האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

RECORD BREAKING AND TEMPORAL CLUSTERING

by

FLAVIO TOXVAERD

Discussion Paper # 395 June 2005

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

RECORD BREAKING AND TEMPORAL CLUSTERING

Flavio Toxvaerd*†

This draft: June 2005. Preliminary, comments welcome.

Abstract. Casual observation suggests that athletics records tend to cluster over time. After prolonged periods without new records, a record breaking performance spurs other athletes to increase effort and thereby repeatedly set new standards. Subsequently, record breaking subsides and the pattern repeats itself. The clustering hypothesis is tested for the mile run, the marathon, the world hour record and long jump. For all four disciplines, the null hypothesis of non-clustering is rejected at the 4% level or below. A theoretical rationale for this phenomenon is provided through a model of social learning under limited awareness. The agents are assumed to be unaware of the true limits to performance and to take the current record as the upper bound. The observation of a record breaking achievement spurs the agents to try harder and thus temporarily increase the probability of new records. Subsequently, record breaking trails off and the process is repeated.

JEL Classification: D83, O33.
Keywords: Record breaking, temporal clustering, adaptive learning, limited awareness.

1. Introduction

On May 6, 1954, 25 year-old Roger Bannister was the first in history to run one mile in less than four minutes. With a time of 3 minutes and 59.4 seconds, he had accomplished an astonishing feat, something his contemporaries deemed impossible and which is routinely compared to Edmund Hillary and Tenzing Norgay's conquest of Mount Everest almost exactly a year earlier on May 29, 1953. According to Myers (2002), "for decades it was considered beyond human capacity, virtually in physiological principle, to run a mile inside four minutes". As it turned out, there was no real barrier of four minutes, only a perceived one. But perceptions matter.

A remarkable thing is that before Bannister's achievement, the fastest time remained unchanged for a decade. In contrast, it took a meagre seven weeks for Bannister's record to be broken by John Landy. Such clusters of records on the mile run have occurred repeatedly over the past century and a half and the phenomenon is found also in other disciplines such as the marathon, the one hour cycling race and long jump.

The aims of this paper are twofold. First, the hypothesis of temporal clustering of records is tested for the four mentioned disciplines. For all four, the null hypothesis of non-clustering

*Department of Economics, and Center for Rationality and Interactive Decision Theory, Hebrew University of Jerusalem, Mount Scopus Campus, 91905 Jerusalem, Israel. E-mail: toxvaerd@mscc.huji.ac.il. Phone: +972 (0)2 5883124. Fax:. +972 (0)2 5816071. Web: http://pluto.huji.ac.il/~toxvaerd/.

†I am pleased to acknowledge constructive conversations on the subject of this paper with David Genesove, Alex Gershkov, Chryssi Giannitsarou, Sergiu Hart, Paul Klemperer, Tobias Kretschmer, Marco Ottaviani, Peter Norman Sørensen, Sylvan Wallenstein and Joseph Zeira. I also benefited from comments from seminar participants at the London School of Economics and the Winter Meetings of the Econometric Society in Philadelphia (2005).

is rejected at an appropriate level. Second, having established the clustering phenomenon, a theoretical rationale is provided through a model of social learning under limited awareness.

To test the hypothesis that records cluster over time, methodology developed for the study of epidemics is employed. The so-called *scan statistic* detects a cluster by way of rejecting a null hypothesis of uniformity. For the four disciplines considered, the hypothesis of clustering is corroborated by the test. For the mile run, the marathon, the world hour record and long jump, the null hypothesis of uniformity is rejected at the 4% level or below.

Motivated by the established stylized fact of temporal clustering, I develop a model to explain such a pattern in the progression of records. The model builds on the premise that agents are unaware of the true limits to performance, but that they perceive that the standing record is the maximum attainable. Acting optimally given such perceptions, the existing record is broken only with a very small probability, giving rise to prolonged periods without new records. Eventually though, the record is broken, prompting the agents to revise their perceptions of what is feasible. In turn, this revision spurs subsequent agents to exert more effort, thereby increasing the probability of new records. In this way, an initial record may create a rapid succession of new records, which eventually subsides when perceptions about what is possible catch up with reality.

The progression and time series properties of records in athletics has long attracted the attention of statisticians, who have developed a rich body of research on the subject. A review of the early literature is offered by Glick (1978).[1] Roughly, the early literature focused on the characterization of record data when such records were generated by identically, independently distributed random variables, while more recent literature has relaxed the assumptions of independence and stationarity in several different ways. Yang (1975) considers the effects on records of an increasing population, while Ballerini and Resnick (1985) consider the effects of improving populations. Ballerini and Resnick (1987) study the effects on records of adding a linear trend to a stationary process. While these papers are theoretical in nature, they all illustrate their results with athletics data such as that for the mile run. However, this literature has not identified nor studied the clustering phenomenon. Another stream of research, initiated by Tryfos and Blackmore (1985), has focused on forecasting the future evolution of records. This research agenda has been refined in several different ways, e.g. by Smith and Miller (1986), Smith (1988) and most recently by Carlin and Gelfand (1993). Last, some effort has been devoted to a more applied perspective on records, within the field of sports medicine. Bassett et al. (1999) study the progression of the world hour record and estimate how much of the improvements in distance can be attributed to technological and physiological improvements respectively. Gembris, Taylor and Suter (2002) study the progression of records for a number of different sports events in the German championships and try to separate how much of improvements are due to systematic improvement of the athletes and how much is attributable to chance events.[2] Again, this literature does not touch upon the clustering phenomenon. From the economics literature, two strands of related research should be mentioned. The first is that of rational learning, surveyed in great detail by Chamley (2004). The second is that on informational cycles, as studied by Zeira (1994).

[1] See also Arnold, Balakrishnan and Nagaraja (1998) for a thorough exposition of the theory of records.

[2] Perhaps surprisingly, they find that in only four of 22 disciplines is there evidence of systematic improvement over time.

The rest of the paper is organized as follows. In Section 2, the progression of records for the mile run, the marathon, the one hour cycling race and long jump is described in some detail. The test employed for the statistical analysis of clustering, the scan statistic, is introduced and the results presented. Section 3 offers a model of temporal clustering of records based on the notion of social learning under limited awareness. Section 4 is devoted to a broader discussion of alternative explanations of the clustering phenomenon and their merits.

2. Stylized Facts of Record Progression

Casual observation of the progression of world records suggest that improvements tend to cluster over time. In particular, some records remain unbroken for prolonged periods of time, but once broken, subsequent records are achieved in relatively rapid succession.

For the purpose of illustration, the progression of records in four different disciplines will be studied in some detail. These are the mile run, the marathon, the world hour record and long jump respectively. The choice of disciplines is based on personal taste. In principle, the model offered in this paper applies equally to any sport in which the objective is well defined and easily measurable and whose record progression displays temporal clustering.

