arXiv:cond-mat/0109243v1 [cond-mat.dis-nn] 13 Sep 2001

Antipersistent binary time series

Richard Metzler
Institut für Theoretische Physik und Astrophysik, Universität Würzburg, Am Hubland, D-97074 Würzburg, Germany

Completely antipersistent binary time series are sequences in which every time that an N-bit string μ appears, the sequence is continued with a different bit than at the last occurrence of μ. This dynamics is phrased in terms of a walk on a DeBruijn graph, and properties of transients and cycles are studied. The predictability of the generated time series for an observer who sees a longer or shorter time window is investigated also for sequences that are not completely antipersistent.

I. INTRODUCTION

The analysis and generation of time series has been of interest to physicists in different fields: it yields insight into the dynamics of chaotic systems [1,2] and is used to study such diverse systems as the climate and the heartbeat [3], linguistics [4], and, in recent times, financial markets [5,6]. The prediction of time series generated by complex physical systems can be of immense importance, as evidenced by the efforts put into improving the weather forecast. One of the most common aspects of time series is long-term memory, or persistence. This paper, however, will deal with the opposite behaviour, namely antipersistence.

Antipersistence is the tendency of a time series to show, at one point, the opposite behaviour than at some point in the past. This concept is usually applied to continuous time series [7], but has received little attention due to the dearth of natural phenomena that exhibit antipersistence (for an exception, see [8]).

In this paper, antipersistence in binary time series is defined as follows: a sequence of bits is antipersistent on a scale of N if a string μ of N bits length is likely to be followed by the opposite bit than the last time that μ appeared. This property was observed, for example, in some parameter regimes of the conventional Minority Game [9,10] and in various alternative models based on the idea of the Minority Game [11].

This paper will first deal with completely antipersistent time series from the point of view of discrete dynamical systems. Properties of cycles and transients will be studied in Sections III and IV, respectively. Concepts and results from graph theory turn out to yield much insight.

In Section V, the sequence is treated as an antipredictable sequence for a special prediction algorithm, and antipersistence on one time scale is related to properties on other time scales. Then a more general stochastic model is introduced, and the effects of stochasticity on the previous results are discussed.

II. THE MODEL

An infinitely long, completely antipersistent binary time series of bits $s_\tau \in \{0,1\}$ can by definition be generated in the following fashion: at time t, the history μ_t denotes the binary representation of the last N bits $s_{t-N+1},\ldots,s_t$. For each of the 2^N possible histories μ there is an entry $a_t^\mu \in \{0,1\}$ in a decision table A_t, which also depends on t. At each time step,

- a new bit is generated by taking the table entry corresponding to the current history μ_t (which I will also refer to as "the pattern"): $s_{t+1} = a_t^{\mu_t}$;
- the history is updated: $\mu_{t+1} = (2\mu_t + s_{t+1}) \bmod 2^N$; i.e., all bits are shifted one position to the left (multiplication with 2), the newly generated bit is added, and the oldest (most significant) bit is dropped (division modulo 2^N);
- the table entry $a_t^{\mu_t}$ that was used for making the decision is changed, such that the sequence will be continued with the opposite decision when the pattern μ_t occurs the next time: $a_{t+1}^{\mu_t} = 1 - a_t^{\mu_t}$. All other entries remain unchanged.

This last point is especially important. It means that the table entries are dynamical variables, and the state of the dynamical system is determined by the current pattern *and* the state of the decision table. Fig. 1 shows two steps of the described dynamics of such a system for $N = 2$.

t=0 t=1 t=2

$\mu_0 = \mathbf{00} \to s_1 = \mathbf{0} \to \mu_1 = \mathbf{00} \to s_2 = \mathbf{1} \to \mu_2 = \mathbf{01} \to s_3 = \mathbf{0} \to \ldots$

$A_1^{00} = \mathit{1}$ $A_2^{00} = \mathit{0}$ $A_3^{01} = \mathit{1}$

μ	A_0^μ
00	**0**
01	0
10	1
11	1

μ	A_1^μ
00	***1***
01	0
10	1
11	1

μ	A_2^μ
00	*0*
01	**0**
10	1
11	1

FIG. 1. An example of a decision table with $N = 2$ and two steps of the dynamics. Boldface numbers indicate the current history and the table entry used for continuing the sequence; italic numbers denote the last table entry that was changed. The sequence generated in this example is 010 ....

