

One family of events, grown generation by generation.

Hawkes Processes

Self-Exciting Point Processes in Seismology, Epidemiology, Insurance and Finance, with Python

Callum Whitlow

ODIN PRESS

Hawkes Processes

Copyright © 2026 by **Callum Whitlow**.
All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Paperback version, First edition: 2026.
ISBN: forthcoming (Bowker allocation pending)
Trim Size: 6×9 inch. Paper color: White.

Source-code copyright. All Python source code listings reproduced in this volume—including those in the Appendix, the chapter bodies, and any auxiliary file distributed alongside this book—are the copyrighted intellectual property of **Callum Whitlow**. No reproduction, redistribution, adaptation, translation into another programming language, or incorporation into any derivative work, whether academic, commercial, or otherwise, is permitted without the prior written permission of the copyright holder. Single-machine execution by the individual purchaser of this volume for the sole purpose of reproducing a figure or table from this book is granted as a limited, non-transferable, non-sublicensable use and does not constitute redistribution.

Published by Odin Press.
`https://odinpress.org`
This book was typeset in LaTeX by Odin Press from a manuscript submitted by the author.

9 8 7 6 5 4 3 2 1 0

Preface

This book began with a number that kept changing its name. Seismologists call it productivity, epidemiologists the reproduction number, bankers and insurers contagion, and people who trade energy call it endogeneity. Each trade measures it from a list of times at which something happened, and each asks the same question of the list: what share of these events were caused by other events? I have fitted the same model in four of those trades, and the number under its different names is one number, the branching ratio of a Hawkes process. The book derives that process once, as a family tree of Poisson processes, and then takes it to the trades.

It takes it to five public records, one for each trade, each kept by the agency that collects it: the earthquakes of the 2019 Ridgecrest sequence from the U.S. Geological Survey, the daily COVID-19 case counts of the Delta wave from the Centers for Disease Control and Prevention, the bank failures listed by the Federal Deposit Insurance Corporation, the hurricane losses in the Storm Events Database of the National Oceanic and Atmospheric Administration, and the price spikes in the daily Henry Hub series of the Energy Information Administration. Every one of them bunches in time, and none of them looks like a Poisson record.

That bunching is where the book's question starts, not where it ends. A record clusters when its events breed events, and it also clusters when its events share a cause that nobody recorded: a season, a recession, a reporting calendar. A Hawkes model fitted with a constant baseline cannot tell the two apart, and it reads a moving baseline as excitation. So each application chapter fits the naive model, then lets the baseline move or sets a common-shock model against it, and keeps only what survives a residual test, a likelihood comparison and a profile interval. Asked that way, the five records do not give one answer. Two of them breed events: an earthquake makes aftershocks and an infection makes cases. Two only share shocks that look like breeding: bank failures that follow regional recessions, and

hurricane losses that follow active seasons. One breeds modestly: a spike in the price of gas makes another spike somewhat more likely for a few days, and no more. How a list of times can tell these cases apart, and where it cannot, is the question the book puts to every record.

The book is written for working analysts and graduate students in statistics, applied probability, seismology, epidemiology, insurance and finance who want each result derived, fitted to public data and tested before it is believed. The reader should know the Poisson process and the exponential distribution, conditional expectation and martingales, maximum likelihood, Laplace transforms, generating functions, convolutions, eigenvalues and Brownian motion; the last chapter also uses stochastic differential equations and Itô's formula. Everything else is built where it is used, and the results too long to prove here are stated precisely and cited.

Every number in the book is computed, and every figure drawn, by a Python program printed in Appendix A, one program for each chapter. The data are public-domain works of the United States Government, and the programs fetch them from the agencies' own addresses when no copy is at hand. A reader can download the records again, run the code and check each number, or point the same code at a newer copy of a record and see whether the answer has moved.

Plot the counts first, and let the baseline move before believing the excitation.

