Landscape dynamics of El Farol attendees

Canan Atilgan, Güven Demirel[†], Ali Rana Atilgan*

School of Engineering and Natural Sciences, Sabanci University, Tuzla 34956 Istanbul, Turkey
[†]*Department of Industrial Engineering, Bogazici University, Bebek 34342, Istanbul, Turkey*

*****Corresponding author:

atilgan@sabanciuniv.edu
telephone: +90 (216) 483 9525
telefax: +90 (216) 483 9550

Landscape dynamics of El Farol attendees

Abstract

Paradigm of the El Farol bar for modeling bounded rationality is undertaken. The memory horizon available to the agents and the selection criteria they utilize for the prediction algorithm are the two essential variables for agent strategies. The latter is enriched by including various rewarding schemes during decision making. Playing with the essential variables, one can manoeuvre the overall outcome between the comfort level and the endogenously identified limiting state. The distribution of algorithm clusters varies considerably for short memory. This affects the long-term aggregated dynamics of attendances. A transition occurs in the attendance distribution at the critical memory where the correlations of the attendance deviations take longer time to decay. A larger part of the crowd becomes more comfortable while the rest of the bar-goers still feel congestion for long memories. Introducing direct local interactions within the attendees by forming different types of networks, we create extremes in the attendance distributions. Delayed feeding of the data to the agents or their inclination of to absorb failure by insisting on unsuccessful algorithms introduces significant correlations, hence predictability of attendances. We additionally explore the extent of agents' manipulation, achieved by modulating the threshold in accordance with the correlations in the data.

Introduction

Data is ubiquitously available to each of us. We have been inductively forming our beliefs and perceptions, and adaptively building our expectation formation patterns by observing the past and present events taking place around us. The interplay between what is deduced from the aggregated average behavior and instantaneous forecasts of individuals based upon what is accessible to them shapes the evolutionary path of society's trend.

Inductive methodologies allow agents to construct their own decision making mechanisms from information allocated to all in common. Being an example of this approach, agent-based methodology adopts simple behavioral rules and allows a coordinated equilibrium to be an emergent property of the system, instead of modeling the system as if everyone's actions and beliefs were synchronized a priori with everyone else's (1, 2).

Arthur's El Farol bar problem brings up a paradigm (3). A total number of N players must decide independently whether to attend the bar or not. If a player forecasts that the total attendance will exceed the comfort level, L, she will not show up, otherwise she will go. We are interested in the way the players predict the total attendance and its long-term characteristics. The problem is based on the agents' knowledge of the overall attendance history, but the individual actions are not known. Yet, a collective behavior emerges from the system, whereby the average attendance is bounded on two sides by the threshold imposed from the outside, and the randomness that would take over the system in the absence of the threshold.

The problem has been investigated from many perspectives (4-7). In particular, Johnson and coworkers have explored the variation of volatility with the pool sizes available to the agents, as well as the effect of various selection schemes on attendance (8). Challet et al. have proposed a statistical mechanical model on the problem that becomes exact in the limit of randomness (9). It has also been analytically shown that the problem is mean reverting in nature; therefore, the average attendance will remain within the region $[N/2, L]$ (10). Yet, a controlled study that explores the effect of the parameters that may influence the output of the system has not been performed. Amongst these are (i) the different types of algorithms utilized by the agents, (ii) the strategy employed to select algorithms from this pool, and (iii) the memory horizon for which attendance data is available to the agents. In this study, we systematically study the effect of each of these parameters on the average attendance recorded. Probability distributions and the correlations of the absolute attendances are determined. We find that the emergent behavior in these systems is not only limited by an upper and lower bound, but also that it follows predictable patterns within these limitations once a critical value of short memory is surpassed.

