Branching processes, tipping points and phase transitions

David Aldous

October 26, 2017

2 types of model.

Some probability models of real-world phenomena are "quantitative" in the sense that we believe the numerical values output by the model will be approximately correct. At the other extreme, a **toy model** is a consciously over-simplified model of some real-world phenomenon that typically attempts to study the effect of only one or two of the factors involved while ignoring many complicating real-world factors. It is thus "qualitative" in the sense that we do not believe that numerical outputs will be accurate.

An example is the **Galton-Watson branching process** model I will describe. This is a textbook topic in STAT 150 (a first course in stochastic processes), but I want to emphasize the "just supercritical" formula (2). Another example from STAT 150 is the **M/M/1 queue** which I'll mention at the end.

This **Galton-Watson branching process** model is used as a toy model in many different settings. To have a concrete language, we talk about "individuals" and "offspring". To visualize individuals and offspring, you can either imagine asexual reproduction or look only at males or only at females in a two-sex species like humans.

The model is that there is a probability distribution $\mathbf{p} := (p_i, i = 0, 1, 2, \ldots)$ and that each individual in a generation has a random number of offspring in the next generation, this number being picked from $\mathbf{p}$ independently for different parents.

By default we assume the process starts with 1 individual in generation 0; so there is some random number $Z_n \geq 0$ of individuals in each generation $n = 0, 1, 2, 3, \ldots$. There are two logical possibilities for what might happen in the long run:

- either "extinction" meaning $Z_n = 0$ for all large n
- or "survival", meaning $Z_n \geq 1$ for all n.

One of the highlights of an undergraduate course in stochastic processes is the following theorem.
Write μ and σ for the mean and s.d. of the number of offspring.
Theorem. (a) If $\mu < 1$ then $\mathbb{P}(\text{extinction}) = 1$.
(b) If $\mu > 1$ then $\rho = \mathbb{P}(\text{extinction}) < 1$ and is the solution of the equation

$$\rho = \Phi(\rho) \tag{1}$$

where Φ is the probability generating function defined by

$$\Phi(z) = \sum_{i=0}^{\infty} p_i z^i.$$

Keep in mind that the "independence" assumptions are tantamount to assuming there is no "interaction" between individuals and that there are no external constraints on population size – both assumptions are unrealistic in almost all imaginable real-world contexts.

I won't repeat the textbook derivation of the Theorem, but I will derive an interesting approximate formula for a particular setting. The cases $\mu < 1$, $\mu = 1$, $\mu > 1$ are called *subcritical, critical, supercritical*. I want to consider the "just supercritical" case where $\mu > 1$ but $\mu - 1$ is small.

$$\text{For a just supercritical Galton-Watson process, } \mathbb{P}(\text{survival}) \approx \tfrac{2(\mu-1)}{\sigma^2}. \tag{2}$$

[do calculation on board]

This is often not mentioned in textbooks, so let me give
Derivation of formula (2). Textbook facts about the probability generating function for the random number X of offspring are

$$\Phi(1) = 1, \quad \Phi'(1) = \mu, \quad \Phi''(1) = \mathbb{E}[X(X-1)] = \sigma^2 + \mu^2 - \mu \approx \sigma^2$$

the approximation holding because $\mu \approx 1$.
We want the survival probability $\bar{\rho} = 1 - \rho$. The equation in the Theorem, $\rho = \Phi(\rho)$, can be rewritten in terms of $\bar{\rho}$ as $h(1-\bar{\rho}) = 0$, where $h(x) = \Phi(x) - x$. Consider the series expansion: for small x,

$$h(1-x) \approx h(1) - xh'(1) + \frac{1}{2}x^2 h''(1).$$

Since $h(1) = 0, h'(1) = \mu - 1, h''(1) \approx \sigma^2$ the rewritten equation becomes

$$0 \approx -\bar{\rho}(\mu - 1) + \frac{1}{2}\bar{\rho}^2\sigma^2$$

and solving for $\bar{\rho}$ gives the stated formula (2).

Keep in mind that the "independence" assumptions are tantamount to assuming there is no "interaction" between individuals and that there are no external constraints on population size – both assumptions are unrealistic in almost all imaginable real-world contexts. This is why I call it a "toy model".