Callum Whitlow

Contents

Preface vii

List of Figures xiii

List of Tables xvii

1 Why Events Cluster 1
- 1.1 Five Records of Events . . . 2
- 1.2 The Poisson Benchmark . . . 8
- 1.3 What the Records Show . . . 13
- 1.4 History, Intensity and Compensator . . . 21
- 1.5 The Hawkes Process . . . 26
- 1.6 One Number, Five Names . . . 34

2 The Family Tree 39
- 2.1 Immigrants and Offspring . . . 40
- 2.2 Stability and the Stationary Rate . . . 44
- 2.3 Second-Order Structure . . . 53
- 2.4 Time Rescaling and Residuals . . . 59
- 2.5 Simulating Two Ways . . . 64
- 2.6 The Likelihood and Its Recursion . . . 69

3 Aftershocks 75
- 3.1 The Ridgecrest Catalogue . . . 76
- 3.2 Omori–Utsu Decay . . . 82
- 3.3 ETAS: Magnitudes as Marks . . . 86
- 3.4 Fitting by Likelihood and by EM . . . 92
- 3.5 Residuals and the Missing Small Quakes . . . 98
- 3.6 Forecasting Rates . . . 103

4 Epidemics and the Reproduction Number 111

4.1 Daily Counts and Their Weekly Diary . . . 112
4.2 Hawkes in Discrete Time . . . 116
4.3 The Reproduction Number Is the Branching Ratio . . 121
4.4 The Reporting Kernel . . . 130
4.5 A Finite Population . . . 140

5 Bank Failures and Mutual Excitation 149
5.1 The Failure Record . . . 150
5.2 Mutual Excitation . . . 156
5.3 Estimating the Branching Matrix . . . 163
5.4 Contagion or Common Shock? . . . 170
5.5 Rolling Windows . . . 181

6 Clustered Claims and Ruin 187
6.1 The Storm-Loss Record . . . 188
6.2 Ruin with Poisson Claims . . . 194
6.3 Claims That Breed Claims . . . 200
6.4 Ruin under Clustered Claims . . . 208
6.5 The Loading That Clustering Demands . . . 215

7 Markets at the Edge of Criticality 223
7.1 Spikes as Events . . . 224
7.2 Endogeneity Against a Moving Baseline . . . 228
7.3 Nearly Unstable Processes . . . 238
7.4 Heavy Tails and Roughness . . . 246
7.5 One Number Across Five Trades . . . 253

A Source Code 261
A.1 Chapter 1 — Why Events Cluster . . . 261
A.2 Chapter 2 — The Family Tree . . . 268
A.3 Chapter 3 — Aftershocks . . . 275
A.4 Chapter 4 — Epidemics and the Reproduction Number . . . 282
A.5 Chapter 5 — Bank Failures and Mutual Excitation . . 289
A.6 Chapter 6 — Clustered Claims and Ruin . . . 296
A.7 Chapter 7 — Markets at the Edge of Criticality . . . 303

References 311

Index 337

About the Author **339**

Chapter 1

Why Events Cluster

Five public records sit behind this book. The U.S. Geological Survey lists the earthquakes near Ridgecrest, California. The Federal Deposit Insurance Corporation lists the banks that have failed since 1934. The National Oceanic and Atmospheric Administration lists the storms that did damage. The Energy Information Administration prints the daily price of natural gas at Henry Hub, and the Centers for Disease Control and Prevention counted COVID-19 cases by day. Each is a list of times at which something happened, and in each the times bunch: a mainshock and a crowd of smaller quakes after it, a decade of bank closures and then long quiet, storm losses by season, winter runs of price jumps, an epidemic in waves. The book puts one question to every record. What share of its events were caused by other events?

The model that answers it is the Hawkes process, in which every event raises the rate of further events for a while. Its one summary number, the branching ratio n, is the mean number of events that one event causes directly. Seismologists call it productivity, epidemiologists the reproduction number, bankers and insurers contagion, market people endogeneity. Below one, cascades die out and the record settles to a steady rate; at one or above, the expected rate grows without bound.