The model can also be considered from a graph-theoretical perspective: each pattern μ corresponds to a node on a directed DeBruijn graph of order N (see Fig. 2 for an example of such a graph). Each node obviously has two edges entering it, coming from the two possible predecessors, which I denote ${}^0\mu$ and ${}^1\mu$. For example, if $\mu = 1100$, the possible predecessors are ${}^0\mu = 0110$ and ${}^1\mu = 1110$. Each edge also has two outgoing edges, leading to the two possible successors μ^0 and μ^1 (for $\mu = 1100$, the successors are $\mu^0 = 1000$ and $\mu^1 = 1001$). The graph is connected, since one can reach each node from any other node in a maximum of N steps by taking the appropriate exit edges.

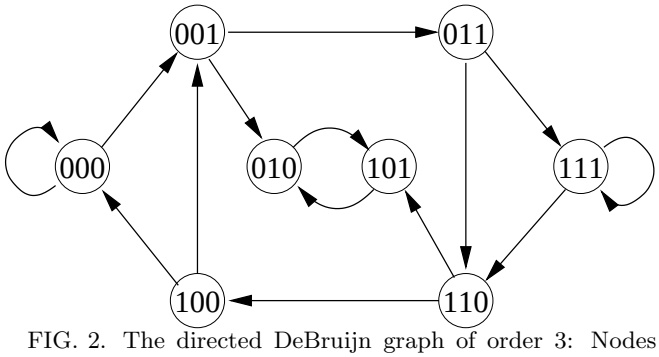


FIG. 2. The directed DeBruijn graph of order 3: Nodes represent binary strings of length 3, edges lead to strings that are generated by shifting the current string one position and adding either 0 or 1 as the new least significant bit.

In our model, at any time, only one of the exits leaving each node is labelled "active" – the one corresponding to the table entry a_t^μ. A time step consists of travelling from the current node to the next along the active exit, then "burning the bridge", i.e. labelling the previously active exit inactive and vice versa.

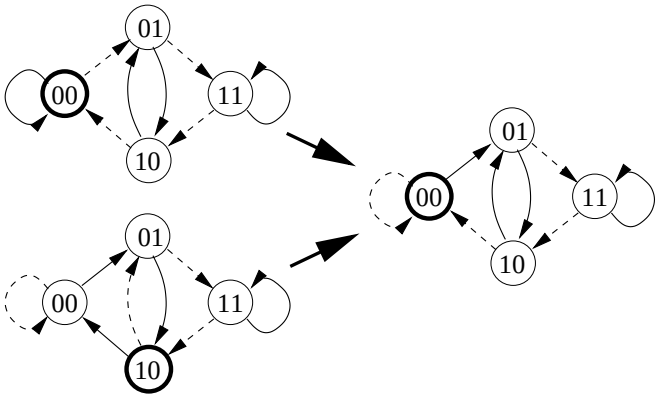


FIG. 3. Example for an irreversible situation on a graph of order 2. Active exits are denoted by solid lines, inactive ones by dashed lines. The bold circle indicates the node currently visited. Both the upper left configuration (which happens to be part of a cycle, and corresponds to the example in Fig. 1) and the lower one (which is part of a transient) lead to the configuration on the right.

III. PROPERTIES OF CYCLES

The introduced dynamics is deterministic, and the combined system of pattern and table has a finite number $\Omega = 2^N \cdot 2^{(2^N)}$ of different states, so the dynamics necessarily leads into a cycle eventually. The dynamics is irreversible: if a currently visited node has two inactive entrances, it is impossible to tell which path the system took to get to its current state (for an example, see Fig. 3). This means that not every state can be part of a cycle, so we will have to consider the necessary conditions for being in a cycle. I will show, step by step, that all cycles are of length $2 \cdot 2^N$ and touch all nodes exactly twice.