The data set employed for each discipline is the longest available time series which includes the exact dates of the records. The data on the mile run, the marathon, long jump and the world hour record were obtained from the British Milers' Association, the International Olympic Committee (IOC) and the International Cycling Union (UCI) respectively. They are included in the appendix for completeness. Throughout, the starting date of a series is taken to be the date of the first recorded observation, while the end date is 4 September, 2005.[3]

First, consider the evolution of records for the mile run. The first recorded time included in the data set was 4:52.0, achieved by Cadet Marshall in 1852. Since then, the record time has come down by more than a minute, to Hicham El Guerrouj's standing record of 3:43.12, achieved in 1999. Over the period, the record was equalled or bettered a total of 48 times, with an average inter-record time of 3 years.[4]

The distribution of new records over time is shown in the first panel of Figure 1. Inspection of the plot shows that the evolution has been less than even, with some periods seeing no improvements at all, while other periods show evidence of a flurry of activity. During the eight-year period 1915-1923, the record remained unbroken, while in the three-year period 1942-1945, it was broken six times, with three new records in 1942 alone. After that, the record remained unbroken for almost a decade, until Bannister's historical achievement in 1954 and Landy's improved time about seven weeks later.

Turning to the progression of records for the marathon, the best time has decreased a dramatic 50 minutes over the past century, from John Hayes' time of 2:55:19 in 1908 to Paul

[3]The available data sets are lists of dates on which records where equalled or bettered. To make the plots and statistical tests of this paper, the lists of dates where converted to series with one day as the time unit, with any particular day scoring 1 if the existing record was equalled or bettered on that day and 0 otherwise. There were no incidences of more than one record on any one day. The Matlab code used for the tests is available upon request.

[4]A standing record was equalled only twice over the period.

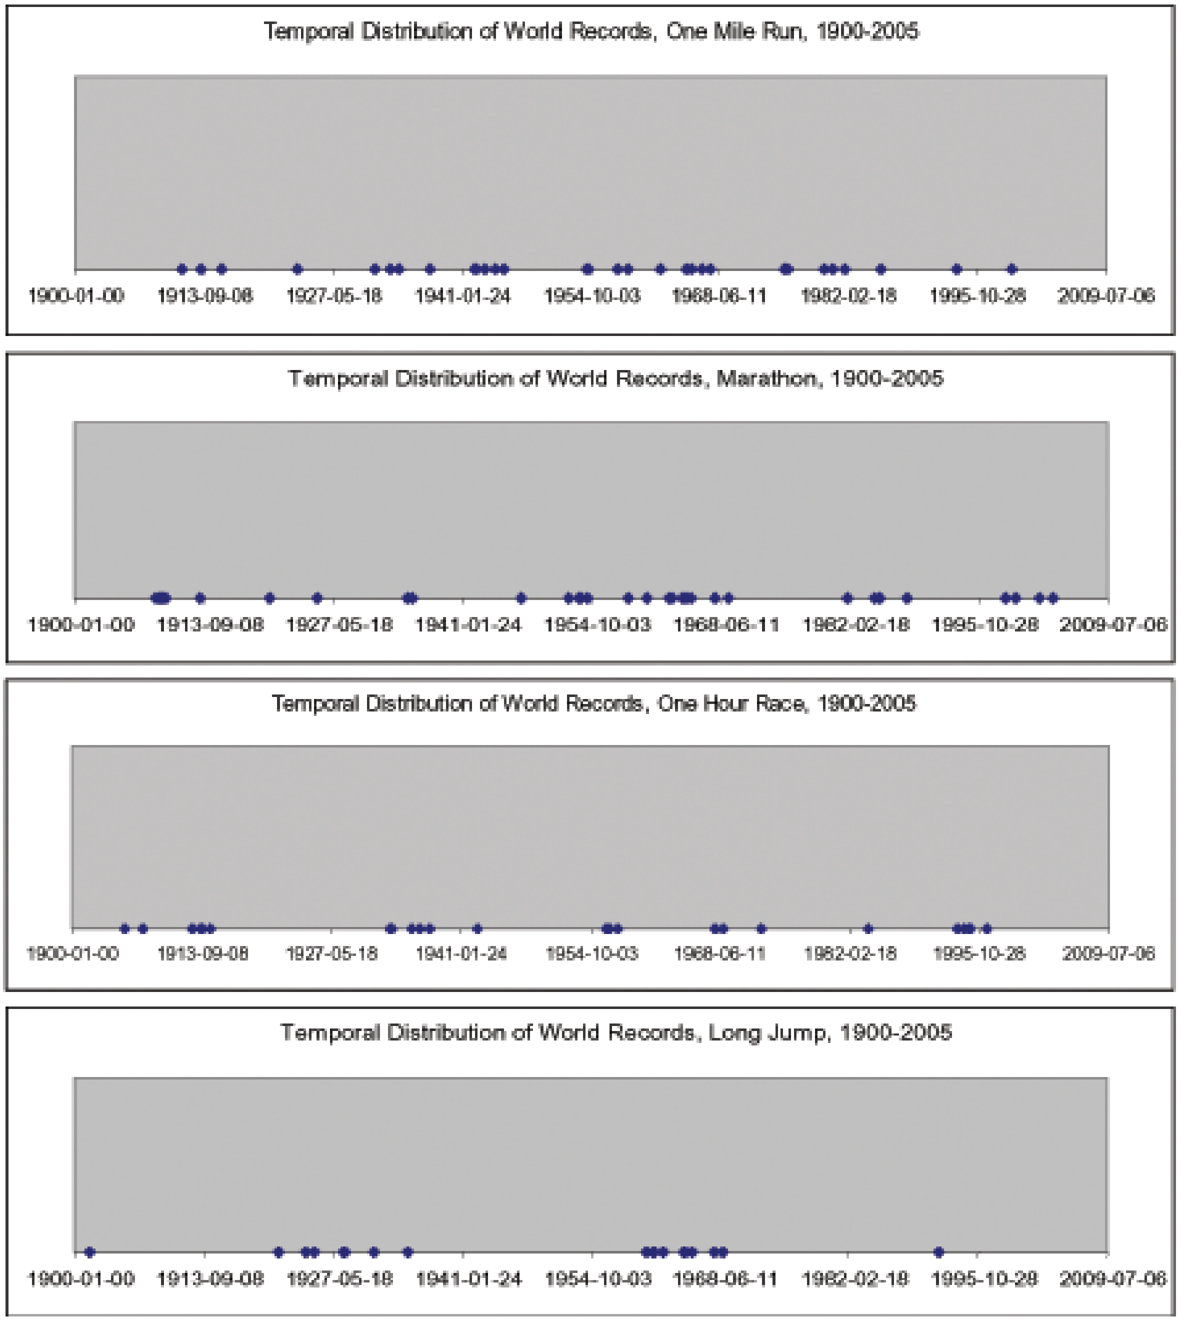

Figure 1: Temporal Distribution of Records (selected disciplines).