Agent-based modeling of the El Farol bar

We use a system of N potential attendants. Each week we seek to find the number of people who will actually attend the bar, a_o. The attendance data of the last m weeks $[a_m, a_{m-1}, \ldots, a_1]$ are available to the agents. Each attendant, j, takes a decision to attend or not by making a prediction, q_j, on the possible value of a_o, using one of the many available algorithms. The algorithm pool here consists of (i) point-wise hypothesis – the agent uses the attendance data of the k^{th} previous week ($1 \leq k \leq m$) as his prediction, $q_j = a_k$; (ii) arithmetic average – the agent uses the average of the last k ($1 < k \leq m$) weeks as his prediction, $q_j = \frac{1}{k}\sum_{i=1,k} a_i$; (iii) weighted-average – the agent uses a weighted average of the last k ($1 < k \leq m$) weeks, where the more recent a week's data is, the larger weight it has, $q_j = \sum_{i=1,k} \frac{2ia_i}{k(k+1)}$; (iv) trend – the agent makes a least squares fit to the last k weeks' data ($1 < k \leq m$), and uses its extrapolation to the following week as his prediction, bounded by [0,100]. There are, therefore, a total of $4m - 3$ algorithms available, where m is the memory of the system. Once the agent makes his prediction q_j, he will not attend if the prediction is larger than a previously set comfort level, L.

The algorithm pool available to each agent affects the results. In one scenario, all algorithms in the pool may be available to all agents. However, note that alternatively, a sub-pool specific to each agent may be assigned *a priori*, and the agents may only choose from within this subset. Johnson and coworkers have shown that this choice significantly affects the volatility, where a minimum value of the volatility is obtained for a given fraction of the pool size, and it has a larger value for lower or higher fractions; maximum volatility is recorded for a fraction of one (8). The sub-pool may consist of a particular type of strategy (i) – (iv) above, or a mixture of those. The union of all the sub-pools should then define the total of the algorithms. The algorithms are initially assigned to the sub-pools using a uniform distribution. The extent to which the memory of each agent extends into the past is equivalent to m in all these cases. In what follows, one property we investigate in detail is how m affects the outcome. Since the fraction that gives the minimum volatility depends on m, we choose to make all the algorithms available to all the agents; i.e. the sub-pool fraction is one. Volatility is therefore at its maximum for the entire memory horizon.

Another property that significantly influences the results is the way the agents pick future algorithms from their sub-pool. This may be done in one of a multitude of ways: In the simplest of the applications, each agent hangs on to his current algorithm as long as it is successful, and changes it as soon as he fails in his prediction (denoted by ***scheme I*** here). Alternatively, various rewarding schemes that evaluate the success of the algorithms in retrospect may be adopted. In one scheme, the agent re-evaluates the predictions of all of his algorithms, and picks the one that would have been successful *and* provides the closest value to a_o (***scheme II***). In another scheme, the agent keeps a log of the success of his algorithms, by giving a point to all algorithms that would have succeeded at the current step. The agent then picks the one that has the highest cumulative score as his next predictor, or randomly selects between algorithms in case of equivalent scores (***scheme III***).

Obviously, we are not limited by the selection schemes listed above and employed in this study. There are other ways of implementing behavior at this level. One example is insistence, where the agent changes his current algorithm only after a predetermined number of successive failures. Rewarding may also be implemented in alternative ways, such as an intermediate case of schemes II and III, where the cumulative score is kept on only the best predictor, or by having a memory effect on the cumulative score, etc. As we shall see, such a choice significantly affects the fundamental properties of the system.

Effect of memory on mean attendance and distributions of active algorithms

In figure 1 we display the variation of the average attendance with increasing memory for the three schemes discussed above. We observe that there are upper and lower bounds on the attendance: As the amount of information available to the agents increases, the average attendance approaches a value that is equal to or less than the comfort level, L. The ultimate value attained, depends on the choice of the algorithm selection scheme, I, II, or III. The lower bound, on the other hand, tends to $N/2$; i.e. when there is very limited amount of information provided to the agents, their predictions become randomized, irrespective of the comfort level. The change from the lower bound, $N/2$, to the upper bound, occurs in an S-shaped curve; however, for short memories, depending on the selection scheme used, some of the data points deviate from this curve, as shown with the hollow circles in figure 1.

The reason for this deviation becomes clear if we investigate the fraction of algorithms utilized by the agents at different memories. This is shown in figure 2 for the three selection procedures, and the four types of algorithms that may be utilized by the agents. We find that in all cases, the distributions converge once the agents extend their memory horizon past a critical value. That memory value also corresponds to the point where the scatter in the data in figure 1 disappears, and the data track the curves marked by the solid line. For scheme I, since the algorithms are freely and readily changed from the total pool, that position is rapidly reached at a short memory of $m = 4$. For schemes II and III, however, the reaching of the corresponding position occurs at $m = 9$ and the scatter in the data at lower memory values is reflected into figure 1. We therefore mark the approximate expected trace of the data, had they followed the converged distributions in all memory values, by the dashed line.