Let's think of a toy model for the spread of epidemics such as influenza. Each infected person will infect some random number of other people; the mean such number is called the **reproduction number** μ. We can use the previous Galton-Watson process to model the number of cases in the **initial phase**; if $\mu < 1$ the epidemic will not occur; if $\mu > 1$ and there are (at least) several initial cases then there will be an epidemic.
Once the epidemic grows it is natural to work with

$$g(t) = \textbf{ proportion } \text{of population infected.}$$

If we ignore the fact that people recover, and if we assume that all pairs of people are equally likely to have contact, then the rate of growth of the epidemic is roughly proportion to the number of infected-uninfected contacts, and this is most simply modeled by the

logistic equation

$$g'(t) = cg(t)(1 - g(t)).$$

The solution, up to an arbitrary time-shift, is

$$g(t) = \frac{1}{1 + e^{-ct}}.$$

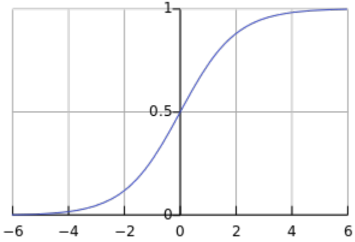

Reality check. In that toy model, 100% of population is eventually infected. In fact, in the annual "seasonal influenza" epidemic in the U.S., typically the percentage of population infected is in the range 5% - 20%.

So what's wrong with the toy model? Many things, in particular

- people recover (infective for about 7 days)

and the population is not homogeneous:

- different levels of partial immunity
- different people have different numbers of inter-personal contacts
- spatial locations matter.

These all affect the eventual proportion of population infected, but the S-shaped curve remains typical.

[next slide]

FluSurv-NET* Laboratory-Confirmed Cumulative Hospitalization Rates (per 100,000), 2012-13 Season

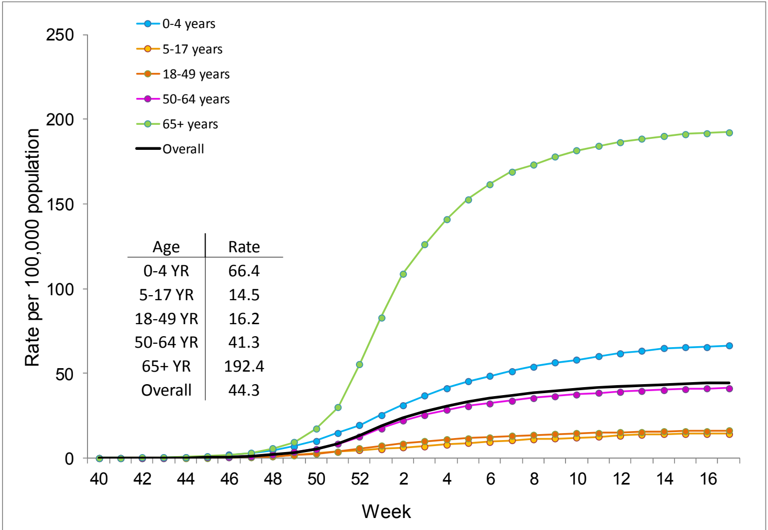

*FluSurv-NET results represent surveillance in the 10 Emerging Infections Program (EIP) states (CA, CO, CT, GA, MD, MN, NM, NY, OR, and TN) and five additional states (IA, MI, OH, RI, and UT)

Bottom line. Regardless of details, models of the spread of a "feature" – epidemics or technology or opinions etc in human society – starting from a few individuals mostly have similar behavior: either the feature dies out quickly, or it starts growing exponentially until reaching some proportion of the population. The exponential growth rate is just

$$\mu = \text{ number of new individuals who obtain the feature from a typical individual }.$$

Tipping points

The phrase “tipping point” was popularized by Malcolm Gladwell’s 2000 book *The Tipping Point: How Little Things Can Make a Big Difference*. [show Google Ngram]. Let me start by asking

- How do people actually use this phrase as a metaphor?
- What is the underlying physics analogy?

I will then show how this phenomenon arises in toy probability models.

In previous years one could use Google to search blogs and here are 8 examples found for the 2014 course.

A Tipping Point for Too Much Talent. August 27, 2014
Can a sports team ever have too much talent?
Those of us following the trades and roster jockeying by National Football League, Premier League and National Basketball Association teams could reasonably assume that the answer is no. But a new study of hundreds of games in several professional sports leagues suggests that, in fact, talent does have a **tipping point**, beyond which too many great players become detrimental to a teams success, a finding with broad implications for coaches at all levels of play, as well as fans and athletes possessing transcendent and more-average gifts.

Ebola tops EU meeting with warnings crisis is at "tipping point".
Oct. 20, 2014
European foreign ministers gather in Luxembourg Monday to try and formalise a joint EU response to combat the Ebola virus amid diplomatic warnings the crisis has reached a "**tipping point**".

