Risk Mitigation Models for a Japanese Railroad

by

Alan Philip Rimm-Kaufman

B.S., Yale University (1991)
M.S., Massachusetts Institute of Technology (1994)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1996

© Massachusetts Institute of Technology 1996. All rights reserved.

Author .....
Department of Electrical Engineering and Computer Science
May 23, 1996

Certified by .....
Arnold Barnett
Professor of Operations Research and Management,
Sloan School of Management
Thesis Supervisor

Certified by .....
Amedeo Odoni
Professor, Departments of Aeronautics and Astronautics and
of Civil and Environmental Engineering
Thesis Supervisor

Accepted by .....
Thomas Magnanti
Co-Director, Operations Research Center, and
George Eastman Professor of Management Science,
Sloan School of Management

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUN 28 1996

ARCHIVES

LIBRARIES

Risk Mitigation Models for a Japanese Railroad

by

Alan Philip Rimm-Kaufman

Submitted to the Department of Electrical Engineering and Computer Science on May 23, 1996, in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Operations Research

Abstract

This thesis considers three hazards faced by a large Japanese railroad, JR East. These hazards are catastrophic earthquakes, rainfall-induced slope failures, and train collisions. We develop models to assess the risk that these hazards pose to passenger operations. We use these models to evaluate the efficacy of various mitigation options. When sufficient historical data are available, we construct statistical models; for rarer hazards, we construct models from "first principles" using physics, geometry, and basic probability theory.

For our investigation of catastrophic earthquakes, we model the number and speed of conventional and high-speed trains that could crash at quake-induced track failures. We use these models for typical and worst-case scenario analysis. For our investigation of rainfall-induced slope failures, we analyze the rainfall alert system currently used by the railroad. We propose alternative systems based on exponentially-averaged rainfall indices. Using one mountainous line as a case-study, we evaluate the performance of these various systems, and we assess the system characteristics responsible for superior performance. For our investigation of train collisions, we analyze data on both collisions and signal overrun incidents. We also develop a conceptual model for collision accidents. Due to intentional redundancy in the safety system, this model suggests improvements to the most reliable component of the safety system provide the greatest result. To integrate these three studies, we propose a method to assess the relative importance of these hazards.

Thesis Supervisor: Arnold Barnett
Title: Professor of Operations Research and Management,
Sloan School of Management

Thesis Supervisor: Amedeo Odoni
Title: Professor, Departments of Aeronautics and Astronautics and
of Civil and Environmental Engineering

Acknowledgments

I wish to express my gratitude to my two advisors, Professor Arnold Barnett and Professor Amedeo Odoni. They have patiently shared their experience and their ideas with me. I admire their approach to our field: creative, direct, applied O.R. Where this thesis shines brightest, their teachings are most evident. I also thank the other two members of my dissertation committee. Professor Daniele Veneziano helped me learn about earthquakes and landslides; my research has been much improved by his expertise. Professor Joe Sussman kept my research "on track" from the rail perspective. I am also grateful for Professor Sussman's work nurturing the MIT-JRE relationship: I am one of the many beneficiaries of his efforts. Personally, this has not been the easiest of years. Beyond fine advisors, these four are also decent *people*, and I am sincerely grateful to them for their compassion.

I thank JR East for allowing me to work on such interesting problems, and generously supporting me while I did so. I also thank the railroad two helpful trips to Japan. Japan is a beautiful country, and I enjoyed these visits. The JR East Safety Research Laboratory has taught me about their system and shared extensive databases. In particular, I wish to thank Mr. Makoto Shimamura for his support. The number of questions he has answered for me by phone, fax, and e-mail is beyond counting.

Several colleagues at the ORC in particular have helped me by serving as "sounding boards" for ideas: Árni Hauksson, Ármann Ingólfsson, Eddie Pinker, and Rob Shumsky. I am also thankful to them for their friendship, as well as to Sungsu Ahn, Mitchell Burman, Elaine Chew, Sougata Datta, Eriko Kitazawa, Kerry Malone, David Markowitz, Michael Miller, Joe Milner, S. Raghavan, Rodrigo Rubio, Sarah Stock, and Beril Toktay. My thanks to Paulette Mosley and Laura Rose for keeping the ORC running so smoothly.

The size of my family doubled while I was MIT. My thanks go out to my in-laws, Sylvia and Alfred ("How do you intend to write a thesis with no data?") Rimm, for their support and love, and similar appreciation to my sibs-in-laws and their kids:

Ilonna Rimm, Joe Madsen, and Miriam and Ben; David, Janet, Danny, and Rachel Rimm; and Eric, Alison, and Hannah Rimm. I couldn't have asked for a better gang.

My family has been a constant source of support, encouragement, and love. For getting me through MIT, and getting me through this year, I thank them: to Mark Sobel and Ellen Murphy, whose friendship has twined through my years like a Passiflora vine; to Susan and Ole Una, for helping me keep it all in perspective; and to Nina and Tom and Danny Weld, for the phone conversations and the laughter. I thank my grandparents, Nathan and Dorothy Sobel, for their non-stop encouragement, expressed in a steady stream of packages, cards, and calls. And I thank my parents, Norman and Essie Kaufman, for teaching me to value education, for urging me to go as far as I could, and for supporting me the entire way.

To my wife, Sara Elisabeth: words are insufficient to express my love. These last two years with you have been the best of my life. Thank you.

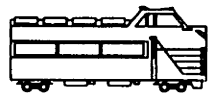

Dedication

This thesis is dedicated to my grandparents, Nathan and Dorothy Sobel.

Contents

1 Introduction **13**

1.1 Literature review 14

1.2 Overview of thesis 18

2 Seismic risk to trains **21**

2.1 Motivation and scope 21

2.1.1 Choosing a damage metric 23

2.2 Simple crash risk model 24

2.2.1 Model formulation 26

2.2.2 Case study: the Tohoku Shinkansen 31

2.2.3 Limitations 37

2.3 Regional crash risk model 48

2.3.1 Model formulation 48

2.3.2 Case study: the Tohoku Shinkansen 62

2.3.3 Case study: the 1995 Kobe quake 68

2.4 Conventional Trains 71

2.4.1 Model formulation 72

2.4.2 Case study: the Yaminote line 75

2.5 Comparing seismic risk: Shinkansen vs. conventional trains 78

2.6 Crash speed model 82

2.6.1 Model formulation 83

2.6.2 Case-study: Shinkansen and conventional trains 85

2.7 Conclusions 87

3 Slope failure risk to trains **91**

3.1 Motivation and scope 91

3.2 The current alert system 92

3.2.1 Description 92

3.2.2 Performance 95

3.2.3 Uncertainty in safety performance measures 100

3.2.4 Derived performance measures 102

3.2.5 Performance of modified versions of current rule 111

3.3 Alternative methodologies 114

3.4 The exponential rain index 117

3.4.1 Performance 118

3.4.2 Performance on national data 127

3.5 The double exponential rain index 132

3.6 Safety vs. disruption tradeoff 137

3.7 Slope failure risk in context 143

3.8 Conclusions 145

4 Collision risk to trains **149**

4.1 Motivation and scope 149

4.2 Collision summaries 150

4.3 Statistical approach 154

4.3.1 Collision data 154

4.3.2 Signal overrun data 156

4.4 Conceptual approach 168

4.4.1 "Normal Accidents" 168

4.4.2 Collisions as normal accidents 170

4.4.3 The triple-stress model 173

4.4.4 Numerical example 191

4.5 Conclusions 197

5 Conclusions **201**

5.1 Comparative risk analysis 201
5.1.1 Relative expected disutility 204
5.2 Contributions of thesis 209
5.3 Directions for additional research 211

Chapter 1

Introduction

JR East is a passenger railroad which provides service across the northern half of the Honshu island of Japan, including commuter service in the Tokyo metropolitan area. By almost any measure, this railroad is among the safest in the world. The railroad was formed in 1987 with the privatization of the Japanese National Railroad; since this time, JR East has provided over 750 billion passenger-kilometers of service with only thirteen passenger fatalities. [Nas95] In terms of passenger deaths per passenger kilometer, JR East is twenty times safer than the US commercial airline industry. [USD93] And these 13 fatalities of ticketed passengers include "non-operational" deaths such as falls from train platforms: since 1987, only one JR East passenger has been killed in a train accident. [RK95] As evident from such performance, the company places a great premium on safety, even maintaining an internal research laboratory dedicated to this purpose.