Also interesting is the fact that the point-wise hypotheses (cycle detectors) are the most frequently used algorithms in all schemes. In scheme I where the algorithms are changed immediately in case they fail, they appear nearly twice as much as all the other types, indicating that they are the more often winners at instantaneous steps. That they make-up ⅔ of the algorithms utilized in scheme II, where the agents change to the predictor with the least error, points to the fact that they are also more precise predictors. Furthermore, in scheme III which uses cumulative rewarding, the competitive edge of cycle detectors gets compounded, and 90% of the algorithms utilized are composed of these. This over-expression of the point-wise hypotheses hints that there are cycles that occur in the data, an observation that may be corroborated by the correlations in the attendance profiles that emerge (see figure 4 in the next section). The remaining types behave almost equally well in scheme I, having values reduced slightly below that expected of random choice (23, 19, and 20% respectively for arithmetic average, weighted average, and trend, as opposed to the expected nearly 25% each); the reductions contribute almost equally to the enhancement in the cycle detectors. However, once rewarding is introduced in schemes II and III, weighted average algorithms are used less. In particular, in III their probability of appearance is nearly zero for all memories. Also, trend algorithms behave slightly better than arithmetic averages in II and III, opposite of that observed in scheme I.

In sum, for a given scheme, beyond a critical size of memory, the attendances fall onto an S-shaped curve. For shorter memories, attendances show large deviations from this curve. The critical memory where the average attendance begins to follow a predictable pattern corresponds to the converged distributions of the algorithm types used from the algorithm pool. The question remains, however, as to the origin of the bounded nature of attendance and its dependence on agent memory.

Effect of memory on attendance distributions and deviation clustering

In this section, we confine our attention to scheme I, and the threshold value $L = 60$. The probability distribution of the attendance for this set is displayed in figure 3. In the case of low memories, the probability distribution of attendance is widely spread over the attendance range exhibiting almost a double hump (blue line). One hump is centered on 70 with low variance and the second one is furnished with very high variance around 40. As the memory takes higher values, the lower end of the distribution moves towards the center and the hump at the lower end becomes more and more pronounced, while the higher end also shifts slightly towards the tail. When the memory reaches 60 (green line) the bump heights are equalized. Hereafter, with the higher memory values (dashed black line and black line are associated with the memory 90 and 150, respectively), the higher end distribution gets shallower loosing its maximum. Finally, the attendance distribution converges when the memory is around 200. The red line in figure 3 is obtained for the case where the memory is 210. The converged distribution is left only with one peak whose maximum is very close to the threshold value with a very thick tail of high attendances.

We note that the "transition" occurring in the attendance distribution resembles the phase transition phenomena frequently observed in nature. The point where the transition occurs in figure 3 (where the double humps are equalized) is also the memory level at which the correlations of the attendance deviations take longer time to decay zero. In figure 4, we present the correlations of the attendance deviations (at the top) and the correlations of the absolute attendance deviations (at the bottom). The most dominant correlations occur in consecutive weeks, a negative value which prevails for all memory levels, representing the immediate pull back to the mean. The correlations between the attendances of every other week, on the other hand, tend to be maximum at the point where the two humps in the distributions are equalized (i.e. the transition point, $m = 60$, in figure 3) and then disappears for very long memories. A positive correlation between the attendances of every third week, in contrast, exist even at $m = 210$.

Exhibited at the bottom of figure 4 are the correlations of the absolute attendance. While an examination of the correlations lends one to assume that correlations die, the correlations of the absolute attendance display decay in the first two weeks and then remains positive and significantly high. This is a manifestation of the intuition that the attendances tend to revert the mean. Yet, figure 4 signifies that "large changes tend to be followed by large changes, of either sign, and small changes tend to be followed by small changes." The structural issues pertinent to deviation clustering have been recently reviewed (11). Note that the non-diminishing correlations between the data is manifested in the success of the point-wise algorithms (figure 2).

Concluding Remarks

Whether the average attendance will converge to the externally provided comfort level or not depends on the algorithm selection procedures of the agents. Changing the algorithm used whenever it fails, irrespective of the past success of the algorithm, and picking up another one randomly (scheme I), drives the average attendance to the comfort level, L, as the agents use more information from the past. Taking into account a merit based stickiness to the algorithms employed in the past (schemes II and III), exhibit considerable deviation from the path that carries the average attendance to L. As shown in figure 1, stickiness not only alters the plateau levels, but also yields large fluctuations at the approach-to-plateau pathways, especially for shorter memory allocations (empty circles).

