Y_i and D_i denote the number of yen and dollars held by the threat-based algorithm after the ith purchase (i.e. after the purchase at rate a_i). Then, by Rule (iii), $D_i + Y_i = a_i/r$. Let s_i be the amount spent on the ith purchase. Then $s_i = D_{i-1} - D_i$ and $a_i s_i = Y_i - Y_{i-1}$. Subtracting the equation $D_{i-1} + Y_{i-1} = a_{i-1}/r$ from the equation $D_i + Y_i = a_i/r$ and applying the definition of s_i we obtain: for $i > 1$, $s_i = \frac{1}{r}\frac{a_i - a_{i-1}}{a_i - 1}$. Similarly, $s_1 = \frac{1}{r}\frac{a_1 - r}{a_1 - 1}$. The competitive ratio r is achieved against the sequence of exchange rate maxima $a_1 < a_2 < \ldots < a_k$ only if $\sum_{i=1}^k s_i \leq 1$, and the best competitive ratio that the threat-based algorithm can attain against this sequence is the unique r for which $\sum_{i=1}^k s_i = 1$. A brief computation yields that this r is given by

$$r = 1 + \frac{a_1 - 1}{a_1}\sum_{i=2}^{k}\frac{a_i - a_{i-1}}{a_i - 1} \tag{1}$$

Thus, to determine the minimum competitive ratio attainable by the threat-based algorithm in an n-day game we maximize r in (1) over all choices of $k \leq n$ and $a_1, \ldots, a_k$ such that $1 \leq a_1$ and $a_k \leq M$.

By straightforward calculus one can show that, for a fixed value of a_1, the maximum is achieved when $k = n$, $a_n = M$ and all the ratios $\frac{a_i - 1}{a_{i-1} - 1}$, $i = 2, 3, \ldots, n$, are equal. This leads to the expression $r = 1 + \frac{a_1 - 1}{a_1}\cdot(n-1)\cdot(1 - (\frac{a_1 - 1}{M-1})^{\frac{1}{n-1}})$. Maximizing this expression as a function of a_1 we find after some manipulation that, for $n \geq 2$, r_n, the maximum value of r, is given by the implicit equation

$$r_n = n \cdot (1 - (\frac{r_n - 1}{M - 1})^{\frac{1}{n}}) \tag{2}$$

The quantity r_n is the competitive ratio of the threat-based algorithm. We shall show in the next section that no on-line algorithm can achieve a smaller competitive ratio.

Against the optimal choice of the a_i it turns out that $s_1 = s_2 = \ldots = s_n = \frac{1}{n}$. Thus, against the optimal adversary, the threat-based algorithm obeys the conventional wisdom of investment advisers by employing a *dollar-cost averaging* strategy, in which an equal number of dollars are invested each day.

It is easy to see that r_n increases with n. If the on-line player does not know the number of days n in advance, the competitive ratio should be taken to be $r_\infty \stackrel{\text{def}}{=} \lim_{n\to\infty} r_n$. The quantity r_∞ satisfies the following implicit equation.

$$r_\infty = \ln\frac{M-1}{r_\infty - 1}. \tag{3}$$

Not surprisingly, r_∞ is equal to the quantity $r*$ that arose earlier in connection with the continuous case.

2.7 Lower Bound

The threat-based algorithm is optimal for variant 2. Let S be any deterministic conversion strategy. Using a simple adversary argument, we show that, for any n, S cannot achieve a ratio smaller than r_n (as in (2)). Let the rates be the worst-case sequence $a_1, a_2, \ldots, a_n$, against which the threat-based algorithm achieves the competitive ratio r_n by spending $1/n$ dollars each day.

(a) If the first day i that S does not spend the same as our strategy it spends too little, then the adversary halts the game, forcing conversion of the remaining dollars at the lowest possible rate. The adversary acquires a_i yen and (see Rule (iii)) S acquires less than a_i/r_n yen.

