Why Do Clocks Move Clockwise?

The Dynamics of Collective Learning

Vivek S Borkar

A mathematical model for collective learning by several autonomous agents is described. This is a stochastic recursion that arises in several disparate fields like statistics, engineering and economics.

A Little Learning is a Dangerous Thing

Why do clocks move clockwise? In other words, how did we learn to build and read clocks in one particular way without any collective decision to do so? The answer is not obvious, because 'counterclockwise' clocks did exist a few centuries ago.

There are other similar questions. How did we agree to drive on the right (i.e., left) side of the road? How did we learn various customs and conventions, social protocols, dressing codes, etc.? How do certain standards get adopted in industry? Why do people sometimes go overwhelmingly to one particular restaurant or for one particular brand of shoes rather than another comparable one?

The common feature in these phenomena is that through purely individual 'learning' by several autonomous agents, a common focus has emerged. To borrow a paradigm from economics, it is as though an invisible hand has ordained a particular outcome.

This, of course, does not imply that the resulting outcome is always the best possible. The two possible conventions for clocks or for driving are completely equivalent, but this need not be the case in social conventions or technological standards. Examples abound, my own favourite being the old British system of measurements my generation had to

Vivek S Borkar is with the Department of Computer Science and Automation, Indian Institute of Science. Disowned by purists of every hue because of his cross-disciplinary proclivities, he turns to *Resonance* for solace.

But whether good, bad or ugly, how was a particular behaviour learnt by these autonomous agents? The one line answer is: They 'urn'-ed it.

suffer through in school. Economics provides further examples. (Economists are quick to point out that their usage of the word 'equilibrium' does not imply approval of the same. As one of them puts it, a corpse is more at equilibrium than us, but we wouldn't trade places with it.)

But whether good, bad or ugly, how was a particular behaviour learnt by these autonomous agents? The one line answer is: They 'urn'-ed it.

Le rouge et le noir

A simple mathematical model for such phenomena is the nonlinear urn. One adds to an empty urn one ball at a time, either red or black. The probability of the $(n+1)$-st ball being red is $p(x_n)$, where x_n is the fraction of red balls at time n and $p : [0,1] \to [0,1]$ is a 'nice' (say, differentiable) function. If $y_n =$ the number of red balls at time n, $x_n = y_n/n$ and

$$y_{n+1} = y_n + \xi_{n+1}$$

where ξ_{n+1} is a $\{0,1\}$-valued random variable whose (conditional) probability of being 1 given the history up to time n is $p(x_n)$. Let $M_{n+1} = \xi_{n+1} - p(x_n), a_n = (n+1)^{-1}$. Then a little algebra leads to

$$x_{n+1} = x_n + a_n(p(x_n) - x_n) + a_n M_{n+1}, \tag{1}$$

where $\{a_n\}$ satisfies

$$\sum_n a_n = \infty, \quad \sum_n a_n^2 < \infty. \tag{2}$$

$\{M_n\}$ are uncorrelated random variables. In fact, the conditional expectation of M_{n+1} given the past up to time n is zero. Thus it acts like a 'noise' sequence added to the discrete iteration

$$x_{n+1} = x_n + a_n(p(x_n) - x_n). \tag{3}$$

We can view (3) as a discretization of the ordinary differential equation (o.d.e.)

$$\dot{x}(t) = p(x(t)) - x(t) \tag{4}$$

with decreasing stepsize $\{a_n\}$. Thus it approximates (3) better and better with time and may be expected to mimick the asymptotic behaviour of the latter. But we have the noise $\{M_n\}$ to contend with. The total noise contribution up to time n is $\sum_{k=1}^{n} a_k M_{k+1}$, whose variance is uniformly bounded by $4\sum_n a_n^2$ which is finite. Thus the total noise contribution till eternity remains bounded. To be precise, the series $\sum_n a_n M_{n+1}$ can be shown to converge with probability one. Hence its 'tail' $\sum_{m=n}^{\infty} a_m M_{m+1}$ goes to zero, making the noise contribution to the above approximation asymptotically negligible.

One colour gets an edge and then it feeds on itself till it completely dominates the population.

This can be made rigorous (see *Box 1*), with the conclusion that (1) and (4) have the same asymptotic behaviour with probability one. As a well-posed scalar o.d.e. with bounded trajectories, (4) must converge to a point. (Figure out why!) This point must be an equilibrium point where the right hand side of (4) is zero, i.e., $p(x) = x$. This equation has at least one solution (why?). *Figure 1* gives instances with one and three solutions resp. Thus (4) and therefore (1) must converge to one of these points with probability *one*, though one cannot say to which.