Information carrying capacity modulates the fractional use of the type of algorithms available in the pool. For short time horizons, it is rather difficult to estimate the distinct distributions. The algorithm clusters, namely, cycle detectors, trend followers, average takers, are shared by the agents sporadically, exhibiting a transient regime as observed in figure 2. As more of the past information is made available to all agents, independent of the rewarding scheme, fractional distributions are collectively balanced. Note, however, that reaching the detailed balance in the distributions does not coincide with reaching the plateau value of the attendance in figure 1, but is rather associated with the regions of small variations from the paths described by the S-shaped curves.

The mean-reverting nature of the process is characterized by the correlation of the attendances in figure 4. Metaphorically, the energy wells associated with different memory horizons of the process may be obtained by the negative logarithm of the attendance distributions. In figure 3, we showed that the attendance distribution experiences a phase transition when the areas of the wells are equalized when memory reaches 60. After the transition, a larger crowd feels the comfort, yet, still many attends without joy.

REFERENCES

1. *Adaptive Agents, Intelligence, and Emergent Human Organization: Capturing Complexity through Agent-Based Modeling* (2002) in *Proc. Natl. Acad. Sci. USA*, Vol 99 suppl. 3.
2. Tesfatsion L, Judd K, eds. *Handbook of Computational Economics*. Vol. 2. (2006).
3. Arthur WB (1994) *Amer. Econ. Rev.* 84: 406-411.
4. Arthur WB (1999) *Science* 284: 107-109.
5. Arthur WB (2006) in *Handbook of Computational Economics*, eds. L. Tesfatsion and K. Judd (North-Holland), ch. 32.
6. Challet D, Marsili M, Zhang Y-C (2005) *Minority Games: Interacting Agents in Financial Markets* (Oxford University Press).
7. Bell AM, Sethares, WA, Buckley, JA (2003) *IEEE Trans. Signal Process.* 51: 875-885.
8. Johnson NF, Jarvis S, Jonson R, Cheung P, Kwong YR, Hui PM (1998) *Physica A* 258: 230-236.
9. Challet D, Marsili M, Ottino G (2004) *Physica A* 332: 469-482.
10. Lus H, Aydin CO, Keten S, Unsal HI, Atilgan AR (2005) *Physica A* 346: 651-656.
11. Granger CWJ, Machina MJ (2006) *J Econometrics* 135: 15-29.

Figure Captions:

Figure 1. Average attendance versus memory of agents for N = 100 agents, and a comfort level of L = 60. Each data point is the result of 50 runs of 2000 weeks. The first 100 weeks of attendance data are not included in the averages to remove transient effects. Error bars are smaller than the data points. Dashed lines approximate the paths that would have been followed at short memories, had the algorithm types followed the converged distribution profiles (see figure 2 and the text for details). Solid lines at longer memories are drawn through the points to guide the eye.

Figure 2. The distributions reached by each of the algorithm types (point-wise, arithmetic average, weighted average, and trend) at different memory values. Note that there are $4m - 3$ algorithms available to the agents at each memory value, m. The sum of the values at each m is equal to one.

Figure 3. Probability distributions of the attendance for different memory horizons. Note the transition observed at the memory, $m = 60$ (green line). The fat tail does not disappear even in the longer memory horizons, indicating a significant amount of uncomfortable bar-goers.

Figure 4. The correlations of the attendance deviations (top) and the absolute attendance deviations (bottom). At the top figure, the correlations disappear in the second week in the attendance deviations. At the bottom, the correlations decay fast in the first two weeks and stay significant at later times.