(b) If the first time that S does not spend the same as our strategy it spends too much, the adversary continues the game but ends it on the first day j after that where the total T_j spent by S so far is less than or equal to our j-day total t_j. We can devise a strategy S' that spends the same total T_j by the end of the jth day, that acquires more yen than S, and still cannot achieve a ratio better than r_n. S' spends the same as our strategy on each of the first $j-1$ days (t_{j-1} in all) and $T_j - t_{j-1}$ (which is $\leq s_j$) on the jth day, and afterwards continues like S.

The generalization of this lower bound to randomized on-line conversion strategies is straightforward. Given the distributions that are used by the randomized strategy, an oblivious adversary has enough information to calculate the expected amounts that the randomized strategy will spend on each trading day.

3 Playing Against a Probability Distribution

Let y be the maximum value that the exchange rate will assume. In this section we consider the case where the on-line player knows the probability distribution of y. Let F be the cumulative distribution function for y; i.e., $F(x) = Pr[y \leq x]$. Let $F^c(x) = 1 - F(x)$. As the exchange rate varies the on-line player's estimate of y also varies; however, we assume that, if the maximum rate observed so far is s, the conditional distribution of y given the history of the exchange rate is obtained simply by excluding all values of y less than s, and maintaining the relative probabilities of values greater than or equal to s; i.e., the conditional distribution has the cumulative distribution function F_s given by: for $x < s$, $F_s(x) = 0$; for $x \geq s$, $F_s(x) = \frac{F(x)-F(s)}{1-F(s)}$.

We assume that the exchange rate varies continuously over time; discontinuous changes only work in the on-line player's favor by providing him with extra information. Let $x*$ be the value of x at which the quantity $mF(x) + xF^c(x)$ is maximized. Then the on-line player can maximize his expected return by exchanging all his dollars at rate $x*$ if the exchange rate ever reaches $x*$, and otherwise exchanging all his dollars at the rate m. His optimal expected return is thus $mF(x*) + x*F^c(x*)$. Define $r(F)$, the *return ratio against* F, as $\frac{E[y]}{mF(x*)+x*F^c(x*)}$.

Theorem 1 *Let $r(m, M, a)$ be the optimal competitive ratio attainable in the continous case, when the parameter values m, M and a are given. Then*

$$r(m, M, a) = \max_F r(F)$$

where F ranges over all probability distributions over the interval $[a, M]$.

The theorem states that the smallest competitive ratio that the on-line player can achieve is equal to the largest return ratio that the adversary can force by choosing a probability distribution for y.

Proof: Consider the two-person zero sum game where the set of pure strategies for each player is the interval $[a, M]$, and the payoff $f(x, y)$ to the first player if the first player chooses x and the second player chooses y is given by: if $x \leq y$ then $f(x, y) = y/x$ else $f(x, y) = y/m$. The interpretation is as follows. The first player is the on-line player and the second player is the adversary. The pure strategy x for the on-line player represents the decision to exchange his dollars as soon as the exchange rate reaches x, and to exchange at the rate m if the exchange rate never reaches x. The pure strategy y for the adversary corresponds to choosing y as the maximum value of the exchange rate. The payoff $f(x, y)$ is the ratio between the return to the adversary and the return to the on-line player when the pair (x, y) of pure strategies is played. A mixed strategy for the on-line player is a probability distribution over the values of x, giving the rate at which he will exchange dollars as the maximum observed value of the exchange rate increases continuously; if G is the cumulative distribution function for this probability distribution, then $G(x)$ represents the fraction of his dollars that the on-line player will exchange at exchange rates less than or equal to x. A mixed strategy for the adversary is a probability distribution for y, the maximum exchange rate; such a distribution can be specified by its cumulative distribution function F.

The min-max theorem of game theory implies that

$$\max_G \min_y E_G[f(x, y)] = \min_F \max_x E_F[f(x, y)] \tag{4}$$

where $E_G[f(x,y)]$ is the expected payoff when the adversary chooses y and the on-line player plays the mixed strategy G, and $E_F[f(x,y)]$ is the expected payoff when the on-line player chooses x and the adversary plays the mixed strategy F. The common value of the two sides of equation (1) is called the *value* of the game.