The 3D printing Tipping Point: Quality up, Cost down. 2 Nov 2014
If 3D printing wasn't on your radar, it certainly should be now. Innovators are capitalizing on new ways to use the technology, creating 3D-printed homes and cars that may be faster and cheaper as the technology continues to improve. Have we reached the **tipping point** where 3D printing quality is going up and its cost is going down?

Apple Pay Is Here: A Tipping Point? Oct. 21, 2014
Apple Pay is now available and with the success of the iPhone 6 there is a real possibility that the payment system will really change. There is a lot going on in the area of "unbundled" banking opportunities and it is really starting to draw in "money" players. Forget what the "old" banking system looks like...think of how an "unbundled" banking system might look like.

Is 2014 the “Tipping Point” for the GMO Labeling Movement?
September 29, 2014
In Oregon and elsewhere, voter sentiments are trending more in favor of consumers’ “right to know.”

Net Neutrality in the U.S. Reaches a Tipping Point. Oct 31 2014
We’ve spent years working to advance net neutrality all around the world. This year, net neutrality in the United States became a core focus of ours because of the major U.S. court decision striking down the existing Federal Communications Commission (FCC) rules. The pressure for change in the U.S. has continued to grow, fueled by a large coalition of public interest organizations, including Mozilla, and by the voices of millions of individual Americans.

Big Cats at a Tipping Point in the Wild. August 7, 2014
With lions, leopards, and other big cat species on a downward spiral, we sit at a **tipping point** when it comes to the conservation of some of the worlds most iconic animals.

Is Syria on the Verge of a Tipping Point? September 8, 2014
The Syrian civil war, which has dragged on for three years, has until now been deadlocked in a bloody war of attrition. However, the forceful emergence of ISIS, renamed the Islamic State (IS), may ironically have opened the door for a change in the conflict.

Is there any unifying concept behind these stories? Wikipedia hasn't figured this out . . . . . .
[show Tipping point (disambiguation)]
The metaphor of **tipping** suggests a rapid switch between two qualitatively very different states. The online Marriam-Webster dictionary says

> *the critical point in a situation, process, or system beyond which a significant and often unstoppable effect or change takes place*

To me there is a subtle distinction between two contexts.

- tipping point as something within a specific observed process causing a change from *before* to *after*
- Change in "parameters" – the background setting – which affects the process being observed.

[board: the previous examples]

In the context of democratic politics a steady shift in public opinion can result at some point in rapid changes due to different governments and laws. This really is like tipping, though is a consequence of human rules (like sports rules) rather unlike any physical phenomenon.

Microprocessor Transistor Counts 1971-2011 & Moore's Law

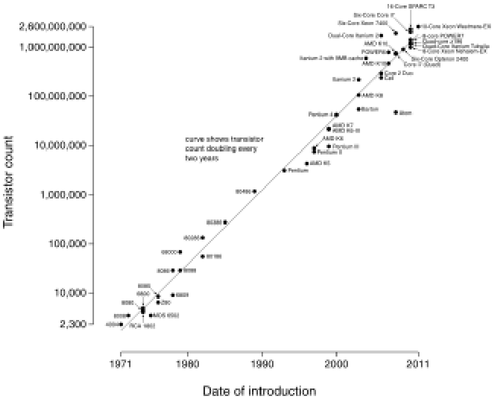

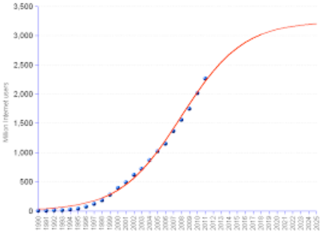
3,500
3,000
2,500
2,000
1,500
1,000
500
0
Million Internet users

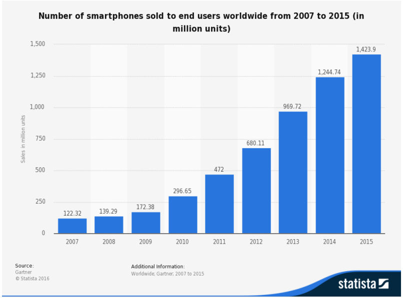
Number of smartphones sold to end users worldwide from 2007 to 2015 (in million units)
1,500
1,250
1,000
750
500
250
0
Sales in million units
122.32
139.29
172.38
296.65
472
680.11
969.72
1,244.74
1,423.9
2007
2008
2009
2010
2011
2012
2013
2014
2015
Source:
Gartner
© Statista 2016
Additional Information:
Worldwide; Gartner; 2007 to 2015
statista