Tergat's standing record of 2:04:55, achieved in 2003. Over this period, the record has been bettered 34 times, with an average inter-record time of 3 years.

Like for the mile run, the decrease in the best time for the marathon has been far from smooth. E.g., 1909 alone saw nine successive new records, while the seven-year period 1913-1920 saw none. Similarly, 1935 saw three new records, while there were none in the following twelve-year period. Last, the sixties saw relatively many new records, seven, while there was not a single new record in the proceeding twelve years. The second panel of Figure 1 shows the temporal distribution of records for the marathon. Visual inspection confirms the uneven distribution over time.

A similar pattern is also evident from the progression of the world hour record in cycling. The world record has increased by more than 20km, from Henry Desgrange's distance of 35.325km in 1893 to Chris Boardman's standing record of an impressive 56.375km, achieved in 1996. Over the period, the record was bettered 32 times, with an average inter-record time of 3.3 years.

The temporal distribution of best times for one hour cycling is displayed in the third panel of Figure 1. Again, the uneven temporal pattern of records found for the mile run and the marathon is apparent. E.g., during the five-year period 1907-1912, the record remained unbroken, but it was then broken four times in the span of one year, in 1913-1914. Subsequently, the record remained unbroken for eighteen years straight, until 1935. Then, in the two-year period 1935-1937, it was broken four times.

Last, consider the progression of records for long jump, displayed in the fourth panel of Figure 1. Over the last century, from Peter O'Connor's record of 7.61m in 1901, the best result has increased by more than two meters, to Mike Powell's standing record of 8.95m, achieved in 1991. Over the period, the record was bettered or equalled 17 times, with an average inter-record time of 5 years.[5] For the long jump, the clustering phenomenon is perhaps more striking than in any other discipline, containing two of the longest standing records in any track and field discipline. O'Connor's 1901 record remained unbroken for two decades, until 1921. The next decade witnessed six new records, while Owens' 1935 record remained unbroken for a quarter of a century. Then, in the 1960's, the record was broken or bettered a total of nine times, while Beamon's 1968 records remained unbroken for more than two decades.

2.1. Statistical Analysis of Clustering. To confirm the evidence provided by visual inspection of the distribution plots of Figure 1, a formal test of the clustering hypothesis will be performed. Before introducing the particulars of the test, it is useful to explain in more detail what a cluster is and what it is not. Consider a number of points on a segment of a line, as here, the dates at which a record was equalled or broken. A number of such points are said to be clustered, if they are located in a relatively small subinterval of the line segment under consideration. That is, they are clustered if the successive distances between them are very small compared to the length of the whole interval. At the other extreme, a number of points are not clustered at all if they are evenly spaced on the line segment. Thus, there is no clustering if there is no particular subinterval of the line segment that has more points than any other subinterval of the same length.

[5] A standing record was equalled only twice over the period.

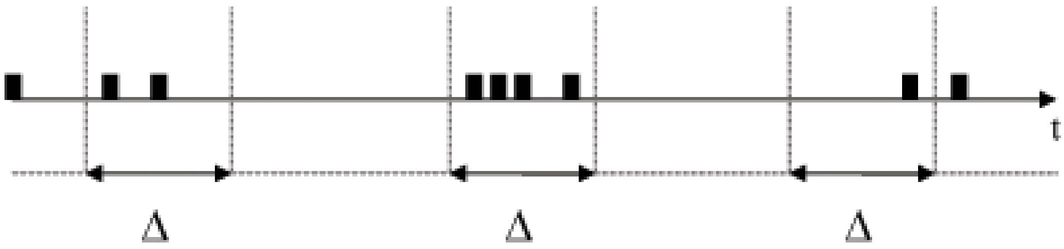

Figure 2: The Scan Statistic's Moving Window.

The next question to ask, given a particular set of observations, is what kind of distribution is likely to have generated the observed pattern. For the case of no clustering, or evenly spaced points, a natural conjecture is that they have been generated by a uniform distribution. Indeed, the uniform distribution shall be the null hypothesis against which the test of clustering will be performed. The alternative hypothesis is simply that the observations are not generated by a uniform distribution. To sum up, the data will be said to be clustered if the null hypothesis of them being generated by a uniform distribution can be rejected at a reasonable level of confidence.

A simple statistical test embodying the approach just described is based on the *scan statistic*. This test was developed by Naus (1965, 1966) and subsequently tabulated by Wallenstein (1980). It is widely used for the detection of epidemics and outbreaks of infectious diseases.

To see how the test works, consider a time interval $[0, T]$, N observations of some event occurring within this time interval and the exact time at which each of these observations occurred. Next, let $[t, t + \Delta]$ be a subinterval with $\Delta < T$ and find the maximum number of events occurring within any such interval as t slides across the range $[0, T - \Delta]$. Let m_Δ be this number of observations. Last, calculate the probability, *under the null hypothesis of uniformity*, of observing $k \geq m_\Delta$ points on an interval of length Δ. Denote this probability by

$$p \equiv \Pr(m_\Delta \geq k) \equiv P(k, N, \delta)$$

where the window $\delta \equiv \Delta/T$ is just a simplifying normalization. For the purpose of illustration, Figure 2 shows how the scan statistic's window of length Δ slides across the time line. The short black line segments illustrate dates on which events occur.

The quantities T and N are given by the data, which leaves δ to be chosen. Given some choice of window δ, the lower p is, the more unlikely it is that the observations are generated by a uniform distribution, lending credence to the alternative hypothesis of temporal clustering of the observations.

It has been shown that it is infeasible to obtain an exact form for $P(k, N, \delta)$ when T and N are large, but the function may be approximated quite well. The statistics literature on such approximations is rich and has proposed very sophisticated and accurate methods (see Glaz, Naus and Wallenstein, 2001 for details). A simple approximation, which will be employed here, is that of Wallenstein and Neff (1987). They show that

$$p = P(k, N, \delta) \approx \left(\frac{k}{\delta} - N + 1\right) \Pr(Y = k) + 2 \Pr(Y \geq k + 1)$$

where

$$\Pr(Y = k) = \left(\begin{array}{c} N \\ k \end{array}\right) \delta^k (1 - \delta)^{N-k}$$

That is, Y has a binomial distribution with parameters N and δ. This approximation is accurate up to three decimal points for low p values, but performs less well for large p values.[6]

Before presenting the results of the test, the choice of scanning window δ deserves some discussion. The scanning window is chosen to be the length of the conjectured cluster and its choice is therefore guided by visual inspection of the distribution plots. For that reason, clustering in the distributions of the different disciplines is tested with different windows.