In this thesis, we address three safety concerns of the railroad: seismic hazard, slope failure hazard, and collision hazard. For each of these hazards, we assess the current level of risk and determine the efficacy of possible mitigation measures. In the last chapter of this thesis, we present a preliminary comparison of these hazards. Each subsequent chapter includes the necessary background for that portion of thesis. This first chapter provides an overview of generally relevant literature and of the thesis.

1.1 Literature review

Except for Chapter 3, which considers the relationship between rainfall and slope failure, much of this thesis lacks direct antecedents in the literature. At the same time, this thesis freely draws ideas from a broad corpus of previous research spanning many disciplines. In the next section, we briefly review this background material. We structure this review in four categories: general rail safety, probabilistic risk assessment for rare events, earthquakes, and risk integration frameworks. We postpone a literature review of statistical methods for slope failure prediction to Chapter 3. We do not present a literature review of train collisions as there appears to have been little research effort in this area to date.

General rail safety

One finds little previous work in the operations research literature on rail safety. The research that has occurred concentrates on the transportation of hazardous materials, either addressing the safest way to route noxious substances across a network [BC88] or estimating the damage from a spill. [GR84] [Swo87] In other disciplines, rail safety has been investigated using statistical and deterministic models. We employ the statistical approach in Chapter 3 to analyze slope failures, and in the first half of Chapter 4 to analyze train collisions. Statistical analyses have also been made for JR East level crossing accidents. [Ing94c] [MA95] Instead of taking a statistical approach, the actual dynamics of a rail accident can be modeled analytically [Kun72a] [Kun72b] or through simulation. [BAC90] Such research can suggest ways to reduce the chances of an accident [SE92] or the consequences of an accident. [Gal93] For our earthquake models, we searched the dynamics literature for models giving the probability a moving train is shaken from its tracks during a quake as a function of train speed; to the best of our knowledge, nothing has been published on this question to date.

The hazards addressed in this thesis arise either from natural processes (earthquakes, slope failures) or from system accidents (collisions). A core assumption

throughout this work is that these natural processes are inherently random (i.e., nature bears no specific grudge against JR East). In contrast, malevolent agents could intentionally target critical railroad systems to maximize damage to the company. Japanese railroads have been victimized by terrorists in the past. The years between 1949 and 1952 were known as the "era of railroad terror" in Japan, with 4,600 incidents of rail sabotage in 1949 alone. [Joh72, p. 81] More recently, the Aum Shiri Kyo cult attacked the Tokyo subway system with nerve gas in 1995. [SK95] However, as such attacks are not random events, assessing and mitigating terrorism risk demands an entirely different methodology than that of this thesis, and thus the terrorist threat remains outside the scope of this work.

Probabilistic risk assessment for rare events

This thesis draws heavily upon the rich field of probabilistic risk assessment. [BKJ92] Within this area, low-probability high-consequence accidents have received particular attention, starting with Rasmussen's seminal study of nuclear reactor safety. [R+75] As in Chapters 2 and 4 of thesis, a central challenge is estimating the frequency of rare events in the absence of sufficient historical data.

One approach is to employ statistical techniques to estimate the rate of a rare event. This approach treats the system as a "black box" that infrequently generates accidents. Using statistical methodology, inferences may be made from the occurrence, non-occurrence, or spacing of accidents. Ingólfsson uses this approach to bound the rate of fatal accidents within the Shinkansen "bullet train" system, which has had no fatal accidents in its thirty years of operation. [Ing94b] The disadvantage of such analyses is they typically produce weak bounds. As with any form of analysis, such methods rely on some assumptions, and approaching sparse data with different assumptions can lead to different results. One interesting demonstration of this occured in a volley of articles in *Nature*, where three sets of authors argued from the same two facts (namely, two serious reactor accidents in 4,000 reactor years of operation) and reached strikingly different conclusions about the frequency of nuclear reactor meltdowns. [IL86] [Edw86] [CO87] When assumptions are clearly stated and

accepted, however, the "black box" statistical approach can be quite useful, allowing one to determine a confidence interval for the rate of a rare event. [SS82] [BM87] [Swa90]

Instead of treating the system as a "black box," more direct models may be constructed to "leverage" the information contained in sparse data. This model-based approach finds application in many areas, such as nuclear reactor accidents, [COV90] [Bie93] oil spill accidents, [GHM94] and liquefied natural gas explosions. [Fai76] We use this approach to analyze seismic risk to trains in Chapter 2.

Earthquakes

More has been written on earthquakes than one could read in a thousand lifetimes. Survey articles offer valuable introductions to the challenges of earthquake engineering and earthquake prediction. [Eco95] [Bla95] More closely related to this thesis, the Kobe earthquake of 1995 provided an unwelcome but timely data point. In the weeks following this disaster, both the technical and popular presses provided extensive coverage of quake damage. Some of these articles also included useful information on rail damage. [Kri95] [Beg95] [JEN95]

In his doctoral thesis, Ingólfsson [Ing94a] investigates the efficacy of unequivocal early warnings to reduce deaths caused by an earthquake. Using both the judgment of experts as well as "first-principle" models, Ingólfsson estimates that short warnings (on the order of one to five minutes) have the potential to halve the death toll in urban areas. With a similar concern, we evaluate the potential benefits provided to the rail system by early warnings before earthquakes in Chapter 2.

In two papers with similar emphasis to our work, Shibata considers the impact of a major earthquake on the Japanese high-speed rail system. [Shi93] [Shi94] Shibata estimates the extent of Shinkansen damage following a major earthquake: track destroyed in 100 locations over a 200 km interval, 15 Shinkansen trains derailed, and hundreds of passenger fatalities. However, we infer from these papers that such estimates reflect Shibata's educated guesses, rather than model-based calculations. Our models in Chapter 2 suggest Shibata's predictions may be overly pessimistic.

In a study closely related to this thesis, Papadimitriou and Veneziano develop an engineering risk model to assess the performance of various early warning systems for the JR East Tohoku Shinkansen. [Pap95] [PV95] They employ an accepted attenuation submodel [KAT84] and an accepted seismicity submodel [Uts74] for Japanese earthquakes, and supplement these with a novel damage submodel for the Shinkansen viaduct. Their focus is on assessing the performance of Shinkansen seismic early warning systems across the entire spectrum of earthquake magnitudes. In contrast, our focus in Chapter 2 of this thesis is on the number of conventional and high-speed trains that would crash at track failures following a major quake (of magnitude similar to 1995 Kobe earthquake), the impact speeds of such trains, and the efficacy of various risk-reduction options.

Risk integration frameworks

Two other research efforts addressing safety in the JR East system are worthy of note. In his masters thesis, Nasser proposes a "safety performance index," consisting of a weighted sum of various accident consequences. [Nas95] His index combines various indicators of risk—delays, injuries, fatalities—into a convenient univariate measure. By observing the trend of this index over time, one can determine if system safety is improving, degrading, or holding steady. Using such a retrospective index to predict future safety trend is difficult, unless one believes the system has reached some sort of a steady state. Moreover, such indices may not address the reasonable allocation of safety expenditures across a portfolio of risk mitigation options.

Bohenblust, Sussman, and Odoni propose a more global approach to risk integration and safety planning, a so-called "safety management plan." [BSO95] A safety management plan has three elements: a corporate statement of safety policy, a corporate organizational structure that accommodates and encourages safety planning, and a risk-based safety analysis. This thesis may be viewed as supporting the last of these elements. Risk-based safety analysis combines risk assessment with cost-benefit analysis to determine a cost-effectiveness allocation of safety resources between various risk reduction options. The models of Chapters 2, 3, and 4 could assist with risk

assessment component, and the preliminary model of Chapter 5 could assist with the integration of these results.

1.2 Overview of thesis

Following this introduction, Chapter 2 analyzes the consequences of a catastrophic earthquake on the JR East conventional and high-speed rail networks. Two models form the core of this chapter. The first model considers the number of trains that could crash at quake-induced track failures, and the second describes the impact speed of such crashing trains. We find that each conventional train caught within the damage region of a major earthquake faces a 1/20 chance of crashing at a track failure. Shinkansen trains travel at higher speeds and thus require longer distances to stop; we find that each Shinkansen train caught with the damage region of a major quake faces a 7/10 chance of crashing at track failure. Further, our crash speed model suggests that trains which crash while braking tend to crash at high impact speeds, averaging more than 66% of full speed.