The way to tie this up with our opening remarks is to view the population share of a particular behaviour as the fraction of balls of a particular colour. In many interesting cases (such as clocks), the $p(\cdot)$ is such that after some initial randomness, one colour gets an edge and then it feeds on itself till it completely dominates the population (success breeds success, money attracts money, etc.).

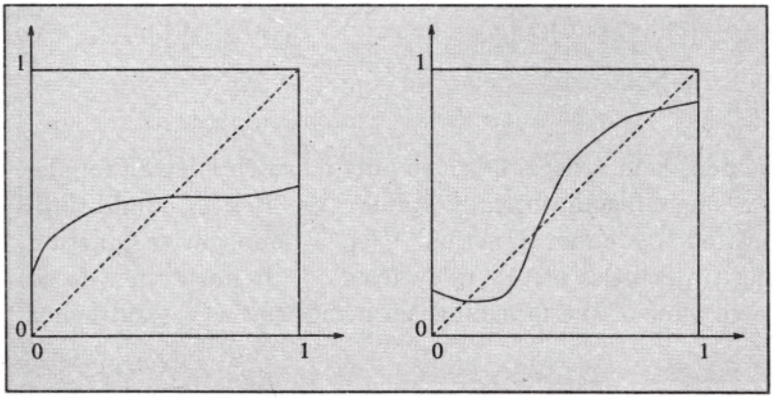

Figure 1.

Box 1. A Dicy, but Discrete Affair.

Einstein did not believe that God plays dice with the universe. We lesser mortals, however, often find it convenient to pretend that he does, because some natural processes are easier to model and analyse as random than otherwise. One such process is the 'martingale', a sequence of real-valued random variables $\{X_n\}$ such that the conditional average of X_{n+1} given 'past' up to n equals X_n with probability one. (Think of the net capital after n plays of a fair game, wherein one gains or loses nothing on average. The term 'martingale' itself originates from gambling.) Among other nice things these processes do, one is that they converge with probability *one* under suitable conditions (such as uniformly bounded absolute moments). Our 'net noise' process $\sum_{m=0}^{n-1} a_m M_{m+1}$ is one such, leading to the conclusion $\sum_{m=N}^{N+k} a_m M_{m+1} \to 0$ with probability one, i.e., noise input is 'asymptotically negligible'.

Changing tracks, consider the discretization of the o.d.e. (4) given by

$$x_{n+1} = x_n + a(p(x_n) - x_n).$$

Define $\overline{x}(t), t \geq 0$, by $\overline{x}(an) = x_n$, with linear interpolation in between. A standard application of the celebrated Gronwall inequality (see Coddington and Levinson, *Theory of Ordinary Differential Equations*, Tata McGraw Hill, 1955) leads to, for $T > 0$,

$$\max_{0 \leq t \leq T} ||\overline{x}(t) - x(t)|| \leq \kappa(T) O(a)$$

where $\kappa(T)$ is a constant depending on T and $O(a)$ stands for 'of the order of a'. (Thus the right hand side tends to zero as 'a' does.) In our problem, however, there are two differences: The decreasing (hence nonconstant) stepsize a_n and the 'noise'. Nevertheles, Gronwall comes to the rescue with an estimate of the type given below: Let $t(n) = \sum_{m=0}^{n} a_m$ and $\tilde{x}(t(n)) = x_n$ (as in (1)) with linear interpolation in between. Then for $T > 0$,

$$\max_{t(n) \leq t \leq t(n)+T} ||\tilde{x}(t) - x(t)|| \leq \kappa(T) O\left(a_n + \max_N \left|\sum_{m=n}^{n+N} a_m M_{m+1}\right|\right).$$

As before, the r.h.s $\to 0$ as $n \to \infty$.

The final leg of the argument is in general hard, so only a sketch is given for the special case when (4) has a single 'asymptotically stable' (see *Box 2*) equilibrium x^*. Then (4) has an associated 'Lyapunov function' $V(.)$, a nonnegative function that strictly decreases along trajectories of (4), away from x^*. It must then do so for $\tilde{x}(.)$ as well 'eventually' in view of the foregoing, ensuring that $\tilde{x}(.)$ (and hence $\{x_n\}$) also converge to x^*.