Some of the proofs that now follow are redundant; on the other hand, they help to understand the properties of the system, and some of them are applicable to generalizations of the problem, whereas others are not.

Let us assume that at time 0, the system is already moving on a cycle of length l. We count the number of times that a history μ has occurred between time 0 and time t by a visit number v_t^μ. Since the definition of a cycle is that after l steps the system must be in the same state again, it is necessary that v_l^μ *is even for all* μ, since the table entries a^μ return to their original state only after even numbers of visits.

Also, *all possible nodes are part of the cycle.* Let us prove this by assuming the opposite, namely that there are some nodes that are not touched by the cycle. Since the graph is connected, there must be unused connections between the part of the graph involved in the cycle and the part that is left out. But as we have seen in the paragraph before, the visit number of each of the nodes that are actually part of the cycle must be at least 2 (larger than 0, and even), so each of its two exits is used, including the one leading to the part of the graph supposedly not included in the cycle. This is a contradiction, so all nodes are involved.

An even stronger statement is possible: the total number of visits to the predecessors ${}^0\mu$ and ${}^1\mu$ of μ must be equal to twice the number of visits to μ, since exactly half of the visits they get are followed by μ, while the other half is followed by the so-called conjugate state $\bar{\mu}$ of μ. (For example, 001 is the conjugate state of 000.) Thus, we have

$$v^\mu = (v^{{}^0\mu} + v^{{}^1\mu})/2 \quad \text{for all } \mu. \tag{1}$$

This can be written as a linear equation for an eigenvector with eigenvalue 1 of a matrix M, with entries $a_{\nu\mu} = 1/2$ if μ is a possible successor of ν and $a_{\nu\mu} = 0$ otherwise. For example, for $N = 2$, the set of Eqs. (1) looks as follows:

$$\left[\begin{pmatrix} \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & 0 \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix} - \mathbb{1} \right] \begin{pmatrix} v_l^{00} \\ v_l^{01} \\ v_l^{10} \\ v_l^{11} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \tag{2}$$

Since the sum of columns in matrix M is always 1, and the individual entries are ≥ 0, and it describes transitions on a connected graph, we can apply results from the theory of stochastic matrices to state that it has one unique eigenvector with eigenvalue 1 [12], and we easily guess that $v_l^\mu =$ const fulfills Eq. (1). That means that in a cycle, *all states are visited with the same frequency.*

The next step is to show that in a cycle, each node is exactly visited twice, i.e., all cycles of length $4\cdot 2^N$, $6\cdot 2^N$, and so on, are in fact two, three or more repetitions of a $2\cdot 2^N$-cycle. Again, assume the system is moving on a cycle. If this cycle were truly longer than $2\cdot 2^N$, there must, at the point $t = 2\cdot 2^N$, be nodes that have been visited three or more times while others have not been visited for the second time – the visit numbers must add up to $2\cdot 2^N$, and if all visit numbers were equal to 2, the cycle would be complete. More specifically, there must be an earlier time when all visit numbers v_t^ν are either 0, 1, or 2, and one node is about to be visited for the third time. It suffices to show that this cannot happen to prove that the cycle cannot be longer than $2\cdot 2^N$.

The third visit to a node μ with $v^\mu = 2$ cannot come from a predecessor (let us say, ${}^0\mu$) with a visit number of $v^{{}^0\mu} = 0$, for the obvious reason that this predecessor has not been visited yet. It also cannot come from a predecessor with $v^{{}^0\mu} = 1$: if $v^\mu = 2$, either it must have been visited before from ${}^0\mu$ (which it cannot – the predecessor has only had one visit so far), or it must have been reached twice from ${}^1\mu$ – this is impossible as well, since it means that $v^{{}^1\mu} \geq 3$. For similar reasons, we can exclude a visit from a node with $v^{{}^0\mu} = 2$: either $v^{{}^1\mu} \geq 3$ as before, or the first visit to ${}^0\mu$ led to μ – then the second cannot. This means that all nodes must receive two visits – thus finishing a cycle – before one of them can be visited for the third time. The question arises why this line of reasoning does not hold true during the transient. The key lies in the observation that the first node to receive three visits is the node where the system was started – it did not get its first visit *from* anywhere on the graph, so the arguments are not applicable.