Chapter 3 considers the JR East slope failure alert system. This system relies on real-time precipitation data to stop trains when the slope failure risk is deemed high. We evaluate the performance of the current alert rule and family of alternative rules. For this purpose, we introduce a model of driver awareness to quantify the relative benefits of slowing versus stopping trains. We also model the probability of a slope failure condition on the value of a "rainfall index," and use this relationship to infer the relative valuation of risk versus delay for an arbitrary alert rule. We find the current rainfall alert system outperforms our family of proposed alternative rules, thus revealing which aspects of the current system are responsible for its superior performance.

Chapter 4 evaluates the risk of train collisions in the JR East conventional train network. The first half of this chapter takes a statistical approach, analyzing historic data on collisions and signal overruns to assess the trend in the JR East collision rate. The safety benefits of an on-going program of improvements to the signaling systems

across the network are discussed. We find the signal overrun rate has dropped 10% per year since 1987, and we estimate this improvement may have "averted" 1.4 collisions since privatization. The second half of the chapter introduces a more conceptual model of collision risk, and uses this model to explore the efficacy of various options to reduce the probability of train collisions. Between driver vigilance, schedule slack, and the signaling system, we find the signaling system provides most of the protection against collisions, and improvements are most beneficial when applied in this area.

Chapter 5 concludes this thesis by examining the relative importance of these three hazards. While collision risk and slope failure risk are relatively well understood, there is considerable uncertainty regarding when the next large earthquake will strike the Tokyo area. We find the importance of the seismic threat relative to the two other hazards is highly sensitive to one's beliefs regarding the annual probability of a major Tokyo earthquake. This chapter also summarizes the principal findings of the thesis and offers suggestions for additional research.

The reader is advised that symbols are recycled throughout the thesis: we find seven different uses for "λ" alone. With each new model, we state the current meaning of each symbol involved.

Chapter 2

Seismic risk to trains

2.1 Motivation and scope

Japan sits atop the junction of the Pacific, Eurasian, and Phillipine tectonic plates, and has accordingly suffered the devastation of major earthquakes throughout its history. The Hokkaido quake of 1730 killed 137,000 people. The Kanto quake of 1923 killed 143,000 people. [Ing94a] The Kobe quake of 1995 killed 5,300 people. [Beg95] And earthquakes will continue to strike Japan: it is not a question of whether Japan will suffer another major quake, but when. As the quakes themselves cannot be stopped, the focus must be upon reducing the damage they cause.

The damages caused by major earthquakes include death and injury, destruction of buildings and property, and disruption of trade and commerce. This chapter considers one very specific type of earthquake damage: moving trains which crash as a result of an earthquake disrupting their tracks. Thankfully, major earthquakes are rare, and the data on this specific type of damage is correspondingly scarce. To understand the likelihood and the extent of such damage, and to investigate what a railroad might be able to do to minimize it, we construct a series of mathematical models that offer insight into these matters. We do not intend these models to serve as detailed engineering studies, to apply to other types of disasters, or to apply to other transportation modes. Rather, in the absence of historical experience, these models offer a preliminary but useful understanding of a specific accident scenario of concern

to a railroad.

We begin our analyses by considering the JR East Shinkansen "bullet train" system. In contrast to the dense maze of conventional rail lines running through and around Tokyo, the high-speed Shinkansen network has a simple topology. This network is essentially linear, comprised of four lines emanating from Tokyo: the 535 km Tohoku line to Morioka; the 360 km Yamagata line to Yamagata; the 330 km Joetsu line to Niigata; and the 1180 km Tokaido/Sanyo line to Hakata. This linear topology simplifies our models.

Section 2.2.1 presents a simple model for the distribution of the number of trains that would crash in the aftermath of a major earthquake. This model aims to capture the most essential factors relevant to the number of crashing trains: the extent of track damage, the density of track damage, the density of trains, the speed of trains, the emergency deceleration rate of trains, and the time at which trains become aware of the quake hazard (either delay or early warning). Section 2.2.2 applies this model to the Shinkansen system as a case study, though the general approach applies to conventional trains as well. Section 2.2.3 discusses the effect of some of the simplifying assumptions of this first model.

Section 2.3.1 extends this basic model into a regional crash risk model. The regional model considers the different operational characteristics of different sections of a rail line, and also includes the effects of tunnels, station stops, and the time of day that the quake occurs. Section 2.3.2 again uses the Shinkansen system to demonstrate the model, employing the Tohoku line between Tokyo and Morioka as a case study. We conclude our Shinkansen analyses by considering the 1995 Kobe quake. The trains of the Sanyo Shinkansen escaped harm by a fortunate accident of timing: the Kobe quake struck early in the morning, before the trains began service. Section 2.3.3 estimates what might have occurred had this quake struck later in the day.

Section 2.4 modifies these crash risk models appropriately to apply to the conventional rail system. As a case study, Section 2.4.2 considers the impact a large earthquake in the metropolitan Tokyo area could have upon the conventional rail

system.

All of the models in Sections 2.2 and 2.3 are concerned with the number of Shinkansen trains or conventional trains that might crash in an earthquake. We next consider a different issue: the speed at which crashing trains would crash. Section 2.6 presents a model for the distribution of crash speed and applies it to both Shinkansen and conventional trains.

Section 2.7 concludes the chapter by discussing the policy implications of these various models.

2.1.1 Choosing a damage metric

There are many damage metrics that a railroad might select to measure the consequences of an earthquake on their rail system. These might include:

- the number of derailed train cars, and the cost to set them back on the rails,
- the number of damaged train cars, and the cost to repair them,
- the duration of service disruption, and the resulting lost revenues,
- the kilometers of damaged track, and the resulting repair cost,
- the passenger, employee, and third-party injuries; and accompanying costs,
- the passenger, employee, and third-party fatalities; and accompanying costs,
- the erosion of public confidence in the safety of the rail system, and the loss of corporate face.

These damage metrics would be of direct interest to a railroad. However, they present formidable modeling challenges. Modeling damage to individual train cars demands a detailed understanding of the crash dynamics and a car's structural properties. [Gal93] Modeling damage to a railroad's public image requires quantifying many intangible societal and individual concerns. Modeling the human death and injury from train crashes is made very hard by the great variability observed in the accident

record. [Bin93] [B+93] (As an extreme example, the high-speed French TGV train derailed on December 21, 1993 while traveling 180 mph; there were no fatalities and only two of the 170 passengers sustained light injuries. [Knu94])

We introduce two proxy measures of damage to sidestep these difficulties:

- the distribution of the number of moving trains traveling between stations that crash in the aftermath of an earthquake,
- and the distribution of the impact speed of such trains.

We opt to estimate the number of moving trains that would crash in an earthquake, and not extend our models to passenger fatalities.

Though an earthquake may damage trains in maintenance yards and depots, such trains neither carry passengers nor travel at high speeds. And though an earthquake could damage a passenger-laden train paused at a station platform, such passengers face little additional risk by virtue of being on board a train during a quake. (One could even argue it is safer to be on a stationary train when an earthquake strikes than at work in a high-rise office building.) We believe that of most concern to a railroad are accidents involving passenger-carrying trains crashing at high speed. Accordingly, we restrict our models to in-service trains traveling on the main line between stations.

2.2 Simple crash risk model

We now estimate the chance that a major earthquake harms a moving train. Though our general approach applies equally well to the Shinkansen network and the conventional rail network, the particulars differ slightly. For expository convenience, Sections 2.2 and 2.3 focus on Shinkansen trains, and Section 2.4 addresses conventional trains.

How might a quake damage a moving Shinkansen train? A quake might destroy the track in front of an approaching train, or shake a moving train from its tracks. In urban or industrial areas, damaging fires and explosions could follow an earthquake. Buildings might topple onto trains or destroy tracks. In coastal areas, tsunamis (tremendous sea waves sometimes caused by marine quakes) could sweep away trains

and tracks. Fortunately, some of these horrific possibilities appear unlikely. Most of the Shinkansen track runs between cities, running through suburban and rural areas. As only a small portion of the Shinkansen network is inside urban areas, only a small portion of the network faces danger from buildings, fires, or explosions. Similarly, most of the Shinkansen track is inland, safely out of a tsunami's reach. As such, we shall not deal with collapsing buildings, tsunamis, fires, and explosions in this chapter.