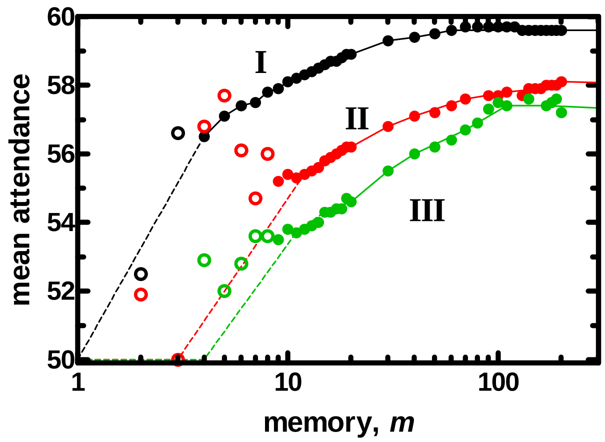

Figure 1

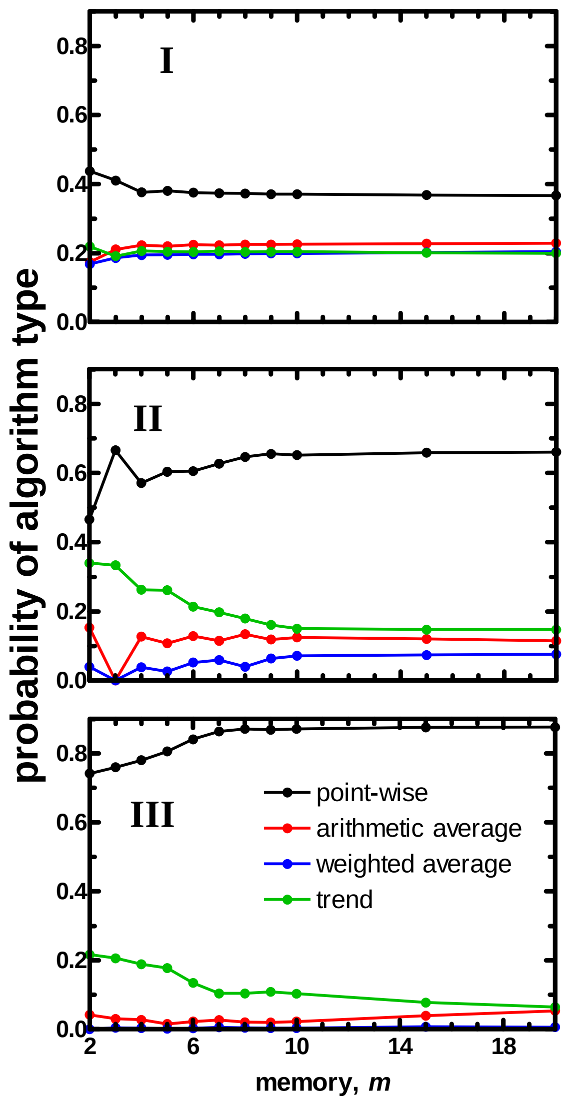

probability of algorithm type
I
II
III
0.0
0.2
0.4
0.6
0.8
point-wise
arithmetic average
weighted average
trend
2
6
10
14
18
memory, *m*

Figure 2

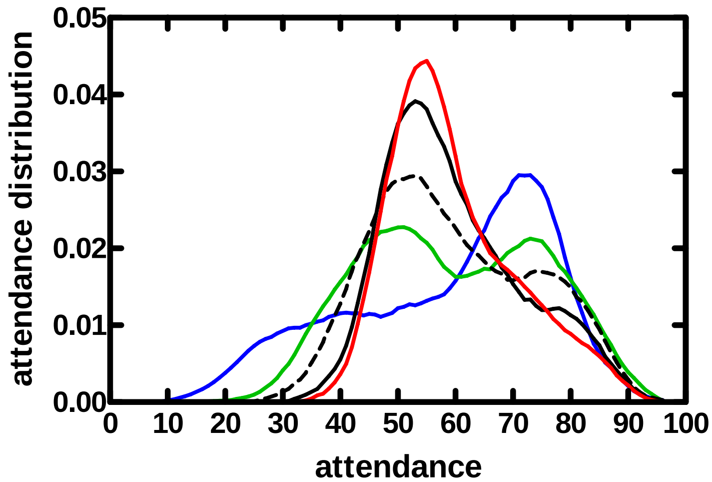

0.05
0.04
0.03
0.02
0.01
0.00
attendance distribution
0
10
20
30
40
50
60
70
80
90
100
attendance

