ETAS-trophic Failures: Fit, Classification, & Forecasting

Philip B. Stark

Department of Statistics
University of California, Berkeley

Geller-Fest Workshop
Big Data in Geosciences: From Earthquake Swarms to Consequences of Slab Dynamics
25–27 May 2017

"Chances are 50/50 that it's a total crapshoot, but it could go either way."

Earthquake poker: metaphor not physics

- Earthquakes occur "as if" in a casino game whose rules are embodied in some mathematical model known to the seismologist.
- Like saying that there is a deck of seismology cards.
  - The deck contains some blank cards and some numbered cards.
  - In a given region, in every time interval, a card is dealt from the deck.
  - If card is blank, no earthquake.
  - If card has a number on it, there's an event with that magnitude.

What's the game?

Different models make different assumptions about # cards of each type there are, the shuffling, whether drawn cards are returned to the deck, etc.

- One extreme: big earthquake cards distributed fairly evenly (the characteristic earthquake model).
- Another extreme: cards thoroughly shuffled; after each draw the card is replaced and the deck is re-shuffled (tantamount to the Poisson model).
- In between: deck shuffled less than thoroughly (e.g., high cards tend to be followed by low cards—aftershocks), cards are not replaced (modeling stress accumulation or stress release), deck not re-shuffled between draws.
- ETAS: Draw w/ replacement. #s distributed according to GR, but if draw card w/ #, extra draws sooner; draw freq depends on recency of drawing #d card.

What's the game?

Different models make different assumptions about # cards of each type there are, the shuffling, whether drawn cards are returned to the deck, etc.

- One extreme: big earthquake cards distributed fairly evenly (the characteristic earthquake model).
- Another extreme: cards thoroughly shuffled; after each draw the card is replaced and the deck is re-shuffled (tantamount to the Poisson model).
- In between: deck shuffled less than thoroughly (e.g., high cards tend to be followed by low cards—aftershocks), cards are not replaced (modeling stress accumulation or stress release), deck not re-shuffled between draws.
- ETAS: Draw w/ replacement. #s distributed according to GR, but if draw card w/ #, extra draws sooner; draw freq depends on recency of drawing #d card.

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

ETAS

Combines Poisson, Gutenberg-Richter Law, Omori Law

- Earthquakes occur at random
- Poisson background; heterogeneous in space, homogeneous in time
- Magnitudes IID G-R
- After an event, chance of children follows Omori law
- Omori intensity depends exponentially on parent's magnitude
- Magnitudes of children IID G-R
- Children can have children (but events have 0 or 1 parents)
- Example of *linear marked Hawkes process*

Earthquake probabilities

Why do we think earthquakes have probabilities?

- Standard argument:
  $M = 8$ events happen about once a century in locale.
  Therefore, the chance is about 1% per year.
- But rates are not probabilities.
- Probabilities imply rates in repeated random trials; can estimate the probability from empirical rate.
- Having an empirical rate doesn't make something random.

Earthquake probabilities

Why do we think earthquakes have probabilities?

- Standard argument:
  $M = 8$ events happen about once a century in locale.
  Therefore, the chance is about 1% per year.
- But rates are not probabilities.
- Probabilities imply rates in repeated random trials; can estimate the probability from empirical rate.
- Having an empirical rate doesn't make something random.

Earthquake probabilities

Why do we think earthquakes have probabilities?

- Standard argument:
  $M = 8$ events happen about once a century in locale.
  Therefore, the chance is about 1% per year.
- But rates are not probabilities.
- Probabilities imply rates in repeated random trials; can estimate the probability from empirical rate.
- Having an empirical rate doesn't make something random.

Earthquake probabilities

Why do we think earthquakes have probabilities?

- Standard argument:
  $M = 8$ events happen about once a century in locale.
  Therefore, the chance is about 1% per year.
- But rates are not probabilities.
- Probabilities imply rates in repeated random trials; can estimate the probability from empirical rate.
- Having an empirical rate doesn't make something random.