Evidence from the recent Kobe disaster suggests that an earthquake could harm a moving train by shaking it from its tracks or by disrupting its tracks. Due to a fortunate accident of timing, the early-morning Kobe quake struck before the railroad began its daily operations, and thus did not harm any moving trains. The violent ground motion, however, did toss or topple several parked conventional trains off their tracks in a depot. This quake also twisted conventional track in numerous places, and shattered the Shinkansen viaduct in several spots near both ends of the Kobe tunnel. [SRL95l] We believe that had the Kobe earthquake struck later in the morning, when the trains were running, some moving trains would have crashed from being thrown from their tracks or from derailing at failed track. (Section 2.3.3 considers what effect this quake might have had on the Sanyo Shinkansen system.)

Thus, the Kobe experience suggests an earthquake may harm a moving Shinkansen train by destroying its track or by shaking it off its track. What is the relative likelihood of these situations? Railroad engineers at JR East suggest that the motion required to damage the Shinkansen track is far less than the motion required to derail a moving Shinkansen train. [Pap95] If true, trains are more likely to crash by running into broken track, because the viaduct that supports the track would break long before there was sufficient ground motion to toss trains off their tracks. We accept the judgment of the JR East engineers and thus make the following strong assumption: we assume that viaduct damage poses the primary risk to moving Shinkansen trains, and that shaking risk may be a second-order effect. Accordingly, the models described in Sections 2.2 and 2.3 do not include shaking risk to Shinkansen trains.

2.2.1 Model formulation

The conceptual framework of our model rests on a simple observation. If a moving Shinkansen train needs x kilometers to stop, and if there is damaged viaduct within the x kilometers in front of this train, then this train will be incapable of stopping before reaching the damaged viaduct section. In this case, the train will crash at the viaduct damage site while still traveling at some speed. (We use the phrase "crash" as a synonym for "derail at damaged track while moving with some speed" throughout this chapter; depending on the train's speed and location, the consequences of a "crash" could range from a harmless jolt to a catastrophic accident with loss of life.) We wish to determine the probability of a crash, and how many trains might face this risk. In Section 2.6, we also determine how fast such trains might be traveling when they strike damaged viaduct.

We can use elementary physics to determine the distance needed to stop a moving train. If a train is traveling at constant speed v_c m/s, and then begins braking with constant deceleration a m/s/s, the train will stop after traveling $v_c^2/2a$ meters. [HR74, p. 34] The Shinkansen emergency brakes do produce nearly constant deceleration at a rate of $a_e = 0.75$ m/s/s, making this a reasonable model. [PV95] We use the term "braking distance" to describe the distance needed to stop a train under emergency braking.

The duration of an earthquake can range from an instant to several minutes. For this initial model, we ignore the temporal extent of a quake and assume an earthquake and its accompanying damage occur instantaneously. (Section 2.2.3 examines this assumption.) For convenience, we anchor our time origin at this instant, defining the time of the quake as $t = 0$.

Let $\tau > 0$ denote the time at which the train begins emergency braking. As we have chosen our time origin as the instant of the quake, τ is also the delay time between the moment the quake strikes and the moment the train begins emergency braking. Assume the train has constant speed v_c m/s at the moment of the quake. For the moment, assume the track and viaduct ahead of this train are free from damage. In this case, the train travels τ seconds at speed v_c m/s before it begins

braking, covering $v_c\tau$ meters. The train then brakes with constant deceleration rate a_e m/s/s, traveling an additional $v_c^2/2a_e$ meters before stopping. We use the term "exposed stopping distance" for this length $x = v_c\tau + v_c^2/2a_e$. As the train needs x meters to stop (taking into account the delay before braking), any viaduct failures within this distance expose the train to the risk of a crash, and thus we say the train is "exposed" to risk over this distance.

If, by some means, the train could receive advance warning of the quake, it could begin emergency braking before the quake struck. By starting to brake before the quake strikes, the train can have a smaller exposed stopping distance. (Of course, the total distance to stop the train remains the same; the train just gets a "head-start" on braking before the quake.) By starting to brake before the quake, a train reduces its risk of encountering viaduct damage. We model such an early warning by allowing $\tau < 0$ and modifying the exposed stopping distance formula. By starting to brake at time $\tau < 0$, before the quake strikes at time zero, the train can slow from speed v_c to speed $\max\{0, v_c + a_e\tau\}$ with no risk. (Recall $\tau < 0$, so $v_c + a_e\tau < v_c$. The maximum function ensures that the train speed remains non-negative, sensibly asserting that a train would not back up after an emergency stop.) Thus, in the case of early warning, the train is traveling at $\max\{0, v_c + a_e\tau\}$ m/s when the quake strikes, and needs an additional $(\max\{0, v_c + a_e\tau\})^2/2a_e$ meters to stop. We combine the early warning case and the delayed braking case to produce a formula for x, the exposed stopping distance (Figure 2-1).

$$x = \begin{cases} v_c\tau + \frac{v_c^2}{2a_e}, & \text{if } \tau \geq 0 \\ \frac{[\max\{0, v_c + a_e\tau\}]^2}{2a_e}, & \text{if } \tau < 0 \end{cases} \tag{2.1}$$

We next present a highly simplified model for viaduct damage caused by an earthquake. We model locations of viaduct damage as discrete points following a spatial Poisson process of rate λ over an interval of length $2r$ (Figure 2-2). We intend r to denote the radius of the region of potential damage, so $2r$ represents this region's diameter, or length. In other words, we model viaduct failures as isolated points, not as damaged intervals; we assume the length of the "region of potential damage" is $2r$;

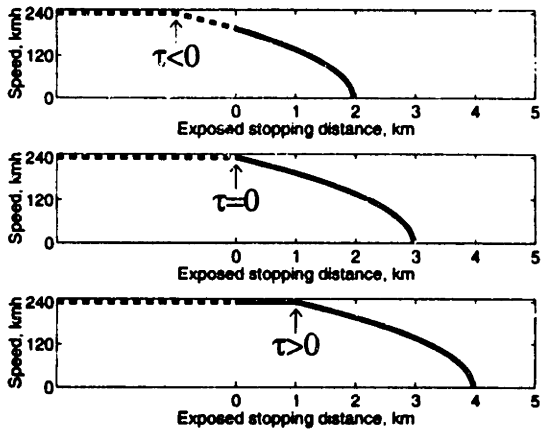

Figure 2-1: Three examples of the effect of τ on the emergency stopping distance and on the braking curve. The light dotted vertical line on each plot denotes time zero, the instant of the quake. The train begins braking at time τ, indicated by the vertical arrow. When $\tau < 0$, the train enjoys some risk-free deceleration before the quake strikes, reducing its exposed stopping distance. When $\tau > 0$, the delay in braking increases the train's exposed stopping distance.

we assume the number of failures in the region of potential damage follows a Poisson distribution with mean $2\lambda r$ failures; and we assume (conditional on the number of failures in the region) that the these failure points are uniformly distributed across the region of potential damage.

There is much to criticize about this simple damage model. Surely the rate of viaduct failure varies with distance from the quake's epicenter and depends on local geological conditions. Further, the Poisson assumption implies that each meter of viaduct fails or survives independently of its neighbors. This seems implausible: we suspect that if a quake destroys the viaduct at some point, nearby points are much more likely to fail as well. We evaluate these criticisms in Section 2.2.3, and accept the simplified viaduct damage model for our initial crash risk model.