More generally one considers a vector iteration

$$X_{n+1} = X_n + a_n(h(X_n) + M_{n+1}) \tag{5}$$

with positive scalars $\{a_n\}$ satisfying (2), $h(\cdot)$ a 'nice' function and $\{M_n\}$ a 'noise' sequence as before. This arose in statistics in the early fifties and was dubbed the 'stochastic approximation algorithm'. It has since been used variously in statistics (nonlinear regression, ...) and electrical engineering (adaptive control and filtering). Its recent resurgence is as a paradigm for learning algorithms in artificial neural networks and as models of learning by economic agents. In these, the attractive feature of (5) has been its usually low resource requirement per iteration and its incremental nature – it makes only small changes at each time. For economists, the former captures 'bounded rationality' of the economic agents, the latter their inertia. For engineers, the former is an engineering reality, the latter a way to buy stability of the algorithm at the expense of its speed. *Box 2* displays some mathematical details about (5).

Box 2. Five Easy Pieces

Here are some further mathematical tidbits about (5). All hold under 'suitable conditions' that shall remain unmentioned.

- An equilibrium point is asymptotically stable if nearby trajectories remain nearby and converge to it. The algorithm converges to any such point with a strictly positive probability, whereas it avoids an unstable equilibrium with probability one.
- In higher dimensions, an equilibrium point is not the only possible attractor (or limiting set) for the o.d.e. associated with (5), therefore for (5). (In fact, these attractors can be quite 'strange'.) An attractor is said to be chain recurrent if it consists of chain recurrent points defined as follows: A point x is chain recurrent if for each $\epsilon > 0$ there is a finite chain of points $y_0 = x, y_1, y_2, \ldots y_n = x$ such that the trajectory starting at y_i ends up within ϵ of y_{i+1} for each i. (Intuitively, these are points that would be mistaken for periodic if exact measurement were not possible). With probability one, (5) converges to some chain recurrent set.

Box 2 continued...

- If (5) is implemented in a distributed manner, different processors compute different components of (5). Then only a few components may get updated at each step and there can be communication delays among processors. Worse, the processors may have different clocks (asynchronism) and may choose different stepsizes $\{a_n\}$. (These are very realistic conditions in economic models.) Nevertheless, one can show that the algorithm tracks the o.d.e.

$$\dot{x}(t) = Q(t)h(x(t))$$

where $Q(t)$ is a diagonal matrix with nonnegative diagonal entries that add to one. The latter can be viewed in a sense as relative frequencies of updates of the component in question.

- If some components use stepsizes $\{a_n\}$ and others use $\{b_n\}$ with $b_n/a_n \to 0$, the former move 'faster' than the latter. The limiting behaviour mimicks a 'singular o.d.e.' with two time scales. The faster component sees the slower one as almost static while the latter sees the former as essentially equilibrated. This can be used to advantage for algorithms that have two loops, the outer one requiring the near-convergence of the inner one for each iteration (e.g., algorithms that alternate between optimization and averaging).

- Though we summarily dismissed the noise, it does contribute fluctuations that can be analysed. While the o.d.e. gives the average or 'typical' behaviour, the fluctuations, on suitable rescaling, approximate a stochastic differential equation (s.d.e.) which then gives us useful information about the fluctuations. One may also add external noise to improve the algorithm, as in simulated annealing. If so, the original o.d.e. should be replaced by an appropriate s.d.e.

Games People Play

An important arena for learning models has been game theory. One considers a population of agents who interact ('play games') with each other, receiving payoffs as a function of their own and others' strategies. Based on the observed payoffs, each agent makes incremental changes in his own strategy.

The oldest of such models is 'fictitious play' for two person games, wherein each agent plays the best response to the current (time) average behaviour of the other. This fits the above framework and leads to an o.d.e.

$$\dot{x}(t) = f(x(t)) - x(t)$$

where $f(\cdot)$ is the so called 'best response' map.

After years of 'static' refinements of the equilibrium concept, the economists have moved to the dynamic models of disequilibrium.

More recently, a more popular paradigm has been the 'replicator dynamics'. Here $x(t) = [x_1(1), \ldots, x_d(t)]$ satisfies the o.d.e.

$$\dot{x}_i(t) = x_i(t)\left[\sum_j a_{ij}x_j(t) - \sum_{j,k} x_j(t)a_{jk}x_k(t)\right].$$

One views $x_i(t)$ as the population share of strategy i at time t and a_{ij}, the payoff on playing i if the adversary plays j. Thus the rate of increase of $x_i(t)$ is proportional to its current payoff $\sum_j a_{ij}x_j(t)$ minus the population average of the payoff $\sum_{j,k} x_j(t)a_{jk}x_k(t)$. This equation originates in evolutionary biology, where it models 'phylogenetic learning', i.e., the (passive) adaptation of species under selection pressure. Economists have adopted it as a model of 'ontogenetic learning', i.e., aggregate behaviour due to individual (active) learning by the agents and recovered it as a limiting case of appropriate models of individual learning.