Using all previous conclusions, the cycles turn out to be solutions to a well-studied combinatorial problem, and the number of different cycles can be found in the literature: since, during a cycle, each string μ appears exactly once followed by each of ist successors, the cycle is a sequence in which each $N+1$-bit pattern, i.e. each node on the DeBruijn graph of order $N+1$, appears exactly once. Such a sequence is a *Hamiltonian circuit*, also known as a *full cycle*, on the $N+1$-graph. For a review on the properties of these cycles, consult Ref. [13]. One of the earliest and most central results on the topic is the number of different cycles, which is $2^{2^N-(N+1)}$ [14].

All possible full cycles on the $N+1$-graph can be generated by the antipersistent walk on the N-graph: write down the desired sequence starting at some arbitrary point, look for the first occurrence of each N-bit pattern μ, and set the corresponding table entry to the bit that follows it. Starting the antipersistent walk at the first pattern of the desired sequence, the antipersistent walk will reproduce it.

Since all cycles are of length $2\cdot 2^N$, a total of $2\cdot 2^N \times 2^{2^N-(N+1)} = 2^{(2^N)}$ states is part of a cycle. As mentioned before, the total number of possible states is $\Omega = 2^N\cdot 2^{(2^N)}$, which means that a fraction of $2^{(2^N)}/\Omega = 2^{-N}$ of possible states is part of a cycle.

IV. PROPERTIES OF THE TRANSIENT

The length τ of the transient is the number of steps taken before the system enters a cycle. The distribution of transient lengths is less accessible to analytical approaches, but easy to measure in computer simulations, either by complete enumeration for small systems or by Monte Carlo for larger ones. The following picture emerges, as seen in Fig. 4:

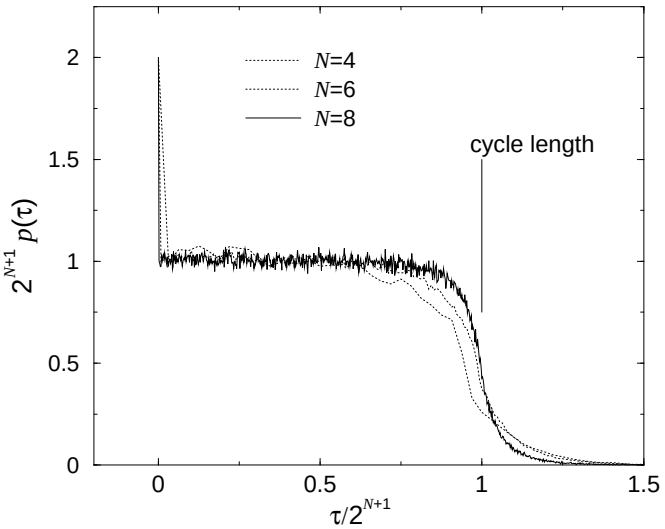


FIG. 4. Distribution of transient lengths τ, rescaled by the cycle length 2^{N+1}.

The probabilty for transient length $\tau = 0$ is just the probability of hitting a cycle right away, and thus the fraction of state space filled with cycles. As mentioned, this is equal to 2^{-N}.

The probability distribution is more or less flat for $1 \leq \tau \leq 2^{N+1}$, the cycle length. From normalization constraints, it follows that $p(\tau) \approx 2^{-(N+1)}$ in that range.

Near $\tau = 2^{N+1}$, there is an exponential drop reminiscent of a phase transition, which gets steeper with increasing N. Even for small N, no transients longer than 2^{N+2} have been observed.