The test results are summarized in Table 1. For the windows chosen for each discipline, the null hypothesis of uniformity is rejected at the 5% level, suggesting statistically significant clustering. The results of the tests confirm the apparent pattern of clusters seen in the distribution plots of Figure 1.

A couple of reflections on the scan statistic are in order. First, the scan statistic ignores, by construction, the exact distribution of those observations that are not included in the largest cluster and is thus unimodal in nature. In particular, a data set with two identical but non-overlapping clusters of a given size could yield the same p-values as a data set with just one cluster of the same magnitude, but with the rest of the observations distributed uniformly over the interval without the cluster. Therefore, the scan statistic is not able to detect multiple clusters or cyclicality in the data. This is pointed out by Molinari, Bonaldi and Daurés (2001), who suggest an alternative approach based on Monte Carlo experiments. For a constructed data set, they show that the presence of multiple clusters can actually mask statistical significance if the scan statistic is employed. Specifically, they show that it is possible that a data set is clustered, but that the scan statistic may fail to pick this up because of peculiarities of the distribution of observations. In a sense, this means that the scan statistic is conservative, so that if it shows clustering of a data set, this same data set would also be clustered according to the method suggested by Molinari et al. (2001). The second comment on the use of the scan statistic relates to the use of the uniform distribution as the null-hypothesis of non-clustering. The uniform distribution is not employed because it is believed to be an accurate description of the record sequence, which it will not in general be for arbitrary underlying distribution of the outcomes. In fact, visual inspection of the data already strongly suggests that it is not generated by a uniform distribution. Rather, it is used as the *definition of non-clustering*. As such, the null-hypothesis of uniformity is a very intuitive way of showing that the data is clustered.

In conclusion, the evidence favors the hypothesis that records cluster over time. This suggests two key observations. First, a cluster is evidence that athletes are not permanently at the frontier of what is humanly possible (they may, of course, be close to what they perceive to be the absolute frontier). The dichotomy between absolute limits to performance

[6] For $k \geq N/2$ and $\delta \leq 1/2$, the given formula is exact.

Discipline	Start date	N	Δ	p-value
One Mile	2 Sep. 1852	49	15	0.0036899
			30	0.0145200
			45	0.0321410
Marathon	24 July 1908	35	350	0.0080598
			400	0.0130830
			450	0.0199370
One Hour	11 May 1893	33	50	0.0211610
			60	0.0301780
			70	0.0406810
Long Jump	5 Aug. 1901	18	700	0.0261660
			750	0.0315550
			800	0.0375460

Table 1: p-values for the Scan Statistic.

and perceived ones is recognized by sports scientists. Myers (2002) paraphrases the world expert on endurance sports Tim Noakes and states that "however exhausted a runner may feel, he is likely to be a fair way from exhausting all his physical potential. The trick is to close the gap [...]".[7]

The second observation is that, once a long-standing record is broken, this event seems to shift the entire distribution of outcomes upwards, thereby increasing the probability that subsequent attempts will set new standards. A case in point is the breaking of the four minute mile barrier. John Landy, who was the first to break Bannister's record, said of his thinking before the breaking of the four minute mile: "I honestly felt, certainly after I'd run half a dozen 4.2s, that there was a bit of a barrier there". After the fact, he thought that "if he [Bannister] can run as fast as that, so can I".[8]

These two observations together suggest that athletes act in an environment with informational imperfections and that they engage in some kind of learning. Furthermore, they suggest that the effect of this learning is that athletes increase their effort, thereby inducing an upward shift in the distribution of outcomes.

There are two stylized facts of the time series that a satisfactory theory should match. The first is that after a period with no new records, when a standing record is eventually broken, new records come in relatively rapid succession. The second is that after a period with relatively many new records, the occurrence of new records subsides and gives way to a new period of tranquility.

Which kind of underlying process could one expect to yield patterns of records consistent with those actually observed? It has long been recognized that the sheer abundance of new records over time makes iid processes particularly unsuited to model real record progressions in sports. Models with deterministic trends, i.e. independently but not identically distributed

[7] According to Bannister, "though physiology may indicate respiratory and circulatory limits to muscular effort, psychological and other factors beyond the ken of physiology set the razor's edge of defeat or victory and determine how close an athlete approaches the absolute limits to performance" (quoted from Myers, 2002).

[8] Quoted from interview with John Landy, The Sports Factor, 30 April 2004: The Mile of the Century.

observations only seem to fare marginally better. That leaves models with some sort of dependence. Unfortunately, such models are very difficult to solve in general, apart from in a few cases (see Arnold et al., 1998).

Before embarking on the details of the model, a discussion of athlete's objectives is in order. It is far from clear how to model the payoffs from participating in an athletic event, as participants (in multi-player events) do not have the same goals, even if competing within the same event. For example, in a marathon, some runners seek to win on the day, others may seek to be within the top three, others may wish to improve their personal records and last, some runners may actively seek to break records or be the fastest of the year. Also, if one believes that clustering in different disciplines reflect the same underlying forces, the details of the model should not be so specific to the particulars of one discipline that it is at odds with those of the other. In particular, since clustering is found in the one hour cycling record (which is by definition a one man event), tournament incentives do not seem to be a crucial consideration when modeling the evolution of records over time. Also, if one were to focus on tournaments, it is not clear how such tournaments should be related *over time* and how payoffs from different tournaments are interlinked. For all these reasons, a pragmatic approach has been taken in which the agents are decision makers whose payoffs depend only on the outcome of their own efforts. While this is certainly a simplification of reality, it does capture some important aspects of reality and yields patterns of record breaking reminiscent of those observed in the data.

3. A Model of Social Learning with Limited Awareness

The following model is one where agents act in an environment the details of which they do not fully grasp. It is inspired by the anecdotal evidence from the history of the mile run and the marathon.[9] In particular, the agents are unaware of the true limits to performance, but learn adaptively about these by observing the outcomes of predecessors' efforts.

3.1. The Basic Setup. Time is discrete and the horizon is infinite. In each period, a new agent appears who lives for just that period. Denote by agent t the agent who lives in period t.[10] In period $t = 1, 2...$, agent t exerts effort e to produce an outcome x drawn from a continuously differentiable distribution $G \equiv G(x_t; e, z_t)$ on some bounded interval $[0, z_t]$, which allows a probability density function g.[11] In the first period, the outcome is drawn from the interval $[0, z_1]$, so that $z_1 \equiv \underline{z}$ defines the lower bound of possible feasibility frontiers for the outcome. Similarly, I assume that there is an upper bound $\bar{z}$ of possible feasibility frontiers for the outcome. A first assumption is made:

A1 $\underline{z} \leq z_t \leq z_{t+1} \leq \bar{z}$ for all t.