Since the left-hand-side of (1) is equal to $r(m,M,a)$ and the right-hand side is equal to $\max_F r(F)$, the proof is complete. □

Let us explicitly compute an optimal probability distribution F for the adversary in the case where $m = a$. Consider the problem of maximizing $E[y]$ subject to the constraint that $mF(x*) + x * F^c(x*) \leq z$, where z is a parameter to be specified later. The constraint is equivalent to the following condition: for all x, $mF(x)+xF^c(x) \leq z$, or, equivalently, $F^c(x) \leq \frac{z-m}{x-m}$, for $z \leq x \leq M$; also, by definition, $F^c(x) \leq 1$ for all x and $F^c(M) = 0$. Since $E(y) = m + \int_m^M F^c(x)dx$ the constraint requires that $E[y] \leq z + \int_z^M \frac{z-m}{x-m}dx = z + (z-m)\ln\frac{M-m}{z-m}$. By calculus, the ratio of $z + (z-m)\ln\frac{M-m}{z-m}$ to z is maximized when $z = m\ln\frac{M-m}{z-m}$. Let r denote $\max_F r(F)$. Then r is determined by evaluating the ratio of $z+(z-m)\ln\frac{M-m}{z-m}$ to z when $z = m\ln\frac{M-m}{z-m}$. A brief calculation shows that $r = \ln\frac{M-m}{rm-m}$. From Section 2.5 we see that this is the same implicit equation that defines the optimal competitive ratio $r* = r(m,M,m)$. This confirms that $r* = \max_F r(F)$, as Theorem 1 predicts. The probability distribution for the adversary that maximizes the return ratio is concentrated on the interval $[rm, M]$. On the half-open interval $[rm, M)$ the distribution has a cumulative distribution function given by the equation $F^c(x) = \frac{rm-m}{x-m}$. The distribution function is discontinuous at M, where the mass $\frac{rm-m}{M-m}$ is concentrated.

3.1 Some practical issues

3.1.1 Improved strategy against a clumsy adversary

The simple conversion strategy above is overly pessimistic since it fixes r at the time of the first purchase based on the assumption of a worst-case sequence of rates and does not change it thereafter. However, whenever the exchange rates deviate from the worst-case sequence, we can strictly improve the off-line to on-line ratio by recalculating the achievable r. At the start of each trading day, the on-line player has some number D of dollars and some number Y of yen. The player knows the number n' of days remaining, and is given an exchange rate a. The player acts as if the current day were the first trading day of an n'-day trading period in which the adversary starts with one dollar and the player with D dollars and Y yen. By arguments similar to those used above, an expression for r is determined and maximized over the *remaining* rates. The amount s to spend is given by

$$s = \frac{a - r\cdot(Y+D)}{r\cdot(a-1)}$$

where

$$r = \frac{a + (a-1)\cdot(n'-1)\cdot\left(1 - \left(\frac{a-1}{M-1}\right)^{\frac{1}{n'-1}}\right)}{a\cdot D + Y}.$$

3.1.2 Errors in assumptions about the exchange rate sequence

What happens if the on-line player's assumptions about the rate sequence turn out to be incorrect? Consider for example, the case in which the upper bound on possible exchange rates is underestimated. Say $M' = cM$ is the true upper bound where $c > 1$. The on-line player will play with $r = r_n$ as defined in (2). In the worst possible rate sequence, $a_{n-1} = M$ and $a_n = M'$. The on-line strategy will spend all the available dollars before the last day is reached and will acquire at least M/r_n yen. On the same sequence, OPT acquires M' yen. The attainable ratio is thus

$$M'/(M/r_n) = cr_n$$

or c times the ratio the player was aiming for.