V. ANTIPERSISTENCE ON DIFFERENT TIMESCALES

Consider the following prediction algorithm: an observer looks at N_{obs}-bit strings from a binary time series, writes down the bit that followed the pattern in the appropriate entry of his decision table (which, of course, has $2^{N_{obs}}$ rows), and predicts that when that pattern occurs the next time, it will be followed by that same bit written in his table.

This algorithm, which could be labelled "blind reliance on recent experience", works well for persistent time series, and it is similar in spirit to what we all do instinctively – similar situations usually lead to similar consequences. An interesting quantity is the success rate $s(N, N_{obs})$ of this prediction algorithm when it predicts a completely antipersistent walk on the N-graph. For simplicity's sake, we consider the long-time limit, in which both the generator and the observer move on a cycle.

If the observer looks at the same time window as the generator ($N_{obs} = N$), it is obvious that the success rate will be 0 – since each pattern is continued with alternating bits on each visit. In that sense, the generated time series is *anti-predictable* for this specific prediction algorithm (for more on anti-predictable sequences, see [15,16]).

For an observer with a slightly larger window, the picture changes: as mentioned above, the antipersistent cycle corresponds to a Hamiltionian cycle on the $N+1$-graph, which is completely persistent and predictable with 100% accuracy. For even larger N_{obs}, the antipersistent cycle looks like a closed path which includes only a fraction of $2^{N-(N_{obs}+1)}$ of nodes on the N_{obs}-graph. Prediction is again 100% reliable, and the observer does not even need all of his storage capacity to handle the cases that occur.

If the observer has a shorter time window than the generator, more than one of the generator's patterns will affect the same table entry for the observer. For example, an $N-1$-bit pattern ν corresponds to either of the N-bit strings 0ν or 1ν, both of which occur twice in the N-cycle, each time followed by a different successor. The success rate of the predictor depends on the sequence in which these combinations occur; if each permutation of $0\nu0$, $0\nu1$, $1\nu0$ and $1\nu1$ has the same probability, the success rate for all patterns is the average over the different permutations. Figure 5 shows that this average is 1/3 for $N_{obs} = N-1$.

For $N_{obs} = N-2$, all permutations of eight combinations of predecessors and successors have to be taken into account – a task best left to computer algebra programs, which yield $\langle s(N, N-1)\rangle = 3/7$, in excellent agreement with simulations (see Fig. 6. Larger differences in the time window are beyond even the scope of computer programs; however, it can be argued that for larger $N-N_{obs}$, the visits to the N_{obs}-nodes become more and more random, and $s(N, N_{obs})$ will tend to 1/2.

permutation				e.r.
0μ0	0μ1	1μ0	1μ1	4/4
0μ0	0μ1	1μ1	1μ0	2/4
0μ0	1μ0	0μ1	1μ1	2/4
0μ0	1μ0	1μ1	0μ1	2/4
0μ0	1μ1	1μ0	0μ1	4/4
0μ0	1μ1	0μ1	1μ0	2/4
		total:		16/24

FIG. 5. Possible permutations of predecessors and successors of a pattern μ on a cycle. The error rate, given in the last column, is the rate of flips between 0 and 1 in the sequence of successors.

VI. STOCHASTIC ANTIPERSISTENCE

All of these observations relied on the fact that the generator is on a cycle with well-known properties. It is thus interesting to ask how stable these results are if the sequence is not completely antipersistent. The simplest generalization is to introduce a probabilty p for changing the table entry/ exit when visiting a node: $p = 1$ reproduces the completely antipersistent walk; $p = 0$ is equivalent to using a constant (quenched) decision table, and $p = 1/2$ generates a completely random time series.

A first intuitive guess would be that even a small deviation from deterministic dynamics completely destroys all predictability: after all, on a path of length 2^{N+1}, there are on the average $(1-p)2^{N+1}$ occasions where the the sequence is continued persistently, thus leaving the cycle. Indeed, a single "error" is usually enough to move the system from one cycle to another; however, much of the local structure remains untouched. It turns out that the functions $s(N, N_{obs}, p)$ of prediction rates converge for large N (meaning roughly $N > 12$) to a set of curves that depend only on p and $N - N_{obs}$, which is displayed in Fig. 6.