[9] For an account of the history of the mile run, see Bryant (2004).

[10] In fact, the assumption that a new agent appears in each period is inessential. As will become clear later, under the assumed perceptions of the agents, all the results carry through unaltered if the same agent is allowed to attempt multiple times.

[11] Note that although the distribution function changes over time (since the support changes over time), it is assumed that for all t, the distribution function comes from the same family (i.e. the functional form is the same). For notational simplicity, the time subscript on G will be suppressed, but it is implicit that at time t, the distribution G is parameterized by time t through the frontier z_t.

This assumption reflects the notion that over time, improvements in training regimes, diets and technology increase what it is possible to achieve. Imposing an upper bound on z_t reflects the fact that, for example, in running there is a time below which it is humanly/biologically impossible to run a prespecified distance.

Some further assumptions on the distribution function follow:

A2 $G(x; e, z) \leq G(x; e', z)$ for all x and z and $e \geq e'$.

Assumption A2 simply states that the distribution G is stochastically increasing in effort (i.e. it shifts in the sense of first-order stochastic dominance). This means that higher effort leads to higher outcome in expectation.

A3 $G(x; e, z) \leq G(x; e, z')$ for all $x \in [0, z']$, e and $z \geq z'$.

Assumption A3 is seemingly similar to A2, requiring that the distribution be stochastically increasing in the possibility frontier *on the shared domain*. Note that A3 contains an important implicit assumption about how the agents of the model update beliefs. In this model, the agents' information is very restricted and they do not update in Bayesian fashion. Rather, they conjecture and correctly so, that the distribution has a particular functional form. What is unknown to them is a parameter of this distribution, namely the upper limit of its support. In a Bayesian model, Bayes' rule dictates precisely how the posterior beliefs are formed. But this approach is not applicable in this model since the agents do not, by assumption, hold beliefs about the sequence $\{z_t\}$. Therefore, an alternative updating rule must be specified. What is assumed here is that the agents maintain the perception that the distribution of outcomes belongs to a given family of functions and remains so after the agents become aware that the frontier has expanded. Furthermore, it is assumed that once they become aware that the frontier has expanded, they roll over probability mass to accommodate for this expansion in the way specified in A3.

A4 $g(x; e, z) > 0$ for all $x \in [0, z]$ and e.

Assumption A4 means that the probability density is strictly positive for all feasible outcomes, i.e. the outcome distribution has full support.

Finally, define the record $\bar{x}_t$ at time t as being the maximum of all outcomes up to time $t-1$, i.e.

$$\overline{x}_t \equiv \max\{x_1, ..., x_{t-1}\}$$

By definition, the sequence of records is increasing over time, i.e. $\overline{x}_t \leq \overline{x}_{t+1}$ for all t.

3.2. The Agent's Problem. Given effort e and outcome x, the agent's utility is given by the separable function

$$u(x) - \psi(e)$$

where the utility of outcome u is strictly increasing and concave and the disutility of effort ψ is strictly increasing and convex.[12] In order to write up the agent's maximization problem,

[12] One may rewrite the problem and specify the agent's utility as $u(r(x))$, where $r(x)$ is the expected rank in some tournament setting and $r' > 0$.

some assumptions on the available information need to be made. I assume that agent t observes only the outcomes of predecessors' efforts $\{x_1, ..., x_{t-1}\}$ and knows the functional form of the family of distribution functions $\{G(x; e, z)\}_{z\in[\underline{z},\overline{z}]}$. He does not know that the frontier of feasible outcomes is expanding, nor its current position z_t. Instead, it is assumed that he perceives the outcome to be drawn from the distribution $G(x; e, \overline{x}_t)$. That is, the agent is not aware of and does not take into account, the possibility of outcomes in the interval $[\overline{x}_t, z_t]$. The agent's problem is then

$$\max_e \left\{ \int_0^{\overline{x}_t} u(x)g(x; e, \overline{x}_t)dx - \psi(e) \right\}$$

A last assumption is imposed on the problem:

A5 $\int_0^{\overline{x}_t} u(x)g''_{e\overline{x}_t}(x; e, \overline{x}_t)dx + u(\overline{x}_t)g'_e(\overline{x}_t; e, \overline{x}_t) > 0.$

Assumption A5 is technical in nature and is basically a sufficient joint condition on u and g for optimal effort to be monotone in the perceived frontier. In essence, this assumption ensures that "when potential rewards are higher, optimal effort is higher". While a more primitive set of assumptions yielding monotone optimal effort would be desirable, the search for such has been greatly complicated by the fact that virtually all existing theory on decisions under uncertainty and associated comparative statics assume no changes in the support of the random variables and focus exclusively on changes in the distribution.

The optimal level of effort e^*, given current perceptions, is implicitly given by the first order condition

$$\int_0^{\overline{x}_t} u(x)g'_e(x; e^*, \overline{x}_t)dx = \psi'(e^*)$$

where subscript e denotes a partial derivative with respect to effort.

3.3. The Dynamics of the Model. Next, turn to the dynamics of the model. Suppose that for some t, $z_t > \overline{x}_t$ and that for some small $\varepsilon > 0$,

$$|G(x; e^*(\overline{x}_t), z_t) - G(x; e^*(\overline{x}_t), \overline{x}_t)| < \varepsilon \tag{1}$$

where $e^*(\overline{x})$ denotes the optimal choice of effort given the perception that the frontier is given by $\overline{x}$. Then

Proposition 1. *Under assumptions A1-A5 and (1), social learning under limited awareness generates temporal clustering of records.*

Proof: Supposition (1) means that there is a very small probability that the outcome will be beyond the perceived frontier, given the optimal level of effort given current perceptions. Note that this does not necessarily imply that the interval $[\overline{x}_t, z_t]$ be small, but that the outcome distribution has a thin upper tail. Since the probability that the outcome falls in this range is small, the probability that the agents become aware that the record can be broken is small. There can therefore be prolonged periods of time where the agents keep acting under the perception that the true frontier is $\overline{x}_t$. During this period, the true frontier may have expanded further, beyond the current level z_t. At some point though, since the true distribution of outcomes has positive probability in the range in question (because of A4), the record *will* be broken. Formally, there exists some finite time s defined by

$$s = \min\{\tau : x_\tau \geq \overline{x}_t.\} \geq t$$

But then, since $x_s = \overline{x}_s \geq \overline{x}_t$, assumption A5 implies that $e^*(\overline{x}_s) \geq e^*(\overline{x}_t)$. In turn, assumptions A2, A1 and A3 imply that for all $x \in [0, z_t]$,

$$G(x; e^*(\overline{x}_t), z_t) \geq G(x; e^*(\overline{x}_s), z_t) \geq G(x; e^*(\overline{x}_s), z_s)$$

The first inequality in particular implies that

$$G(\overline{x}_t; e^*(\overline{x}_t), z_t) \geq G(\overline{x}_t; e^*(\overline{x}_s), z_t) \tag{2}$$

The second inequality implies that

$$G(\overline{x}_s; e^*(\overline{x}_s), z_t) \geq G(\overline{x}_s; e^*(\overline{x}_s), z_s) \tag{3}$$

for all $x \in [0, z_s]$, where it should be noted that $G(\overline{x}_s; e^*(\overline{x}_s), z_t) = 1$ for all $x \in [z_t, \overline{x}_s]$ if $z_t \leq \overline{x}_s$.