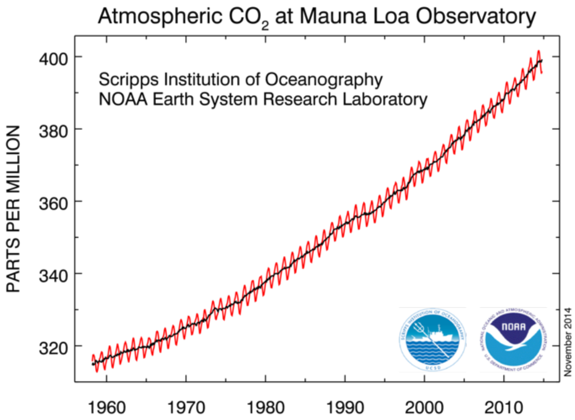

Atmospheric CO_2 at Mauna Loa Observatory
Scripps Institution of Oceanography
NOAA Earth System Research Laboratory
PARTS PER MILLION
400
380
360
340
320
1960
1970
1980
1990
2000
2010
November 2014

Digression: the word **parameter**.

In undergrad Statistics we often use the word in the context of fitting data to a distribution – Normal(μ, σ^2) distribution for heights. In serious scientific modeling, a parameter is better regarded as something relevant in the "background setting" which is measurable. Examples in this lecture:

- mean number of offspring
- mean number of individuals infected with flu by one individual
- cost/quality of 3D printing
- percentage of voters with a given opinion
- microprocessor capacity
- atmospheric CO_2
- temperature or pressure
- quantity/dryness of combustible material in a forest
- arrival rate of customers

Phase transitions. To the question "what is H_2O?" the simplest answer is "water". We all know water exists in three forms – ice, liquid water, steam – but we intuitively perceive liquid water as the "normal" form. That is, we view ice as "frozen water" rather than viewing water as "melted ice". But to Physics each form is equally natural; the particular form depends on the values of several parameters, primarily temperature, but also pressure and presence of impurities. *Phase transition* refers to the fact that the transition between forms occurs at very specific temperatures (0° and 100° Celsius).

Boiling water or making ice cubes are everyday instances of observable phase transitions, but for our purposes they are atypical in that we can easily and directly control the relevant parameter (temperature).

Note also that "phase transition" refers to the general phenomenon rather than to a particular ice cube. To illustrate the theme of this lecture let me move on to another example.

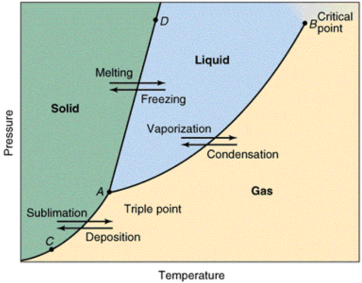
D
Critical
B
point
Liquid
Melting
Freezing
Solid
Pressure
Vaporization
Condensation
A
Gas
Triple point
Sublimation
Deposition
C
Temperature

Untended campfires. Imagine starting a campfire in the woods and then leaving, with no subsequent human intervention. What might happen? Well, two possibilities come to mind. The campfire might die out or it might start a large forest fire. Which happens depends on a lot of factors – temperature and humidity and wind, amount of combustible material near the campfire, etc – which again we call *parameters*.

Now imagine there are quite a few irresponsible people leaving campfires; some will cause forest fires and others won't, depending on the parameters in the particular instance. Mathematical scientists envisage the phase transition abstractly as the boundary between the parameter values which make a forest fire unlikely and the values which make it likely. More concretely, in a climate like California with mild wet winters and hot dry summers there will be a roughly predictable time of year when the forest fire result becomes likely, and this time of year is the "tipping point" for forest fires.

Here are two key points to note in this example. First, the "tipping point" refers to the general phenomenon of forest fires, not to a particular fire. Second, the relevant parameters have very different natures. Some are outside human control (weather), others are part of human social/economic life (laws against starting fires; the Forestry Department clearing undergrowth to help prevent forest fires), while others are very specific to the particular instance (combustible material nearby).

Just as **tipping point** is now widely used in non-technical writing, so is **phase transition** now used throughout the mathematical sciences with some broader meaning. Let me try to describe the central idea, which I will call **Phase transition tipping point** (PTTP).

- A system which fluctuates but (in what we perceive as an existing or "normal" situation) is long-term stable (dynamic equilibrium).
- The way it fluctuates is governed by some parameter; with some other parameter values the system would be unstable or have a qualitatively different equilibrium.
- There are slow changes in parameter (caused by external factors).
- The PTTP occurs when the parameter crosses the boundary between one stable equilibrium and another.

Epidemics are a basic example – is μ greater than or less than 1? Here is another basic "probability model" example.