The limit values for $p = 1$ have been explained above, and they are approached continuosly for $p \to 1$. For $p = 0.5$, the curves intersect at $s = 0.5$ – no prediction beyond guessing is possible. For small p, all curves converge to 1: the system is dominated by short loops in which only a small fraction of the possible states participate, and those are predicted with high accuracy.

Interestingly, between $p = 0.5$ and roughly $p = 0.85$, all shown curves are below 0.5, meaning that even observers with longer memory predict the sequence with less than 50% accuracy. I will give an analytical argument why this is the case for $N_{obs} = N+1$. An $N+1$-bit pattern ν is a combination of an N-bit pattern μ and one of its predecessors, let us say ${}^0\mu$, whereas the companion state $\hat{\nu}$ is a combination of μ and the other predecessor ${}^1\mu$. A visit to either ν or $\hat{\nu}$ switches the exit of μ with probability p. Consider two subsequent visits to ν, with

some number l of visits to $\hat{\nu}$ between them. The probability $s(p, N, N+1)$ of continuing with the same bit after these two visits is a sum of two probabilities: either the exit of μ was switched upon leaving ν the first time and then switched an odd number of times during the l visits to $\hat{\nu}$, or it was not switched the first time and switched an even number of times in between. Given p and the probability $\pi_l(p, N)$ of having l intermediate visits to $\hat{\nu}$, one then obtains by basic combinatorics

$$s(p, N, N+1) = \sum_{l=0}^{\infty} \frac{1}{2} \pi_l(p, N)[1 + (1-2p)^{l+1}]. \quad (3)$$

Unfortunately, $\pi_l(p, N)$ does not seem to be analytically accessible for general p. It can be measured in simulations, and the accuracy of Eq. (3) verified (see Fig. 6); also, for $p = 1/2$, since the system does a completely random walk on the graph, one gets the simple distribution $\pi_l(1/2, N) = 2^{-(l+1)}$. Assuming that this distribution does not change discontinuously near $p = 1/2$, Eq. (3) yields the approximation $s(1/2+\delta p, N, N+1) \approx 1/(2+2\delta p) \approx (1/2)(1-\delta p)$. This is obviously $< 1/2$ for $\delta p > 0$, i.e., $p > 1/2$.

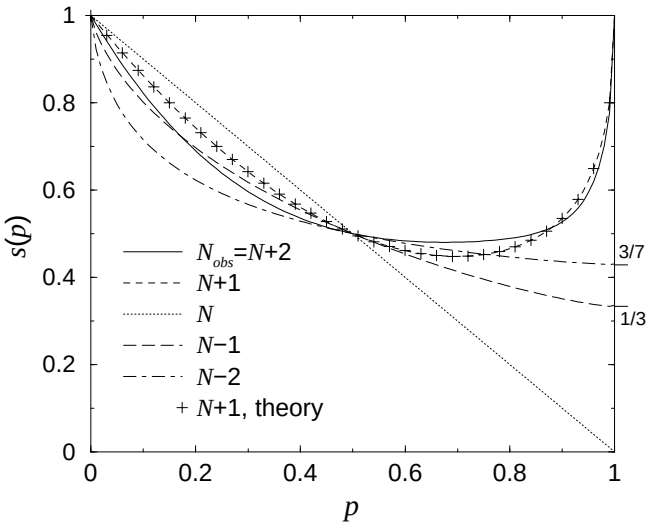


FIG. 6. Success rate of an observer keeping a table of the recent occurrence of N_{obs}-bit strings and the respective following bit, for $N = 16$. p is the probability of flipping the exit in the generating graph. The symbols labelled "theory" were calculated using Eq. (3) with approximate probabilities $\pi_l(p, N)$ taken from the simulation itself.