This chapter builds the instrument and the doubt that goes with it. It delivers the conditional intensity $\lambda(t)$ and the likelihood written through it; the measurement that none of the five records is Poisson; the proof that this alone does not show self-excitation; and the Hawkes process with its stationary rate $\mu/(1-n)$. It runs in six steps: the records as point processes,

the Poisson benchmark, what the records show, history and intensity, the Hawkes process, and the branching ratio under its five names.

1.1 Five Records of Events

Let $0 < t_1 < t_2 < \cdots$ be the times at which a public record says that something happened, and let $N(t)$ be the number of those times in $(0, t]$.

That pair of objects is the whole subject of the book. The mathematics does not care whether the times are earthquakes, bank closures or price jumps; it cares whether they are independent of one another, or whether some were caused by others. The book asks that question of five public records, one for each trade, and answers it the same way each time: with a model for the rate of the next event given everything recorded so far, fitted by likelihood, tested by its residuals, and read only after the baseline rate has been allowed to move.

The question has a price in every trade. A record whose events breed events has fatter tails in its counts than a Poisson model allows, and an aftershock forecast, a capital model or a hospital plan built on the Poisson model is short in exactly the weeks that matter. A record whose events merely share a common cause, a season, a regime, a reporting calendar, has the same fat tails, and a model that reads them as breeding overstates how far one event spreads to the next. The second error is the one the data hide best: a self-exciting model fitted with a constant baseline to a record whose rate moves with the season reads the season as excitement, and the mean number of offspring per event that it returns can sit near one when the truth is far below it [1].

The book is built around that problem. Chapter 2 proves what the model is, a family tree of Poisson processes, with its stability, its second-order structure, the time-rescaling test of a fitted intensity, its simulation and its likelihood. Chapters 3 to 7 take one record each, and each record forces one new piece of the model: magnitudes as marks for the Ridgecrest aftershocks (Chapter 3), whole days and a reporting calendar for the epidemic counts (Chapter 4), several interacting types for bank failures in four regions (Chapter 5), a ruin calculation for clustered storm claims (Chapter 6), and the critical limit $n \to 1$ for spikes in the price of gas (Chapter 7). Each of those chapters runs the same three steps: it fits the naive model with a constant baseline and says what that fit seems to show; it lets the baseline move, or sets a common-shock model against it; and it reports

what survives, with a residual test, a likelihood comparison and a profile interval. The rest of this section fixes the words and describes the five lists.

Definition 1.1 (Point process)**.** A simple point process on $[0, \infty)$ is a random sequence of times $0 < t_1 < t_2 < \cdots$ with finitely many points in every bounded interval. No two events share a time.

A point process is the record itself, read as a random object. The word simple forbids two events at one instant, which is true of earthquakes timed to the millisecond and false of banks closed on the same day; records that break it are dealt with below. A point process is equally well described by the number of its events in each stretch of time.

Definition 1.2 (Counting process and inter-event times)**.** The counting process of a point process is $N(t) = \#\{i : t_i \le t\}$, right-continuous with unit jumps and $N(0) = 0$, and $N(a, b] = N(b) - N(a)$ counts the events in $(a, b]$. The inter-event times are the gaps $\Delta_i = t_i - t_{i-1}$, with $t_0 = 0$.

The counting process and the list of times carry the same information. In one direction,

$$N(t) = \sum_{i \ge 1} \mathbf{1}\{t_i \le t\},$$

and in the other the i-th time is the moment the count first reaches i,

$$t_i = \inf\{t \ge 0 : N(t) \ge i\}.$$

The two maps are inverse to each other because N is right-continuous and climbs by exactly one at each t_i, so $\{t : N(t) \ge i\}$ is the half-line $[t_i, \infty)$. We move between the two views without comment [2, 3]. The gaps add back to the times, $t_k = \Delta_1 + \cdots + \Delta_k$, so a record is also a sequence of waiting times.