Queueing theory is one of the major classical topics in applied probability, included in a typical first course on stochastic processes. Curiously, the "phase transition" analogy is not emphasized in textbooks.

A *queue* is exemplified by a store with one person (the *server*) at the checkout stand. There is a certain demand for checkout service; hypothetically, if there are 25 customers in a certain hour and on average they each require 2 minutes service time, then the server will be busy for 50 minutes in that hour, which we interpret as a *traffic intensity* of $50/60 = 5/6$. If the traffic intensity increased to $8/6$, then after an hour we would expect to see (if nothing else changed) a waiting line in which the last customer would have to wait around 20 minutes.

Of course in practice this rarely happens, for two obvious reasons. Potential customers seeing a long line might well forego shopping there; and to avoid that possibility, store management will ensure that an extra server will be available when needed.

This simple story illustrates two points. The mathematical point is that when the traffic intensity (conventionally denoted ρ) is greater than 1 the server cannot possibly keep up, and the wait line would grow longer and longer; but when $\rho < 1$ then the server must be idle (no customers) for a proportion $1 - \rho$ of the time, implying that the waiting line must become empty fairly frequently.

These two qualitatively different behaviors of a system are the sign of a PTTP, and in a case like this with only a single parameter, the PTTP $\rho = 1$ is called the *critical value*. The second point is that variations in the traffic intensity are caused by features of the broad human social world – which affects the times of day and of week we seek to shop – external to the "system" within the store.

Mean customer wait time. The qualitative phase transition above does not depend on explicit modeling of the queue, but any more quantitative result does, and a natural quantity to consider is the mean customer wait time before being served, in the case $\rho < 1$. You can see intuitively why the details matter by considering extreme cases. If customers were scheduled by appointment then one could arrange that waiting was never necessary; whereas for a jumbo jet load of passengers arriving simultaneously at immigration, there will inevitably be a wait line.

So this is where probability enters the story – we need some probability model. The simplest model is called the *M/M/1 model* whose essential features are that arrival times are assumed to be as a Poisson process ("purely random") and that service times are assumed to have Exponential distribution.

Within this model, a formula whose derivation can readily be found in textbooks is

$$\text{mean number customers in front of you} = \tfrac{\rho}{1-\rho} \tag{3}$$

from which it follows that

$$\text{mean wait time until your service starts} = \tfrac{\rho}{1-\rho} \times A \tag{4}$$

where A is the mean service time per customer.

Let me repeat that the validity of such formulas depends on the precise assumptions, which would typically be unrealistic! Part of the technical side of queueing theory is to find analogous formulas in other, hopefully more realistic, models.

Should you feel guilty about delaying minor tasks? Here is a thought-provoking consequence of formula (4). Suppose your list of tasks you need/want to do per week is longer than time available per week. Rational plan: first decide to allocate a fixed length of time per week to doing these tasks, then prioritize tasks and put on your "to do" list only those important enough – cut off so the average time needed per week equals the time you have allocated.

This scheme sounds good, but is ruined by randomness. Your list of tasks will be a queue with $\rho = 1$, so formula (4) says the mean waiting time until a typical task is completed will become infinite!

[discuss further]

Let us look at branching processes and toy models of epidemics from this viewpoint. Recall the **Galton-Watson branching process** model, usually stated in terms of "individuals" and "offspring".

Write μ and σ for the mean and s.d. of the number of offspring. There is a PTTP at critical value $\mu = 1$, in the following sense.
If $\mu < 1$ then $\mathbb{P}(\text{survival}) = 0$.
If $\mu > 1$ then $\mathbb{P}(\text{survival}) > 0$.

There is a conceptually subtle point here; we are making a statement about the process, not about a realization of the process.

Scientists have thought extensively about potential tipping points involving climate change, although details are uncertain.

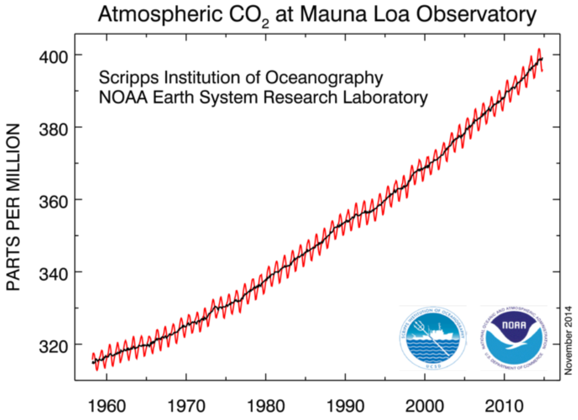

[show Wikipedia and paper]