Thought experiment 1

Chance of death

You are in a group of 100 people.

One person in the group will die in a year.

What's the chance you will die in the next year?

Thought experiment 2

Chance of name

You are in a group of 100 people.

One person in the group is named "Philip."

What's the chance your name is "Philip?"

What's the difference?

Ignorance does not create chance

- If the mechanism for deciding who will die is to pick one at random and shoot him/her, then the chance is indeed 1%.
- But if the mechanism is to shoot the tallest person, there's no "chance" here: You are the tallest, or not.
- *Ludic fallacy* (Taleb)

Rabbits

The Rabbit Axioms

1. For the number of rabbits in a closed system to increase, the system must contain at least two rabbits.
2. No negative rabbits.

Rabbits contd.

Freedman's Rabbit-Hat Theorem

You cannot pull a rabbit from a hat unless at least one rabbit has previously been placed in the hat.

Corollary

You cannot "borrow" a rabbit from an empty hat, even with a binding promise to return the rabbit later.

Rabbits contd.

Freedman's Rabbit-Hat Theorem

You cannot pull a rabbit from a hat unless at least one rabbit has previously been placed in the hat.

Corollary

You cannot "borrow" a rabbit from an empty hat, even with a binding promise to return the rabbit later.

Applications

- A rate is not a probability estimate unless the rate comes from a random phenomenon.
  (Gotta put a rabbit in the hat.)
- Testing whether the process appears to be random using the *assumption* that it is random cannot prove that it is random.
  (Can't borrow a rabbit from an empty hat.)

Applications

- A rate is not a probability estimate unless the rate comes from a random phenomenon.
  (Gotta put a rabbit in the hat.)
- Testing whether the process appears to be random using the *assumption* that it is random cannot prove that it is random.
  (Can't borrow a rabbit from an empty hat.)

What would make the casino metaphor apt/useful?

1. Physics of earthquakes might be stochastic
2. Stochastic models might provide a compact, accurate description of earthquake phenomenology
3. Stochastic models might be useful for predicting future seismicity

B. Luen, 2010. PhD Dissertation, UC Berkeley

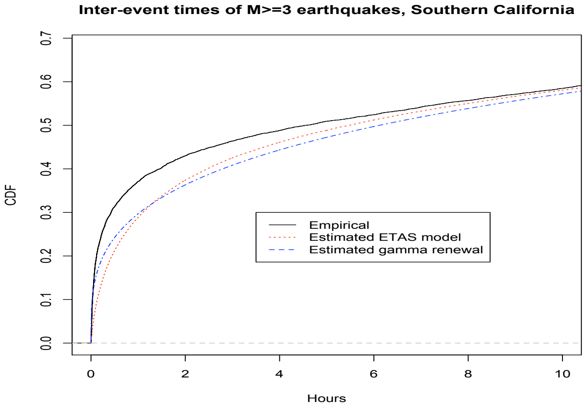

Figure 4.1: Cumulative distribution functions of inter-events times attached. The empirical inter-event distribution (SCEC catalog of Southern Californian $M \geq 3$ earthquakes, 1984-2004, $n = 6958$) is significantly different from both the fitted ETAS and gamma renewal models (in both cases, the P-value is less than 0.00001 for a test using the Kolmogorov-Smirnov test statistic). Empirically, there are more inter-event times under 2 hours than either fitted model would suggest. Beyond 12 hours, the difference in empirical distributions is small (not pictured).

ETAS Parameter estimates & simulations

- Often unphysical for real data (each event expected to have infinitely many children)
- Simulations can have burn-in times of order 10^5y

Sornette & Utkin, 2009. *Phys. Rev. E 79*, 061110

- Simulate ETAS seismicity
- Use ETAS to classify event as background or child (aftershock)

  - *Rather unreliable . . .*
  - *estimated rates of exogenous events suffer from large errors . . .*
  - *branching ratio . . . badly estimated in general*
  - *high level of randomness together with the long memory makes the stochastic reconstruction of trees of ancestry and the estimation of the key parameters perhaps intrinsically unreliable*

Prediction: Automatic Alarms, MDA, & ETAS

- Automatic alarm: after every event with $M > \mu$, start an alarm of duration τ

  No free parameters

- Magnitude-dependent automatic alarm (MDA): after every event with $M > \mu$, start an alarm of duration τu^M

  1 free parameter: u

For both, adjust fraction of time covered by alarms through τ.