The Shinkansen line consists of double track sharing a single elevated viaduct. That is, trains going in one direction (for example, going north) use one track, trains going in the other direction (going south) use the other, and both tracks rest upon the same concrete roadbed. We assume that the spacing, or "headway," between subsequent northbound trains is a constant distance of Δ meters. We assume the

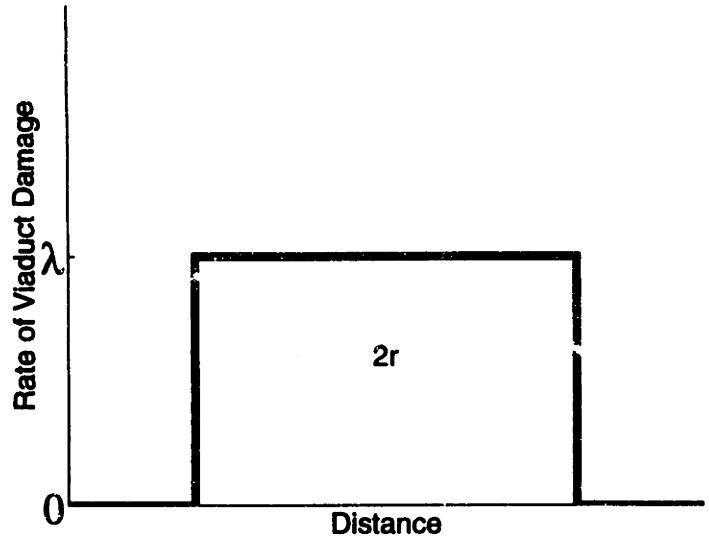

Figure 2-2: Rate of spatial Poisson process model for viaduct damage.

southbound trains also have headways of Δ meters. We assume that at a random instant (in particular, at the instant a quake strikes) the distance from the northbound train to the closest approaching southbound train is uniformly distributed between 0 and Δ meters. This preliminary model ignores the effects of stations, shunting track, the differing speeds of local and express trains, and tunnels.

We now assume an earthquake strikes when the Shinkansen trains are running, and focus our attention on a moving northbound train. We make a second strong assumption: we assume that $\Delta > 2r$, so that at most one northbound train could be inside the region of potential damage when the quake strikes. Assuming the location of the quake is independent of the location of the train, the probability that the quake catches a northbound train inside the region of potential damage is thus $2r/\Delta$.

Equation 2.1 models the exposed stopping distance, x, of such a train caught in the region of potential damage. We assume $x \ll 2r$, or that the exposed stopping distance is small relative to the size of the region of potential damage. Accordingly, we ignore the possibility that a train enters or departs the potentially damaged region while braking to a stop.

Consider a northbound train inside the region of potential damage when the quake strikes. The train requires x meters to stop after the quake strikes. The number of

viaduct failures in the x meters ahead of this train follows a Poisson distribution with mean λx failures. If there are no failures in this distance, this train will stop safely. If there are any viaduct failures in this distance, this train will crash. (Recall that the consequences of a "crash" may vary.) Let A denote the event "a northbound train crashes at a viaduct failure," and let ρ denote the probability of event A. Let B denote the event "the closest approaching northbound train is inside the region of potential damage."

$$\begin{aligned} \rho &= \Pr(A|B) \cdot \Pr(B) \\ &= \left(1 - e^{-\lambda x}\right)\left(\frac{2r}{\Delta}\right), \qquad \Delta \geq 2r \end{aligned}$$

We can repeat a similar analysis for the closest southbound train approaching the region of potential damage. Like the northbound train, the probability an earthquake catches a southbound train inside the region of potential damage is $2r/\Delta$. Like the northbound train, the southbound train will also crash if it encounters any viaduct failures in its stopping distance. However, the probability that a southbound train inside the region of potential damage encounters a viaduct failure is not independent of the probability that a northbound train inside the region of potential damage encounters a viaduct failure, because the northbound and southbound track share the same viaduct. For example, given that the northbound train crashes, there must be at least one viaduct failure in the potentially damaged region, thus increasing the chance the southbound train crashes.

This preliminary model ignores this dependence: we assume the probability that the southbound train crashes is identical to and independent of the probability that the northbound train crashes. (Section 2.2.3 discusses this assumption.)

$$\begin{aligned} \rho &= \Pr(\text{n'bound train crashes}) = \Pr(\text{s'bound train crashes}) \\ &= \left(1 - e^{-\lambda x}\right)\left(\frac{2r}{\Delta}\right) \end{aligned} \tag{2.2}$$

With these assumptions, the number of Shinkansen trains crashing at viaduct

failures in the aftermath of an earthquake follows a Binomial$(2, \rho)$ distribution.

$$\Pr(k \text{ trains crash}) = \binom{2}{k} \rho^k (1-\rho)^{2-k}, \qquad k = 0, 1, 2.$$

We believe that a high-speed crash of even one Shinkansen train would constitute a major catastrophe for a railroad. Accordingly, we focus our attention on p, the probability of one or more crashes.

$$p = \Pr(\text{at least one train crashes}) = 1 - (1-\rho)^2 \tag{2.3}$$

Equations 2.1, 2.2, and 2.3 comprise our preliminary model. Together, they estimate the probability that at least one Shinkansen train crashes due to viaduct damage in an earthquake that strikes when the trains are running. This model has six parameters: a_e, the emergency braking rate of a Shinkansen train; v_c, the cruising speed of a Shinkansen train; Δ, the headways between successive Shinkansen trains on the same track; r, the radius (half the length) of the region of potential damage; λ, the spatial rate of viaduct damage; and τ, the time the Shinkansen trains start to brake. The model is far from perfect; we have noted several of its deficiencies above. However, this simple framework can offer considerable insight.

2.2.2 Case Study: northern region of Tohoku Shinkansen

To demonstrate the simple crash risk model, we will assume a major Kobe-like earthquake strikes in the northern region of the Tohoku line, near Sendai. In contrast to operations near Tokyo, Shinkansen operations in this region are characterized by fewer trains per day, greater distance between trains, greater distance between stations, and higher operating speeds. (The regional model of Section 2.3 considers the entire line, including the Tokyo area.)

We first estimate values for the six parameters and use the simple crash risk model for a direct calculation (Table 2.1). Perhaps the most unambiguous parameter value is a_e, the emergency deceleration rate. From engineering documents for the Tohoku

Parameter	Value	Description
a_e	0.75 m/s/s	emergency braking rate
v_c	240 km/h	cruising speed
Δ	70 km	headways between trains
r	15 km	radius of potential damage region
λ	0.5 failures/km	average rate of viaduct failure
τ	0 secs	delay before braking

Table 2.1: Base-case parameters for simple crash risk model.

Shinkansen, [PV95] we determine $a_e = 0.75$ m/s/s. (This deceleration rate might differ for other models of Shinkansen trains.) The headway between successive trains on the same track, Δ, and v_c, the train cruising speed, vary across the Shinkansen network. To represent service in the northern part of the Tohoku line, near Sendai, we use $\Delta = 70$ km and $v = 240$ km/h. With these values, a stationary observer would see a train pass by (going either way) about every 9 minutes, matching the actual service schedule in this region.

Choosing values for the remaining three parameters is more challenging. For this base-case calculation, we assume all Shinkansen trains begin emergency braking at the instant of the quake: $\tau = 0$. We assume the radius of the region of potential damage is $r = 15$ km, and that the rate of viaduct failures within this region is $\lambda = 0.5$ failures/km. Thus, we assume all of the quake-induced viaduct failures fall within a 30 km region, and the number of such failures in the entire region follows a Poisson distribution with mean 15. We intend these values to represent an earthquake similar to the Kobe disaster, where a large quake struck close to the Shinkansen tracks with no warning. These base-case parameter values are only guesses informed by conversations with railroad engineers and seismologists; we rely on sensitivity analysis to assess their impact on our conclusions.

We use these parameter values to make the following calculations.

$$
\begin{aligned}
x &= \text{exposed stopping distance} \\
&= \frac{v_c^2}{2a_e} = 2.96 \text{ km} \\
\rho &= \Pr(\text{n'bound train crashes}) = \Pr(\text{s'bound train crashes})
\end{aligned}
$$

$$= \left(1 - e^{-\lambda x}\right)\left(\frac{2r}{\Delta}\right) = (0.77)(0.43) = 0.33$$

$$p = \text{Pr(at least one train crashes)}$$

$$= 1 - (1-\rho)^2 = 1 - (.67)^2 = 0.55$$

This model suggests that if a Kobe-like earthquake strikes near the Shinkansen tracks when the trains are running, there is a 0.55 probability that at least one train crashes at a track failure. The northbound and southbound train each enjoy odds of two to one of stopping safely. We do not intend this rough model to offer a precise risk estimate, but it does suggest the order of magnitude: the odds are about even that a major quake striking during operating hours in the northern part of Honshu would cause some Shinkansen trains to crash.