3.2 Variant 3: only fluctuation ratio is known

In this case the player knows only Φ but not the actual bounds on possible exchange rates. A simple observation is that the estimate of the lowest possible rate changes from day to day and is given by a_i/Φ on the ith day.

The best competitive ratio for an n-day game that can be attained by thus revising the estimate of the lowest possible rate is $\Phi(1 - \frac{(\Phi-1)^n}{(\Phi^{\frac{n}{n-1}}-1)^{n-1}})$. It can be shown that as n approaches infinity this number approaches (monotonically from below) $\Phi - \frac{\Phi-1}{\Phi^{\frac{1}{\Phi-1}}} = \Theta(\ln\Phi)$.

4 The Bidirectional Conversion Problem

4.1 Problem statement

The on-line player starts with some D_0 dollars (without loss of generality we may take $D_0 = 1$) and converts back and forth between dollars and yen according to a sequence of exchange rates $E = e_1, e_2, \ldots$ which is revealed on-line. These rates (yen/dollar) must remain in the interval $[m, M]$, but otherwise may rise or fall arbitrarily. When the game ends, the money is all converted to one of the currencies at the present rate. The ratio of the amount accumulated by the optimal off-line algorithm on this sequence of exchange rates to the amount accumulated by the on-line player is the same, no matter which currency they convert to when the game ends. As in the unidirectional conversion problem, the on-line player's goal is to achieve a small *competitive ratio*.

If the number of local maxima and minima in the exchange rate sequence is unbounded, there is no competitive strategy (see lower bound argument in section 4.3). Assume there are k such maxima and minima. We show a lower bound of $(r_\infty)^{k/2}$ and an upper bound of $(r_\infty)^k$ where r_∞ is defined in (3).

4.2 Upper bound

The sequence E of exchange rates will consist of $k/2$ upward runs and $k/2$ downward runs. The optimal strategy for the off-line player is to convert all his dollars to yen at the end of each upward run, and all his yen to dollars at the end of each downward run. We describe how the on-line player can achieve a competitive ratio of r_∞^k.Suppose the first upward run consists of $e_1 \le e_2 \le \ldots \le e_i$, with $e_{i+1} < e_i$. During the first i days the on-line player plays according to our unidirectional conversion strategy, choosing r equal to r_∞. After day i he has D_i dollars and Y_i yen, where $mD_i + Y_i = \frac{e_i}{r_\infty}$. On day $i+1$ he converts all his dollars to yen. Since $e_{i+1} \ge m$ he then has at least $\frac{e_{i+1}}{r_\infty}$ yen at the beginning of the first downward run. He then proceeds similarly during the downward run beginning at day $i+1$, converting yen to dollars during the decreasing run and exchanging any remaining yen on the first day of the next increasing run. Thus, two transactions occur on the first day of any decreasing run: first, the exchange of all dollars to yen, and then the first step of the threat-based strategy with competitive ratio r_∞ for the unidirectional conversion of yen to dollars during the downward run; these conceptually distinct transactions can, of course, be combined into a single transaction. Similarly, on the first day of any increasing run after the first one, two transactions may occur. In each run the ratio between the off-line player's capital and the on-line player's capital increases by at most the factor r_∞, and thus the on-line player achieves the competitive ratio r_∞^k.

This strategy is not optimal. On any upward (downward) run the on-line player can take advantage of the knowledge that, to attain a ratio of r on the following downward (upward) run, the optimal adversary must begin that run with a particular rate. We are currently attempting to use this to get a better upper bound.

4.3 Lower bound

Let k be as described in the last section. Assume that the on-line player knows only m and M (w.l.o.g. $m = 1$ as before). For any n, it is possible for an adversary to force a competitive ratio of $(r_n)^{k/2}$ against any on-line strategy.