Most records carry more than the time. An earthquake has a magnitude, a closed bank has a state, a storm has a bill.

Definition 1.3 (Marked point process)**.** A marked point process attaches to each event time t_i a mark m_i in a mark space $\mathbb{M}$: a magnitude, a damage, a state, the size of a price move. The event times alone form its ground point process.

A marked point process lets the size of an event change what follows it, the one physical fact every seismologist insists on: a magnitude 7 has more

Chapter 3

Aftershocks

A catalogue of earthquakes records the fewest small quakes exactly when there are most of them. In the first hours after a large shock the seismograms are crowded with the shock's own waves and with those of its first aftershocks, and many small events in between are never picked. Nothing in the catalogue marks the gap. A model fitted to such a record without a completeness cut reads the missing quakes as a quiet stretch of the process and builds the quiet into its parameters. That is the problem this chapter starts from, and it does not go away when the model is good.

Chapter 1 wrote the Hawkes process with marks and gave seismology's name for the branching ratio, productivity. Chapter 2 grew the process as a family tree, proved the time-rescaling test and fitted the model by likelihood and by EM, on records that held every event. The Ridgecrest catalogue of 2019, from the U.S. Geological Survey, holds only what its network saw, and its marks change how many children an event has. That forces the first extension of the model: the epidemic-type aftershock sequence (ETAS) model, with a productivity function $\kappa(m)$ and a Gutenberg–Richter law for the magnitudes.

The chapter delivers that model fitted to the 2019 sequence twice, naively and after the incomplete hours are cut, by direct likelihood and by EM, tested by its residuals, and used for a seven-day forecast of the number of aftershocks of magnitude 4 or more, given as a distribution and set against what happened. §3.1 measures the catalogue's completeness and its b-value; §3.2 fits Omori's decay; §3.3 builds ETAS and its branching ratio n; §3.4 fits it two ways; §3.5 tests the fit and repairs it; §3.6 forecasts.

The verdict is short. The triggering is real, and its n depends on where the catalogue was cut.

3.1 The Ridgecrest Catalogue

Chapter 2 fitted Hawkes processes to records that kept every event; an earthquake catalogue keeps only the events its network detected, a marked point process $\{(t_i, m_i)\}$.

Example 3.1 (The Ridgecrest record)**.** The record is the ANSS Comprehensive Earthquake Catalog (ComCat) of the U.S. Geological Survey, queried through its FDSN event service [5] on 6 October 2026 for every event of magnitude at least 2.0 within 100 km of 35.77° N, 117.6° W between 4 July 2018 and 4 July 2020, times in UTC. We keep the events whose type is earthquake, which removes 11 quarry blasts, and stop on 24 June 2020, the day the separate Lone Pine M5.8 sequence began. That leaves 7261 earthquakes with times t_i, in days from 4 July 2018, and magnitudes m_i: a marked point process in the sense of §1.1. It holds a quiet year of background, the Searles Valley M6.4 at 17:33:49 on 4 July 2019, the Ridgecrest M7.1 mainshock at 03:19:53 on 6 July 2019, 33.77 hours later and about 10 km away on a crossing fault of the same system [40], and their aftershocks. Two warnings travel with the file. The magnitudes are the catalogue's preferred values, of mixed types: local magnitudes for most events, moment magnitudes for the largest. And 1853 of the events are automatic solutions that no analyst reviewed, all of them from the first four months of the sequence and all below magnitude 3.

The mixed magnitude types are used as provided, as Savran and colleagues did in their evaluation of forecasts for this sequence [41]. We do not assume one magnitude scale; we check instead that the b-value is stable across cuts, and say where it is not. The quarry blasts are removed as Werner and colleagues remove explosions before modelling the California record [42].

Figure 3.1 shows the record before any model is fitted to it, with the intensity of the repaired fit of §3.4 drawn underneath.