How sensitive is this result to the values of the six parameters? To answer this question, we vary each parameter in turn and observe the effect on p, the probability of at least one crash (Figure 2-3). We see that this base-case risk is not particularly sensitive to our choice of a_e, the emergency braking rate; to v_c, the train speed; to τ, the delay (or early warning) before braking; or to Δ, the headway between trains. In contrast, the risk varies almost linearly with the radius of the damage region, r. The severity of the earthquake causes this effect: due to our choice of λ, the region of potential damage is devastated, and any train caught within this region is almost certain to crash. (Recall the average distance between failures is 2 km, and trains need 3 km to stop.) Thus, the probability of at least one crash is a little less than twice the probability a train is caught in this region, and this varies linearly with r.

$$p = 1 - \left[1 - \left(1 - e^{-\lambda x}\right)\left(\frac{2r}{\Delta}\right)\right]^2 \approx 1 - \left[1 - \frac{2r}{\Delta}\right]^2 = 2\left(\frac{2r}{\Delta}\right) - 4\left(\frac{2r}{\Delta}\right)^2 \approx \left(\frac{4}{\Delta}\right) r$$

By the same reasoning, our risk estimate is relatively insensitive to increases in λ, as the rate of devastation is already severe enough to catch almost every train operating in the damage region. Significant decreases in λ can significantly decrease the chance of one or more crashes.

We may use this simple crash risk model for tradeoff analysis. Suppose the railroad

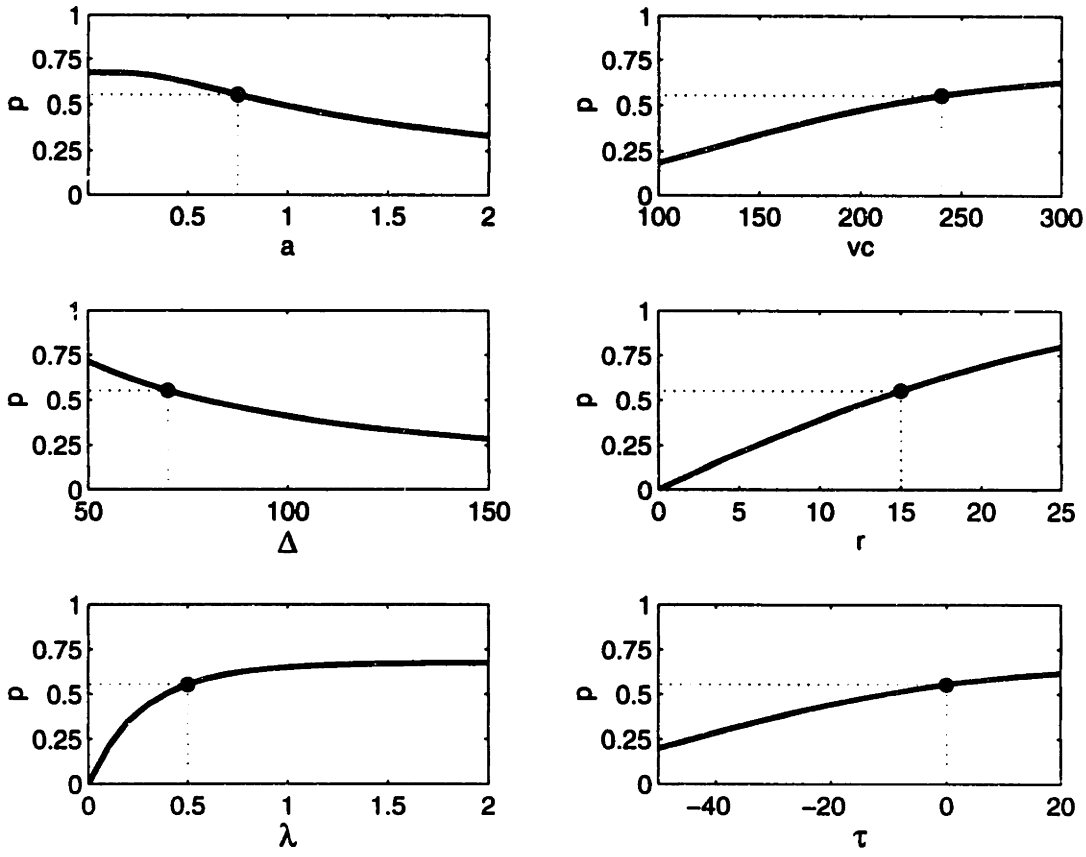

Figure 2-3: Sensitivity of p, the probability of at least one crash, to a_e, the emergency braking rate, in m/s/s; to v_c, the train cruising speed, in km/h; to Δ, the headway between trains, in km; to r, the radius of the region of potential damage, in km; to λ, average rate of viaduct failure, in expected failures per kilometer; and to τ, the delay (or early warning) before braking, in seconds. The black dot on each plot indicates the base-case value.

Increase emergency deceleration rate by 30% to 1.0 m/s/s
Decrease train cruising velocity by 13% to 210 km/h
Increase train headways by 13% to 80 km
Decrease radius of quake damage by 12% to 13 km
Decrease rate of viaduct failures by 24% to 0.38 fail/km
Increase early warning of quake to 11 seconds

Table 2.2: Ten percent risk reductions. Using the simple crash risk model, any one of these changes would produce a 10% reduction in the base-case probability that some Shinkansen trains crash at viaduct failures in the aftermath of a Kobe-like quake.

wished to reduce this risk by 10%, lowering p from 0.55 to 0.495. Hypothetically, the railroad could consider installing faster emergency brakes, operating trains more slowly, or operating trains less frequently. The railroad might attempt to strengthen the viaduct to reduce the amount of track damage, or the railroad might invest in improved early warning systems. Our model allows the comparison of these various possibilities (Table 2.2).

To achieve a 10% reduction in this risk, the railroad could increase the emergency braking rate by 30%. Alternatively, the railroad could reduce the speed of the Shinkansen by 13%. This reduced-speed policy appears economically untenable, as high speed is the primary competitive advantage of Shinkansen transit. Another economically untenable option is a 13% increase in the headways between trains, as this would require the railroad to reduce the frequency of train service. Alternatively, either shortening the spatial extent of quake damage by 12%, or lessening the rate of damage by 24%, would reduce this risk by 10%. The railroad might achieve such reductions by strengthening the viaduct structure. Eleven seconds of warning would also offer a 10% risk reduction. However, there are no reliable methods for earthquake prediction, even for last-minute warnings. [Ing94a] [Bla95] [Eco95] The current early warning system 5 to 10 seconds of warning for inland quakes, and between 20 and 30 seconds of warning for off-shore quakes. [Pap95] Three of the six proposals may be technically and economically feasible: improving the braking system, reducing the expected extent of track damage, and short early warnings. However, we must reiterate that all of these various proposals offer a small return: they reduce a risk estimate 10%, from 0.55 to 0.495.

What if the railroad sought to reduce this risk by a more significant amount, say by 50%? Our model suggests that obtaining such an improvement would be a formidable challenge (Table 2.3). Safely decelerating a train at 2.6 m/s/s, a rate of deceleration roughly 60% faster than the deceleration of a jet airplane landing on a runway, would require advanced braking technology. As Shinkansen trains lack seat belts, such rapid deceleration could cause passenger injuries. Neither running half as many trains nor running trains half as fast is economically acceptable, and there are no instruments available today that accurately provide over 40 seconds of warning before an earthquake. [Ing94a] Only the track-strengthening options appear remotely feasible to halve this risk. Though costly and technically difficult, the most manageable options involve decreasing the expected rate and the expected extent of viaduct failures, either by strengthening the structure or by installing seismic isolation devices.

We also note that viaduct protection, if successful, provides an additional benefit to the railroad. The other options attempt to save trains by keeping trains away from viaduct damage: by stopping trains quicker, by stopping trains sooner, and by spacing trains further apart. In contrast, a track-strengthening policy attempts to minimize viaduct damage directly. Such a policy could result in less viaduct damage in an earthquake, leading to quicker track repairs, and hopefully quicker resumption of service. A rapid resumption of Shinkansen service following a terrible earthquake could prove soothing to a damaged national psyche; it would facilitate a speedy resumption of business and commerce; and it would minimize the railroad's lost revenue. The Shinkansen generates about 22% of JR East's revenues, generating about ten million dollars per day. [Eas93] Each day the Shinkansen remained inoperable after an earthquake would generate losses which would go directly to the bottom line on the railroad's balance sheet.