Inequality (2) means that once a certain record $\overline{x}_t$ has been broken, it is more likely that subsequent outcomes will be above $\overline{x}_t$. In turn, inequality (3) means that the *new* record $\overline{x}_s$ is more likely to be broken in subsequent attempts. This process will continue until the record (i.e. the perceived frontier) is so close to the true frontier that record breaking subsides. When $\overline{x}_t = z_t$, learning stops, while the true barrier z_t expands and again creates a discrepancy between the true and perceived frontier. At some point, the process repeats itself ■

A key supposition in generating temporal clustering of records is that the current record is below the frontier of what is currently feasible. In turn, this also shows that the model allows for the possibility that the breaking of a long-standing record does not induce further record breaking. This can happen if the new best achievement is a substantial improvement on the old record, located very close to (or at) the frontier. Such stand-alone records are also observed in practice, as seen on Figure 1.

A small discrepancy between the motivating example and the model should be mentioned. Before Bannister's record, four minutes on the mile was believed to be the limit, like two hours and ten minutes on the marathon was before Clayton's record. Thus the *existing* records were not perceived to be the limits. This suggests that the perceived frontier at time t should be $\overline{x}_t + \eta$ for some $\eta \geq 0$ rather than $\overline{x}_t$. Clearly, all the arguments of the model go through unaltered with this modification, so η is set to nil for simplicity. Also, this obviates the need to specify, in an ad-hoc way, how η is determined.[13]

Figure 3 represents the process graphically. The solid curve represents the true frontier z, while the dashed horizontal lines are the record $\overline{x}$. Last, the vertical dotted lines represent

[13] Four minutes on the mile, two hours and ten minutes on the marathon, as well as 10,000 on the NASDAQ, seem to acquire intrinsic importance, although of course such beliefs are baseless. Rather, round numbers acquire attention because they are elegant and thus become somehow "focal".

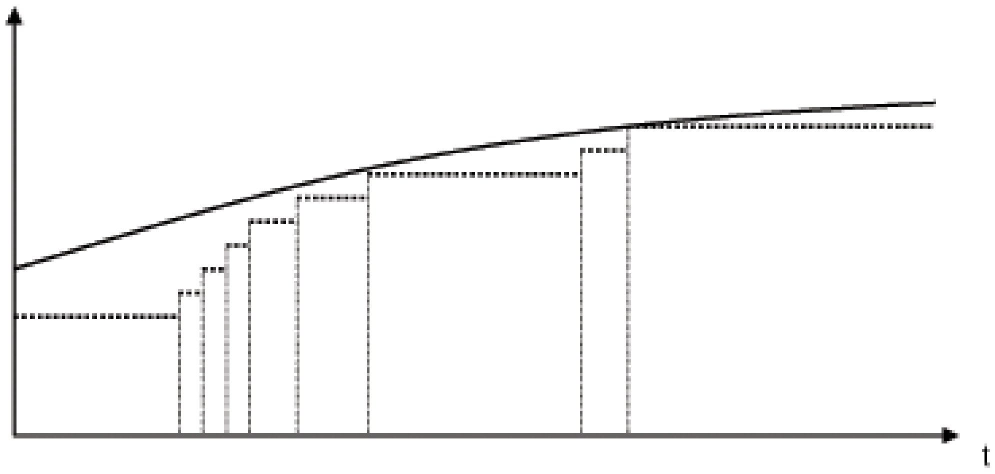

Figure 3: Temporal Clustering with Limited Awareness.

dates at which a record was broken. The temporal pattern shown in Figure 3 is consistent with that of the four disciplines shown in Figure 1. Starting from a best outcome below the feasibility frontier, an initial broken record successively ratchets up the record through the described process until the frontier is reached. The process then stops, until the frontier has expanded further and another record breaking performance starts the process anew.

4. DISCUSSION

In this paper, the phenomenon of temporal clustering of record breaking was studied. The hypothesis of clustering was tested for the mile run, the marathon, the world hour record in cycling and long jump and was in each case found statistically significant.

To explain the phenomenon, a parsimonious model was presented that yields patterns of record breaking consistent with the evidence. The three basic features of the model are *learning from the observation of others*, *monotone optimal effort* and that the *outcome distribution is stochastically increasing in effort*. When a long-standing record is broken, learning alerts other agents that some environmental factor has changed. In turn, monotone optimal effort means that agents, in reacting optimally to what they have observed, try harder than they otherwise would have. Last, since higher outcomes become more likely when effort is higher, record breaking becomes more likely. Eventually, record breaking subsides because the agents have fully adapted to the new environment.

Two possible alternative explanations to the presented one should be mentioned. First is the notion that clusters could be driven wholly by exogenous events. One may argue that the best athletes almost always perform close to their absolute best but that once in a while, there is some exogenous change, such as improved equipment or training routines, that within a relatively short time span simultaneously improves the performance of the athletes, thereby prompting them to break existing records. Once the full potential is reached, given the new conditions, record breaking subsides or trails off until a new exogenous change occurs. While such exogenous changes indeed occur and are important, they seem unable to catch an important element of interaction between athletes' performances and furthermore, run counter to the lore of the different disciplines.

A second interesting question is whether the observed pattern of clustering in the progression of records can be explained through the strategic behavior of athletes, i.e. if clustering could be the equilibrium outcome of some game situation. In this context, it should be noted that athletes, in general, compete for many reasons, not necessarily to break records. After his famous achievement, Derek Clayton declared that "I was hoping to improve my personal best, but I never thought of a possibility of setting the world marathon best. In fact, I did not even think about winning until after half way".[14]

But breaking a record is undoubtedly icing worth having for any athlete and a long-standing one even more so. A case can be made that the longer a record has remained unbroken, the more prestigious it is to be the one to finally break it. This line of thinking is suggestive of a preemption type game. Still, strategic timing of record breaking seems highly unlikely. Athletes have a limited number of years at their physical peak and must perform bearing this fact in mind. Also, such considerations do not suggest why there should be clustering of records after the breaking of a long-standing best. If anything, the incentive to break the new record should be lower than the incentive to break the original one.