- Optimal ETAS predictor: level set of conditional intensity

  4 free parameters: K, α, c, p

ETAS v Auto Alarms: Simulations (Luen, 2010)

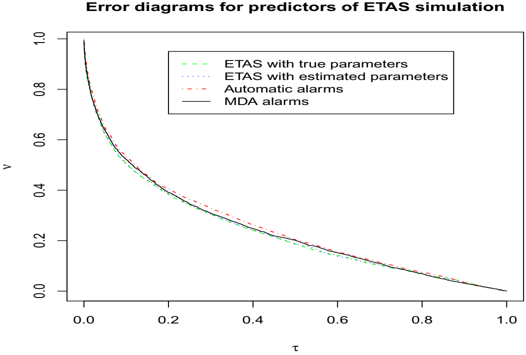

Figure 5.7: Error diagrams for predictors of a simulated temporal ETAS sequence. The parameters used in the simulation were those estimated for Southern Californian seismicity: $m_0 = 3, \mu = 0.1687, K = 0.04225, \alpha = 0.4491, c = 0.1922, p = 1.222$. Models were fitted to a 20-year training set and assessed on a 10-year test set. The ETAS conditional intensity predictor with the true parameters (green dashed line) performs very similarly to the ETAS conditional intensity predictor with estimated parameters (blue dotted line). The magnitude-dependent automatic alarms have parameter $u = 3.70$, chosen to minimise area under the error diagram in the training set. In the test set (solid black line), they perform slightly better than automatic alarms (red dotted-dashed line) and slightly worse than the ETAS conditional intensity predictors. No single strategy dominated any other single strategy.

ETAS v MDA: SCEC Data (Luen, 2010)

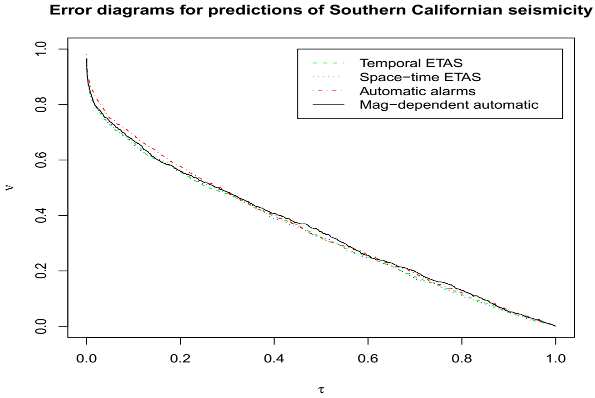

Figure 5.9: Error diagrams for predictors of Southern Californian seismicity. The predictors were fitted to the SCEC catalog from January 1st, 1984 to June 17th, 2004, and tested on the SCEC catalog from June 18th, 2004 to December 31st, 2009. For low values of $\hat{\tau}$, simple automatic alarms do not perform as well as the ETAS predictors. For high values of $\hat{\tau}$, MDA alarms do not perform as well as the ETAS predictors. Note that although success rates are determined for the test set only, predictors used both training and test data to determine times since past events (for simple automatic and MDA alarms) and conditional intensity (for ETAS predictors).

Conclusions: *Quantifauxcation*

- ETAS based on heuristics & metaphors, not physics
- Conflates frequencies with probabilities
- ETAS parameter estimates often unphysical
- ETAS fails goodness of fit test
- ETAS can't identify aftershocks well in ETAS simulations
- Doesn't predict noticeably better than far simpler methods *even in ETAS sims*
- Perhaps shouldn't rely on it
- Popular predictive policing algorithm uses ETAS; claim reliable because it's used in seismology