Though simple in scope, we can draw three important conclusions from this crash risk model. First, this model suggests that if a significant earthquake strikes close to the Shinkansen track when the trains are operating, the odds are high that at least one Shinkansen train will crash at a track failure. While we do not put great faith in

Increase emergency deceleration rate by 245% to 2.6 m/s/s
Decrease train cruising velocity by 47% to 130 km/h
Increase train headways by 121% to 155 km
Decrease radius of quake damage by 55% to 7 km
Decrease rate of viaduct failures by 72% to 0.14 fail/km
Increase early warning of quake to 41 seconds

Table 2.3: Fifty percent risk reductions. Using the simple crash risk model, any one of these changes would produce a 50% reduction in the base-case probability that some Shinkansen trains crash at viaduct failures in the aftermath of a Kobe-like quake. Most of these changes are technologically or economically infeasible; moreover, each of these changes may have very different costs for the railroad.

our exact estimate of this probability, our models do suggest that it is on the order of 0.5, and not 10^{-6} or 10^{-9}. Second, sensitivity analyses indicate that this conclusion is relatively robust, and remains valid for reasonable changes in the values of the model's six parameters. Third, this model suggests that track strengthening measures appear the most feasible means to reduce the likelihood of Shinkansen crashes in the aftermath of an earthquake.

2.2.3 Limitations of simple models

We must verify, however, that these broad conclusions are not mere artifacts arising from the simplifications of this initial model. We also wish to tune these estimates more finely, and we wish to obtain a more accurate understanding of this risk along a rail line. To these ends, this section considers the effects of some of the simplifying assumptions made in the initial model, and Section 2.3 presents a more involved regional crash risk model.

Poisson track damage

The simple crash risk model uses a spatial Poisson process to place track failures within the region of potential damage. This implies each meter of track fails or survives independently of its neighbors, which seems highly improbable.

However, we defend this modeling approach on the grounds that the mean length of a track failure is likely to be small relative to the mean inter-failure distance. In

the Kobe quake, for example, the failed track sections appeared to be 10 to 50 meters long. [SRL95l] Our base-case model assumes $\lambda = 1$ failure/km, for a mean inter-failure distance of $1/\lambda = 1$ kilometer, more than twenty times the failure length observed in Kobe. We thus interpret the Poisson process as placing short runs of damaged track, as opposed to placing individual failures. As the failed runs are short, the correlation between the failure status of any two points along the track drops off quickly with the distance between them.

Shaking risk

The initial model does not include the probability that an earthquake shakes a moving Shinkansen train from its tracks, even though the Kobe experience suggests this is a realistic possibility. (Shaking during the Kobe quake derailed several parked, conventional trains. [SRL95l]) While we would prefer to include crash risk from shaking as well as crash risk from failed track in our calculations, modeling shaking risk is a formidable challenge. How does this risk depend on epicentral distance? Does this risk evolve in time, or in distance traveled? How does this risk depend on track curvature? Studies of track-wheel dynamics provide some insight, but on the whole these are unanswered questions in the railroad engineering and civil engineering literatures. [Kun72a] [Kun72b]

In addition, crash risk from shaking and crash risk from failed track are not mutually exclusive. That is, if A denotes the event "train crashes due to shaking," B denotes the event "train crashes due to failed track," and C denotes the event "train crashes from either shaking or track failure," then $\Pr(C) \neq \Pr(A) + \Pr(B)$, as simply summing these probabilities leads to over-counting. Using an overbar to denote the logical complement of an event, the correct expression for this probability is $\Pr(C) = \Pr(A) + \Pr\left(B|\overline{A}\right)\Pr\left(\overline{A}\right)$. In other words, we would need to model $\Pr\left(B|\overline{A}\right)$, the probability that a train is shaken from its tracks, given that the train does not crash at a track failure. Not only is modeling shaking risk itself a challenge, we see that the problem is thoroughly entwined with the crash risk from track failure problem.

Given such difficulties, we restrict the scope of our models and omit shaking risk. Accordingly, our analyses might underestimate the total crash risk.

Acceleration and deceleration

The simple crash risk model assumes trains cruise at speed v_c km/h, ignoring the fact that trains must spend some distance accelerating up to cruising speed after leaving a station and some distance decelerating to a stop before reaching a station. By assuming a train is always moving at full speed, the simple crash risk model overestimates a train's exposed stopping distance, and thus overestimates the probability that a train crashes at a track failure.

In order of decreasing accuracy, we propose three techniques to address the issue of train acceleration deceleration.

1. Assume the train speed at the instant of the quake, V, is a random variable. Make all calculations conditional on a given value of V, then integrate using the probability density function of V to remove this conditioning.

2. In the next section, we will explicitly address the issue of trains pausing in stations. Trains stopped in stations are not moving, and thus by definition they cannot crash. Suppose the average length of a station stop is t_{stat} seconds; the average time to accelerate up to full speed is t_{acc} seconds; and the average time to decelerate to a station stop is t_{dec} seconds. We have noted that assuming $t_{acc} = t_{dec} = 0$ overestimates the probability of a crash, as it implies that a train leaps to and from full speed instantly. This scenario offers us an upper bound on the crash probability. If we next replace the value of t_{stat} with the value $(t_{stat} + t_{acc} + t_{dec})$ in our model, and again assume $t_{acc} = t_{del} = 0$, we imply that a station protects a train during its acceleration and deceleration periods, as well as when the train is paused within it. (One might imagine an underground station, ensconced within a long tunnel.) This scenario underestimates the crash risk, offering us a lower bound on the crash probability. If these bounds are sufficiently tight, this time re-allocation trick can approximately solve the

acceleration-deceleration modeling issue.

3. Determine an appropriate mean speed, $E(V)$, and use this average in place of v_c throughout the model. Here, we use random variable V to denote a moving train's speed at random location between stations, not at a random instant in time. We know that $E(g(V)) \neq g(E(v))$ for an arbitrary function $g(V)$, though if g is convex, Jensen's inequality guarantees that $E(g(V)) \geq g(E(V))$. [Gal96a, p. 246] By checking the sign of the second partial derivative, one can show that the simple crash model is a concave function of v_c. (The plot in the upper right of Figure 2-3 displays this concavity.) Jensens's inequality implies that substituting $E(V)$ for v_c throughout the model will then overestimate the crash risk. However, this overestimate could be small: if the model is essentially linear in v_c, or if the bulk of the probability mass of random variable V rests upon the value v_c, then replacing v_c with $E(V)$ can offer a reasonable solution to the acceleration-deceleration problem.

As Shinkansen trains spend most of their traveling distance between stations at cruising speed, the regional crash model of Section 2.3 employs the third technique. In contrast, as conventional trains spend much more of their traveling distance between stations accelerating to and decelerating from cruising speed, our model for conventional trains of Section 2.4 employs the first technique.

Quake duration

The initial model assumes that all quake-induced track damage occurs in an instant. This simplification leads to a slight overestimate of the crash risk. Consider a Shinkansen train braking to a stop during an earthquake, and suppose that the track fails in exactly one location within the train's braking distance during the quake. If the train passes the location of this failure before the failure occurs, the train will brake to a safe stop and avoid crashing. That is, by assuming that all track failures occur at the instant of the quake, the initial model does not include the possibility that some track failures occur after the train has passed. Accordingly, the initial model

overestimates the probability of a crash.

We may modify the initial model to determine the magnitude of this overestimate. As before, we model track failures by a spatial Poisson process of rate λ over an interval of potential damage of length $2r$. Consider a train inside the region of potential damage with braking distance x. Let random variable N denote the number of track failures in the train's braking distance; N is Poisson with mean λx. Let random variables $S_1, S_2, \ldots, S_N$ denote the (unordered) locations of these failures; conditional on N, the S_i are uniformly distributed over the spatial interval $[0, x]$ and are mutually independent. We will also assume the train has speed v_c at the time of the quake.