The error rate $1 - s(p, N, N_{obs})$ is a measure of the antipersistence of the time series on the scale N_{obs} – for $N_{obs} = N$, $1 - s$ is completely equivalent to the antipersistence parameter p of the underlying dynamics. However, even if $s(p, N, N+1)$ could be calculated, it would not be possible to recursively apply this function to find the antipersistence on the scales $N+2, N+3$ etc. In other words,

$$1 - s(p, N, N+1) \neq 1 - s(\{1 - s(p, N, N+1)\}, N+1, N+2). \quad (4)$$

It is thus not sufficient to give a single parameter p, or $1-s$, for some N in order to characterize the behaviour of a time series completely and to calculate its predictabilty on other scales of observation. The scale on which the dynamics work is important as well.

VII. RESULTS AND CONCLUSION

I introduced a deterministic algorithm that generates a binary time series that is completely antipersistent with respect to strings of length N. After a short transient, the algorithm runs into cycles of length $2 \cdot 2^N$, in which each string appears exactly twice. These cycles correspond to Hamiltonian paths on a DeBruijn graph of order $N+1$.

The cycle length is much larger than the typical cycle length of a graph with fixed decision tables. This seems typical for antipredictable sequences: sequences that can be predicted with 100% accuracy by some prediction algorithm usually do not require adaptation of the algorithm's parameters, whereas antipredictable sequences explore the combined phase space of the sequence and the generating algorithm, allowing for more, longer, and more complex cycles. In this case, however, the dynamics allow for fairly simple proofs of the properties of the cycles.

Observers that keep track of the most recent occurrence of N_{obs}-bit strings can predict the completely antipersistent cycle with 100% accuracy if $N_{obs} > N$, and with less than 50% success rate if $N_{obs} \leq N$. If the stochasticity is introduced by means of a probability p of flipping the exit edges, the success rate even of observers with $N_{obs} > N$ can drop below 50%, which shows that larger memory does not necessarily give better results. The rate of antipersistence on one scale is not sufficient to calculate the rate for other scales.

VIII. ACKNOWLEDGMENT

I am grateful for discussions with and helpful ideas from Hanan Rosemarin, Andreas Engel, Stephan Mertens, Ido Kanter, Michael Biehl, and Wolfgang Kinzel, and for financial support by the German-Israeli Foundation.

[1] F. Takens, in *Dynamical Systems and Turbulence*, edited by D. A. Rand and L. S. Young (Springer, Berlin, 1981), Chap. Detecting strange attractors in turbulence, pp. 365–381.

[2] T. Sauer, J. A. Yorke, and M. Casdagli, J. Stat. Phys. **65**, 579 (1991).
[3] A. Bunde and J. W. Kantelhardt, Physikalische Blätter **57**, 49 (2001).
[4] I. Kanter and D. A. Kessler, Phys. Rev. Lett. **74**, 4559 (1995).
[5] N. Johnson *et al.*, cond-mat/0105305 (unpublished).
[6] For further references on time series analysis in financial markets, consult the Econophysics homepage, http://www.unifr.ch/econophysics/.
[7] H. Kantz and T. Schreiber, *Nonlinear time series analysis* (Cambridge University Press, Cambridge, UK, 1997).
[8] M. Ausloos and K. Ivanova, cond-mat/0103367 (unpublished).
[9] M. Marsili and D. Challet, cond-mat/0004376 (unpublished).
[10] R. Savit, R. Manuca, and R. Riolo, Phys. Rev. Lett. **82**, 2203 (1999).
[11] G. Reents, R. Metzler, and W. Kinzel, cond-mat/0007351 (unpublished).
[12] W. Feller, *An Introduction to Probability Theory and its Applications* (John Wiley & Sons, New York, 1970), Vol. 1.
[13] H. Fredricksen, SIAM Review **24**, 195 (1982).
[14] C. Flye Sainte-Marie, L'Intermediaire des mathematiciens **1**, 107 (1894).
[15] H. Zhu and W. Kinzel, Neural Computation **10**, 2219 (1998).
[16] R. Metzler, W. Kinzel, L. Ein-Dor, and I. Kanter, Phys. Rev. E **63**, 056126 (2001).