Our argument here follows from the lower bound proof for the unidirectional case. Suppose each player starts with one dollar. Then, regardless of what strategy the on-line player is following, the adversary can construct a sequence of n exchange rates consisting of an upward run followed by an immediate drop of the exchange rate to m, such that, at the end of the sequence, his total holdings in yen and dollars (evaluated at the exchange rate m) will exceed that of the on-line player by at least the factor r_n. Moreover, when the exchange rate drops to m, the adversary will convert all his yen to dollars, and the on-line player can do no better than to follow suit, since he will never have a more favorable conversion rate. Thus, in one upward and one downward run, the ratio of adversary currency to on-line currency can be made to increase by a factor of r_n. This yields a factor $(r_n)^{k/2}$ for the entire game. As n increases, this lower bound approaches $(r_\infty)^{k/2}$.

5 Future work

A striking feature of the unidirectional conversion problem is the conceptual simplicity of the optimal strategy. To attain a given competitive ratio r the on-line player simply defends himself against the threat of dropping the exchange rate permanently to m.

We have identified other problems in which a threat-based strategy is optimal. A typical example is the following problem of *trading on option*. A money trader starts with D_0 dollars and receives a finite sequence of option offers. Each option offer consists of a pair (e, p), where e is an *exchange rate* and $p > 0$ is an *option price*. For any $s \ge 0$ this offer enables the trader to pay ps dollars for the right to later trade s dollars for yen at the exchange rate e. The trader knows in advance that, even without purchasing options, he will always be able to exchange dollars for yen at the rate m, and that the maximum exchange rate he will ever be offered is M. As each offer (e, p) is presented, the trader chooses a corresponding value s and his stock of dollars is depleted by ps. At the end of the sequence of offers the trader converts his remaining dollars to yen by exercising some of the options he has purchased and by exchanging any remaining dollars at the rate m. The optimal off-line strategy is clear; the off-line player will choose that option (e, p) for which $\frac{e}{p+1}$ is greatest and (provided that $\frac{e}{p+1} > m$) will spend $\frac{pD_0}{p+1}$ dollars for the right to exchange his remaining $\frac{D_0}{p+1}$ dollars for yen at the exchange rate e. Under this optimal strategy the off-line player receives $\frac{eD_0}{p+1}$ yen. It can be shown that, if the trader can achieve a competitive ratio of r for this problem, then he can do so using a threat-based strategy of the following form: whenever an offer (e, p) is presented, choose the minimum value of s that will ensure that the trader can obtain $\frac{1}{r}\frac{eD_0}{p+1}$ even under the threat that no further option offers will be made.

It would be of interest to identify other financial games in which a threat-based strategy is optimal. Several problems related to inventory management, equipment replacement and investment are candidates. For example, consider the following *mortgage problem*. The on-line player initially holds a mortgage with interest rate e_0. On the ith day, where i ranges from 1 to n, the bank announces a new interest rate e_i. It is given that, for all i, $m \leq e_i \leq M$. Each day the on-line player must decide either to retain his current mortgage or to refinance at the new rate e_i. Each time he refinances, he incurs a transaction cost T. In addition, for each day in which he holds a mortgage at rate e, he incurs an *interest charge* of e. The player's total cost is the sum of his transaction costs and interest charges. We conjecture that the optimal on-line strategy is a threat-based strategy under which the on-line player refinances if and only if, should he fail to do so, the adversary could force a competitive ratio greater than the target value r by increasing the mortgage rate permanently to M, thus denying the on-line algorithm any further opportunity to refinance.

Acknowledgements

G. Turpin and R. El-Yaniv are very grateful to Allan Borodin for many helpful discussions and comments. R. El-Yaniv also wishes to thank Hisao Tamaki for valuable discussions.

References

[Cover91] T.M. Cover. "Universal Portfolios," *Mathematical Finance*, Vol 1, No. 1 (January 1991), 1-29.

[Ragh91] P. Raghavan. A statistical adversary for on-line algorithms. *On-Line Algorithms*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 79-83, 1991.

[ST85] D.D. Sleator and R.E. Tarjan. Amortized efficiency of list update and paging rules. *Communication of the ACM*, 28:202–208, 1985.