Temporal clustering of events has previously been studied, e.g. by Gul and Lundholm (1995), Grenadier (1996) and Toxvaerd (2004). They all study models in which agents choose, in equilibrium, to act simultaneously or in rapid succession. Importantly though, deciding when to act is wholly different in nature to the decision of breaking a long-standing record. Extremely few athletes can do more than merely dream of breaking a world record, so the idea of strategic timing of such attempts seem somewhat unconvincing.

Last, it should be noted that the model studied in this paper can plausibly be interpreted as a stylized model of technological innovation. An innovation, especially a process innovation, can be seen as simply an improvement upon the existing state of affairs. But an improvement upon the existing is in turn the breaking of a record. An extension of the current work could be to explore this link further and contrast it with existing literature on growth and endogenous innovation waves, such as Helpman and Trajtenberg (1998), Andergassen and Nardini (2003).

[14]Quoted in Nakamura (1967).

REFERENCES

[1] ANDERGASSEN, R. AND F. NARDINI, 2003. Endogenous Innovation Waves and Economic Growth, *unpublished manuscript*.

[2] ARNOLD, B. C., N. BALAKRISHNAN AND H. N. NAGARAJA, 1998. Records, *John Wiley & Sons*.

[3] BALLERINI, R. AND S. RESNICK, 1985. Records From Improving Populations, *Journal of Applied Probability*, 22, 487-502.

[4] BALLERINI, R. AND S. RESNICK, 1987. Records in the Presence of a Linear Trend, *Advances in Applied Probability*, 19, 801-828.

[5] BASSETT, D. R. JR, C. R. KYLE, L. PASSFIELD, J. P. BROKER, E. R. BURKE, 1999. Comparing Cycling World Hour Records, 1967-1996: Modeling with Empirical Data, *Medicine & Science in Sports & Exercise*, 31(11), 1665-1676.

[6] BRYANT, J., 2004. 3:59.4. The Quest to Break the 4 Minute Mile, *Hutchinson*, London.

[7] CARLIN, P. B. AND A. E. GELFAND, 1993. Parametric Likelihood Inference for Record Breaking Problems, *Biometrika*, 80(3), 507-15.

[8] CHAMLEY, C. P., 2004. Rational Herds: Economic Models of Social Learning, *Cambridge University Press*.

[9] GEMBRIS, D. AND J. G. TAYLOR AND D. SUTER, 2002. Trends and Random Fluctuations in Athletics, *Nature*, 417, 506.

[10] GLAZ, J., J. I. NAUS AND S. WALLENSTEIN, 2001. Scan Statistics, *Springer-Verlag*.

[11] GLICK, N., 1978. Breaking Records and Breaking Boards, *American Mathematical Monthly*, 85, 2-26.

[12] GRENADIER, S. R., 1996. The Strategic Exercise of Options: Development, Cascades and Overbuilding in Real Estate Markets, *Journal of Finance*, 51(5), 1653-1679.

[13] GUL, F. AND R. LUNDHOLM, 1995. Endogenous Timing and the Clustering of Agent's Decisions, *Journal of Political Economy*, 103(5), 1039-1066.

[14] HELPMAN, E. AND M. TRAJTENBERG, 1998. A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies, in *General Purpose Technologies and Economic Growth*, E. Helpman (ed), MIT Press.

[15] MOLINARI, N. C. BONALDI AND J.-P. DAURES, 2001. Multiple Temporal Cluster Detection, *Biometrics*, 57, 577-583.

[16] MYERS, N., 2002. Running Through the Mind, *Nature*, 418, 371.

[17] NAKAMURA, K. K., 1967. A History of the Fukuoka International Marathon Championships, *International Association of Athletics Federations.*

[18] NAUS, J. I., 1965. The Distribution of the Size of the Maximum Cluster of Points on a Line, *Journal of the American Statistical Association*, 60, 532-38.

[19] NAUS, J. I., 1966. Some Probabilities, Expectations and Variances for the Size of Smallest Intervals and Largest Clusters, *Journal of the American Statistical Association*, 61, 1191-99.

[20] SMITH, R. L., 1988. Forecasting Records by Maximum Likelihood, *Journal of the American Statistical Association*, 83(402), Applications and Case Studies, 331-338.

[21] SMITH, R. L. AND MILLER, 1986. A Non-Gaussian State Space Model and Application to the Prediction of Records, *Journal of the Royal Statistical Society*, Series B, 48, 79-88.

[22] TOXVAERD, F., 2004. Strategic Merger Waves: A Theory of Musical Chairs, *Working Paper No. 359, Centre for Rationality and Interactive Decision Theory, Hebrew University of Jerusalem.*

[23] TRYFOS, P. AND R. BLACKMORE, 1985. Forecasting Records, *Journal of the American Statistical Association*, 80(389), Applications, 46-50.

[24] WALLENSTEIN, S., 1980. A Test for Detection of Clustering Over Time, *American Journal of Epidemiology*, 111(3), 367-72.

[25] WALLENSTEIN, S. AND N. NEFF, 1987. An Approximation for the Distribution of the Scan Statistic, *Statistics in Medicine*, 6, 197-207.

[26] YANG, M. C. K., 1975. On the Distribution of the Inter-Record Times in an Increasing Population, *Journal of Applied Probability*, 12, 148-154.

[27] ZEIRA, J., 1994. Informational Cycles, *Review of Economic Studies*, 61(1), 31-44.