There are other models in similar vein. The interest is due to a classic problem in economics. The economists have long accepted Nash equilibrium as being the natural equilibrium concept. This is an equilibrium where no agent can unilaterally improve his lot by a strategy change, all else remaining the same. The problem usually is that there are too many candidate equilibria. After years of 'static' refinements of the equilibrium concept, the economists have moved to the dynamic models of disequilibrium, viz., the aforementioned models of learning, hoping to narrow down the choice to those equilibria that arise as an asymptotically stable equilibrium for the same. A further refinement is to add noise to these dynamics and take the small noise limit to identify equilibria that are stable under stochastic perturbations.

Many controversies dissolve away once the problem is couched in mathematical terms.

And Now for Something Completely Different....

Let us now leave aside mathematics and engineering and look at the role of (5), or more generally of mathematical models, in social sciences. At best crude caricatures of reality, are they worth anything?

Let me start by quoting the noted economist Frank Hahn who has observed that many controversies dissolve away once the problem is couched in mathematical terms. Being forced to precisely tag the variables and lay down their hypothesised relationships, either the logical outcome becomes apparent, or an inherent contradiction in the premises gets exposed.

Another advantage is being able to tag specific qualitative phenomena which then become recognisable across the board, in different contexts and in different disciplines. One such instance is the phenomenon of 'locking in'. We have seen that (5) can get 'locked into' an equilibrium from which it will not budge. In economics, this leads to the phenomenon of 'increasing returns' wherein the more you invest, the more you gain, with a spiralling effect. This goes against the conventional wisdom of diminishing returns. Loosely speaking, the latter still applies to the traditional sectors of economy like agriculture, while the former comes to the fore in certain fast moving sectors of modern 'high tech' industry.

An important qualitative insight here is the following: A system locked into an undesirable equilibrium cannot just be nudged away from it, but will have to be forced out, because one is working against restoring forces pulling it back to that equilibrium. If this seems obvious to you, try applying it to your favourite social or technological ill and see if it still looks obvious.

Hahn goes on to suggest that the hostility to mathematisation is a 'sour grapes' syndrome displayed by the mathematically diffident. That is perhaps being a bit harsh, for there are genuine causes of concern. One is the overselling of models by overzealous supporters, who read more into them than what they really have to offer. (Just think of the various 'isms' floating around.) In particular, any quantitative inference must be treated with care.

Another problem is that of having the model display the behaviour you want it to, having subconsciously built it in. For models of learning, yet another potential pitfall is best described by the quote: One learns to itch only where one can scratch. Our models may do the same, leading to a false sense of complacency.

To conclude, my aim has been to give you a glimpse of the exciting interdisciplinary area of learning systems. Granting that they are imperfect models of imperfect systems, one does learn something from them. Exercised with caution, this little learning need not be a dangerous thing. With the intense ongoing activity in this area, one can look forward to an improved understanding of these dynamics in the coming years. Till that happens, we may share the sentiments of the great 20th century philosopher George Harrison, when he said:

> *With every mistake we must surely be learning*
> *still my guitar gently weeps*

Suggested Reading

- **A Benveniste, M Metivier and P Priouret.** ***Adaptive algorithms and stochastic approximations.*** **Springer Verlag. Berlin-Heidelberg, 1990.**
- **F Vega-Redondo.** ***Evolution, games and economic behaviour.*** **Oxford Univ. Press, 1996.**

Address for Correspondence
Vivek S Borkar
Department of Computer Science and Automation
Indian Institute of Science
Bangalore 560 012, India.

Ernst G Straus relates the following anecdote about Albert Einstein:

We had finished the preparation of a paper and we were looking for a paper clip. After opening a lot of drawers we finally found one which turned out to be too badly bent for use. So we were looking for a tool to straighten it. Opening a lot more drawers we came on a box of unused paper clips, Einstein immediately starting to shape one of them into a tool to straighten the bent one. When I asked him what he was doing, he said, "When I am set on a goal, it becomes difficult to deflect me."

The Mathematical Intelligencer. Vol.17.No.2., 1995