The year before the M6.4 holds a scatter of small events; the two large shocks open a dense cloud that thins over the following year, with a late cluster around the M5.5 of June 2020. The lower panel already shows the shape the chapter will fit: a flat background broken by spikes that decay as powers of time.

Figure 3.1 — The Ridgecrest Record.

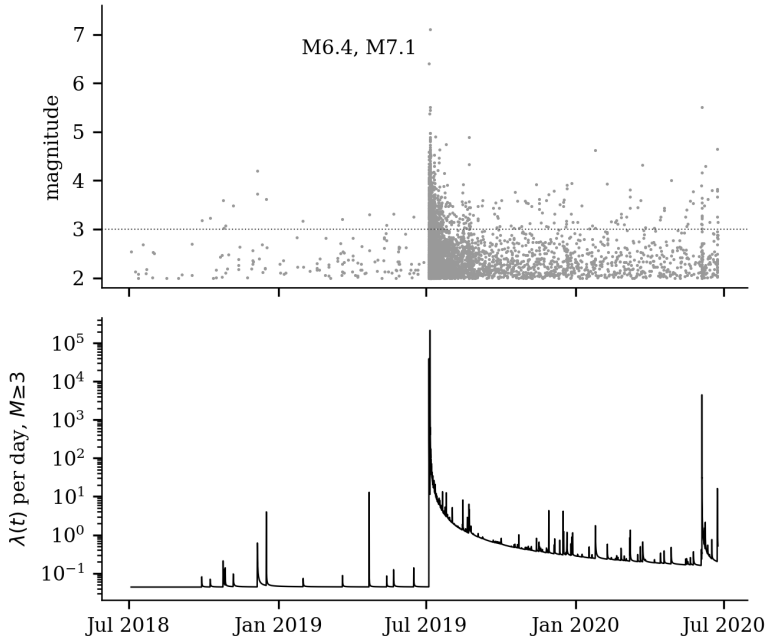


Figure 3.1: Ridgecrest earthquakes as magnitudes against time, with the conditional intensity of the repaired ETAS fit on a logarithmic scale.

Every model in this chapter is fitted to the events above a magnitude cut. Where to put the cut is the first question, and the catalogue answers it differently at different times.

Definition 3.2 (Magnitude of completeness)**.** The magnitude of completeness $M_c(t)$ of a catalogue at time t is the smallest magnitude above which every earthquake of the region at that time is recorded. A constant M_c is a cut that the whole window satisfies: $M_c \geq M_c(t)$ for every t in it.

Too high a cut wastes events; too low a cut lets the incompleteness into every estimate that follows [43]. Woessner and Wiemer compared the methods of estimating it, from the maximum of the non-cumulative frequency–magnitude curve to goodness-of-fit tests and the stability of the b-value [44]; the maximum-curvature method is fast but places M_c too low in bulk data. All of them test the magnitudes above a trial cut against one law.

Some records breed their own events. An aftershock raises the chance of the next one, an infection causes the next infection, a price spike makes another more likely. The Hawkes process is the model of that, and one number, the branching ratio, decides whether a cascade dies out or runs away. Seismologists call it productivity, epidemiologists the reproduction number, bankers and insurers contagion, traders endogeneity.

This book fits that one model to five public records: the Ridgecrest aftershocks, the Delta wave of COVID-19, nine decades of bank failures, thirty years of hurricane losses and the Henry Hub gas price. It asks each the same question: does it really breed events? Fitted naively, all five say yes. Let the baseline move, or test against a common shock, and only some still do.

The main results are proved, every estimate is tested by its residuals and computed two ways, and the Python that draws every figure is printed in full.

Callum Whitlow is an applied probabilist from Wellington who has forecast aftershock rates, priced clustered catastrophe claims and estimated reproduction numbers.

ISBN 13 978-468-41-8614-4
Paperback
9 784684 186144