Progression of World Record: Mile Run		
Result	**Date**	**Name**
4:52.00	02 Sep 1852	Cadet Marshall
4:45.00	03 Nov 1858	Thomas Finch
4:45.00	15 Nov 1858	Vincent Hammick
4:40.00	24 Nov 1859	Gerald Surman
4:33.00	23 May 1862	George Farran
4:29.60	10 Mar 1868	Walter Chinnery
4:28.80	03 Apr 1868	William Gibbs
4:28.60	31 Mar 1873	Charles Gunton
4:26.00	30 May 1874	Walter Slade
4:24.50	19 Jun 1875	Walter Slade
4:23.20	16 Aug 1880	Walter George
4:19.40	03 Jun 1882	Walter George
4:18.40	21 Jun 1884	Walter George
4:17.80	26 Aug 1893	Thomas Conneff
4:17.00	06 Jul 1895	Fred Bacon
4:15.60	28 Aug 1895	Thomas Conneff
4:15.40	27 May 1911	John Paul Jones
4:14.40	31 May 1913	John Paul Jones
4:12.60	16 Jul 1915	Norman Taber
4:10.40	23 Aug 1923	Paavo Nurmi
4:09.20	04 Oct 1931	Jules Ladoumégue
4:07.60	15 Jul 1933	Jack Lovelock
4:06.80	16 Jun 1934	Glenn Cunningham
4:06.40	28 Aug 1937	Sydney Wooderson
4:06.20	01 Jul 1942	Gunder Hägg
4:06.20	10 Jul 1942	Arne Andersson
4:04.60	04 Sep 1942	Gunder Hägg
4:02.60	01 Jul 1943	Arne Andersson
4:01.60	18 Jul 1944	Arne Andersson
4:01.40	17 Jul 1945	Gunder Hägg
3:59.40	06 May 1954	Roger Bannister
3:58.00	21 Jun 1954	John Landy
3:57.20	19 Jul 1957	Derek Ibbotson
3:54.50	06 Aug 1958	Herbert Elliott
3:54.40	27 Jan 1962	Peter Snell
3:54.10	17 Nov 1964	Peter Snell
3:53.60	09 Jun 1965	Michel Jazy
3:51.30	17 Jul 1966	Jim Ryun
3:51.10	23 Jun 1967	Jim Ryun
3:51.00	17 May 1975	Filbert Bayi
3:49.40	12 Aug 1975	John Walker
3:49.00	17 Jul 1979	Sebastian Coe
3:48.80	01 Jul 1980	Steve Ovett
3:48.53	19 Aug 1981	Sebastian Coe
3:48.40	26 Aug 1981	Steve Ovett
3:47.33	28 Aug 1981	Sebastian Coe
3:46.32	27 Jul 1985	Steve Cram
3:44.39	05 Sep 1993	Noureddine Morceli
3:43.12	07 Jul 1999	Hicham El Guerrouj

Progression of World Record: Marathon		
Result	**Date**	**Name**
2:55:19	24 Jul 1908	John Hayes
2:52:46	01 Jan 1909	Robert Fowler
2:46:53	12 Feb 1909	James Clark
2:46:05	08 May 1909	Albert Raines
2:42:31	26 May 1909	Henry Barrett
2:40:35	31 Aug 1909	Thure Johansson
2:38:17	12 May 1913	Harry Green
2:36:07	31 May 1913	Alexis Ahlgren
2:32:36	22 Aug 1920	Hannes Kolemainen
2:29:02	12 Oct 1925	Albert Michelsen
2:27:49	31 Mar 1935	Fusashige Suzuki
2:26:44	03 Apr 1935	Yasuo Ikenaka
2:26:42	03 Nov 1935	Kitei Son
2:25:39	19 Apr 1947	Bok-Suh Yun
2:20:43	14 Jun 1952	Jim Peters
2:18:41	13 Jun 1953	Jim Peters
2:18:35	04 Oct 1953	Jim Peters
2:17:40	26 Jun 1954	Jim Peters
2:15:17	24 Aug 1958	Sergey Popov
2:15:16	10 Sep 1960	Abebe Bikila
2:15:15	17 Feb 1963	Toru Terasawa
2:14:28	15 Jun 1963	Buddy Edelen
2:13:55	13 Jun 1964	Basil Heatley
2:12:12	21 Oct 1964	Abebe Bikila
2:12:00	12 Jun 1965	Morio Shigematsu
2:09:37	03 Dec 1967	Derek Clayton
2:08:34	30 May 1969	Derek Clayton
2:08:18	06 Dec 1981	Rob de Castella
2:08:05	21 Oct 1984	Steve Jones
2:07:12	20 Apr 1985	Carlos Lopes
2:06:50	17 Apr 1988	Belayneh Dinsamo
2:06:05	20 Sep 1998	Ronaldo da Costa
2:05:42	24 Oct 1999	Khalid Khannouchi
2:05:38	14 Apr 2002	Khalid Khannouchi
2:04:55	28 Sep 2003	Paul Tergat

Progression of World Record: One Hour Race		
Result	**Date**	**Name**
35325	11 May 1893	Henry Desgrange
38220	31 Oct 1894	Jules Dubois
39240	30 Jul 1897	Oscar van Den Eynde
40781	03 Jul 1898	William Hamilton
41110	08 Aug 1905	Lucien Petit-Breton
41520	20 Jun 1907	Marcel Berthet
42122	22 Aug 1912	Oscar Egg
42741	07 Aug 1913	Marcel Berthet
43525	21 Aug 1913	Oscar Egg
43775	20 Sep 1913	Marcel Berthet
44247	08 Aug 1914	Oscar Egg
44588	08 Aug 1933	Jan Van Hout
44777	28 Sep 1933	Maurice Richard
45090	31 Oct 1935	Giuseppe Olmo
45325	14 Oct 1936	Maurice Richard
45485	29 Sep 1937	Frans Slaats
45767	03 Nov 1937	Maurice Archambaud
45798	07 Nov 1942	Fausto Coppi
46159	06 Jun 1956	Jacques Anquetil
46394	19 Sep 1956	Ercole Baldini
46923	18 Sep 1957	Roger Riviere
48093	30 Oct 1967	Ferdinand Bracke
48653	10 Oct 1968	Ole Ritter
49432	25 Oct 1972	Eddy Merckx
50808	19 Jan 1984	Francesco Moser
51151	23 Jan 1984	Francesco Moser
51596	07 Jul 1993	Graeme Obree
52270	23 Jul 1993	Chris Boardman
52713	04 Apr 1994	Graeme Obree
53040	02 Sep 1994	Miguel Indurain
53832	22 Oct 1994	Tony Rominger
55291	05 Nov 1994	Tony Rominger
56375	09 Sep 1996	Chris Boardman

Progression of World Record: Long Jump		
Result	**Date**	**Name**
7.61	05 Aug 1901	Peter O'Connor
7.69	23 Jul 1921	Ed Gourdin
7.76	07 Jul 1924	Bob LeGendre
7.89	13 Jun 1925	William de Hart Hubbard
7.90	07 Jul 1928	Edward Hamm
7.93	09 Sep 1928	Silvio Cator
7.98	27 Oct 1931	Chuhei Nambu
8.13	25 May 1935	Jesse Owens
8.21	12 Aug 1960	Ralph Boston
8.24	27 May 1961	Ralph Boston
8.28	16 Jul 1961	Ralph Boston
8.31	10 Jun 1962	Igor Ter-Ovanesyan
8.31	15 Aug 1964	Ralph Boston
8.34	12 Sep 1964	Ralph Boston
8.35	29 May 1965	Ralph Boston
8.35	19 Oct 1967	Igor Ter-Ovanesyan
8.90	18 Oct 1968	Bob Beamon
8.95	30 Aug 1991	Mike Powell