Figure 3

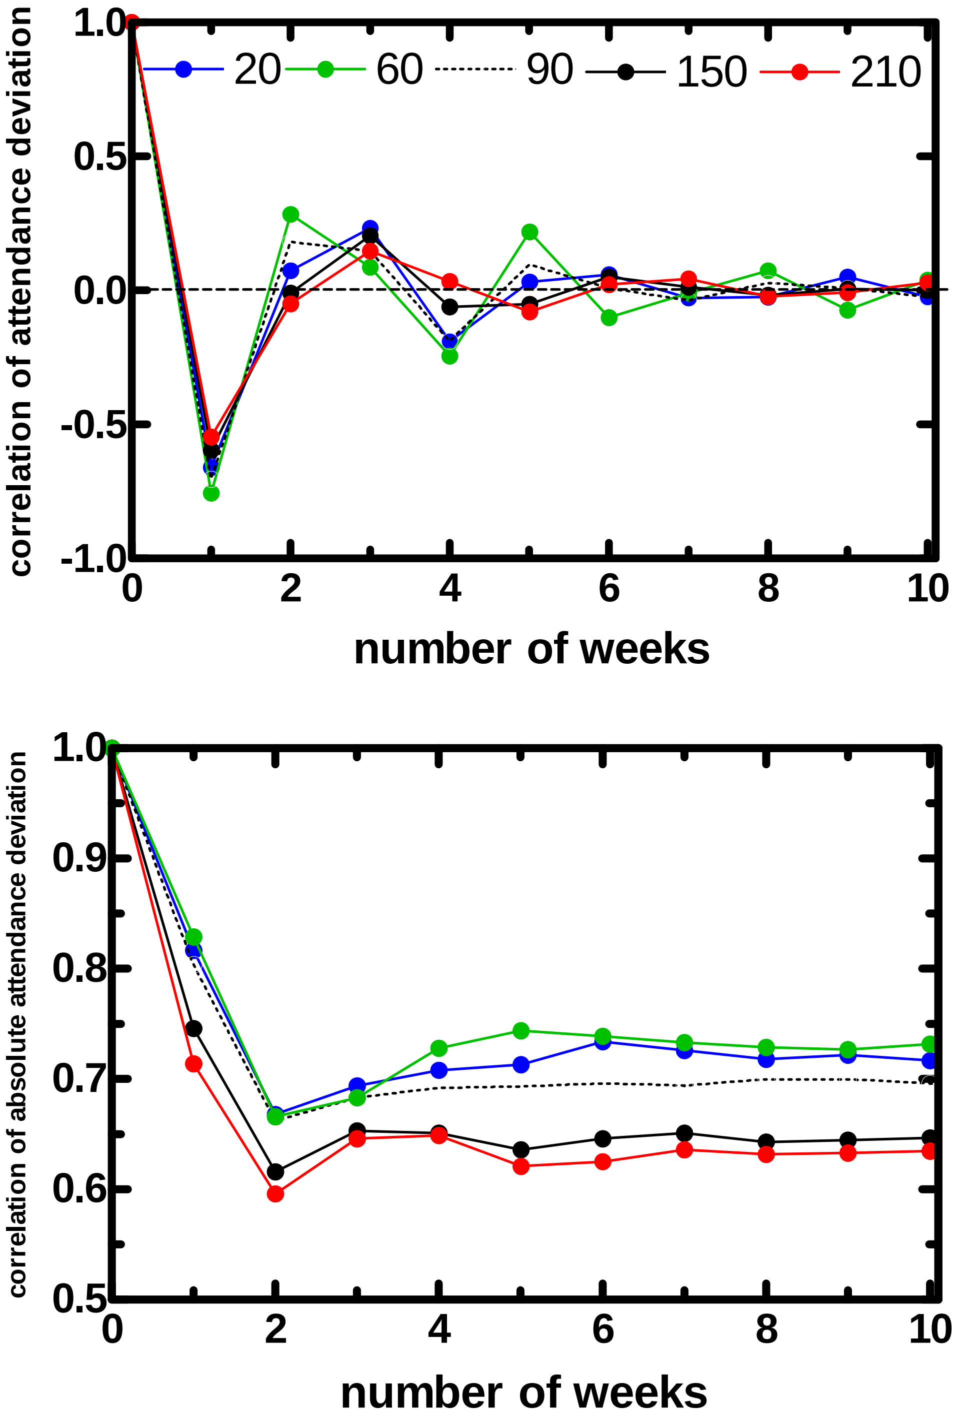
correlation of attendance deviation
1.0
0.5
0.0
-0.5
-1.0
20
60
90
150
210
0
2
4
6
8
10
number of weeks
correlation of absolute attendance deviation
1.0
0.9
0.8
0.7
0.6
0.5
0
2
4
6
8
10
number of weeks

Figure 4