We now introduce the quake's temporal duration into our model. Let time 0 denote the time the quake begins, and let t_d denote the time the quake ends. Let random variables $T_1, T_2, \ldots, T_N$ denote the times of the track failures at locations $S_1, S_2, \ldots, S_N$; we assume the T_i are uniformly distributed over the temporal interval $[0, t_d]$, mutually independent of one another and independent of the S_i. In other words, we assume failures times are independently and uniformly distributed across the duration of the quake, and do not depend on failure locations.

We now consider the northbound train operating inside the region of potential damage when the quake strikes. We assume that failures which occur after the train passes them do not threaten the train. (As before, we treat the train as a point, having no length; we shall subsequently address the small underestimate of risk that ensues.) In contrast, failures occurring before the train reaches them do threaten the train, as these track failures will cause the train to crash.

Assuming the train begins emergency braking at the instant of the earthquake ($\tau = 0$), simple physics indicates that the train will reach location S_i at time $h(S_i)$, where

$$h(s) = \frac{v_c - \sqrt{v_c^2 - 2a_e s}}{a_e}.$$

If $T_i < h(S_i)$, the track failure at location S_i occurs before the train reaches it, and thus threatens the train. As T_i is uniformly distributed over $[0, t_d]$, we can determine

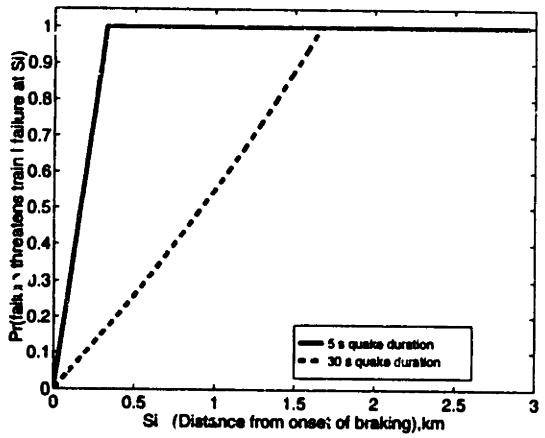

Figure 2-4: Pr(failure threatens train | failure at S_i) vs. S_i, for a $t_d = 5$ sec and a $t_d = 30$ sec quake, using base-case parameters (Table 2.1). Both curves increase, indicating that more distant failures are more likely to occur before the train reaches them. Consider the kink at 0.32 km in the $t_d = 5$ second curve. As the train travels 0.32 km while braking for 5 seconds, and as all quake damage occurs over these 5 seconds, it follows that all failures more than 0.32 km from the train must occur before the train reaches them.

the probability that $T_i < h(S_i)$.

$$\begin{aligned} \text{Pr(failure threatens train} \mid \text{failure at } S_i) &= \Pr(T_i < h(S_i)) \\ &= \frac{\min\{t_d, h(S_i)\}}{t_d} \end{aligned} \tag{2.4}$$

To clarify this model, one may imagine that a quake uses a spatial Poisson process to place track failures across the train's braking distance, then flips a coin for each track failure to decide whether the failure occurs before the train reaches the failure site (threatening) or after the train passes the failure site (not threatening). As the bias of the coin varies with the location, the resulting random spatial process of threatening failures forms a non-homogeneous spatial Poisson process (Figure 2-4).

We employ a rescaling technique commonly used with non-homogeneous Poisson processes by determining M, the expected number of threatening track failures in a train's braking distance. [Gal96a, p. 43] For simplicity, we will assume the train begins braking at the instant of the quake ($\tau = 0$) and spends its entire braking distance inside the region of potential damage. We first calculate M for our original

model, when we model the quake as having no temporal duration ($t_d = 0$). In this case, all failures occur at the instant of the quake, before the train arrives, and thus all failures are threatening failures.

$$M_{orig} = \int_{s=0}^{x} \lambda ds = \lambda x = \lambda \frac{v_c^2}{2a_e}.$$

We next calculate this same expected value for our revised model, when we permit the quake to have a temporal duration $t_d > 0$. In this case, M only counts the expected number of threatening failures, those which occur before the train reaches them.

$$\begin{aligned}
M_{dur} &= \int_{s=0}^{x} \lambda \Pr(\text{failure threatens train} \mid \text{failure at } s) ds \\
&= \lambda \int_{s=0}^{\frac{v_c^2}{2a_e}} \frac{\min\{t_d, \frac{v_c - \sqrt{v_c^2 - 2a_e s}}{a_e}\}}{t_d} ds \\
&= \lambda \left[\int_{s=0}^{(v_c t_d - a_e t_d^2/2)} \left(\frac{v_c - \sqrt{v_c^2 - 2a_e s}}{a_e t_d} \right) ds + \int_{s=(v_c t_d - a_e t_d^2/2)}^{\frac{v_c^2}{2a_e}} ds \right] \\
&= \lambda \left[\frac{3v_c^2 - 3v_c t_d a_e + a_e^2 t_d^2}{6a_e} \right]
\end{aligned}$$

(Of course, $M_{dur} = M_{orig}$ when $t_d = 0$.) Next, we calculate the ratio of these expected values as a function of t_d.

$$g(t_d) = \frac{M_{dur}}{M_{orig}} = \frac{3v_c^2 - 3v_c a_e t_d + t_d^2 a_e^2}{3v_c^2} \tag{2.5}$$

This ratio $g(t_d)$ indicates relative rate of threatening track damage between the two models. Using this relationship, we may model the effect of an earthquake's duration by an appropriate reduction of λ in the original model. That is, for $t_d > 0$, we can use the initial model to estimate the probability that at least one Shinkansen train derails due to earthquake-induced track failures by replacing λ with $g(t_d)\lambda$.

We find that for the base-case parameter values (Table 2.1), the appropriate reduction in λ needed to accommodate $t_d > 0$ is usually quite small (Figure 2-5). For example, if the track damage occurs over a twenty five second earthquake, the appropriate reduction is about 25%, as $g(25) = 0.745$. As discussed previously, such a

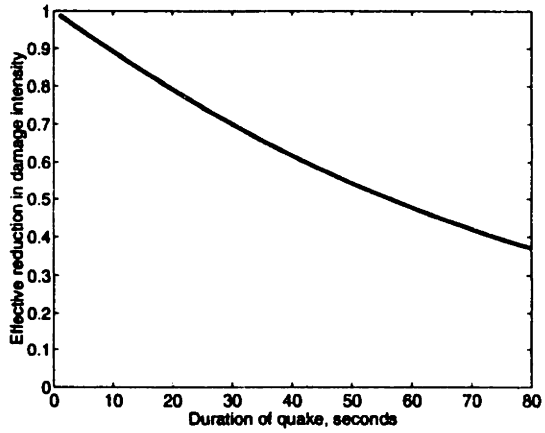

Figure 2-5: Track failure rate adjustment as a function of quake duration, from Equation 2.5. (This figure assumes base-case parameter values from Table 2.1.)

reduction in λ yields only about a 10% reduction in the probability of at least one crash (Table 2.2 and Figure 2-3). We conclude that the reduction in estimated crash risk that arises from modeling a quake's temporal duration is a second-order effect. This analysis justifies the simplifying $t_d = 0$ assumption of the initial model.

Independence of northbound and southbound trains

The initial model uses independent, identically distributed Bernoulli trials to model whether the northbound and southbound trains crash at track failures in the aftermath of an earthquake.

$$\rho = \Pr(\text{n'bound train crashes}) = \Pr(\text{s'bound train crashes}), \quad \text{i.i.d.}$$

However, as discussed earlier, these are not independent events. As the northbound and southbound track rest on a common viaduct, a single viaduct collapse will almost certainly destroy both tracks, endangering both northbound and southbound trains simultaneously. For example, given that the northbound train crashes, we know with certainty that there is at least one track failure within the region of potential damage, which in turn slightly increases the chance that the southbound train crashes. We see these two events are positively correlated, and thus our earlier calculation for the

chance of at least one train crash,

$$p = \Pr(\text{at least one train crashes}) = 1 - (1-\rho)^2, \tag{2.6}$$

is incorrect. We next determine the direction and the magnitude of the error we incur when we assume these events are independent, both analytically and via simulation.

Consider two positively correlated Bernoulli random variables, X_1 and X_2. Let $z = \Pr(X_1 = 1) = \Pr(X_2 = 1)$, and $r = \text{corr}(X_1, X_2)$. Let p_{ij} denote the joint probability mass function of these two random variables: $p_{ij} = \Pr(X_1 = i, X_2 = j)$, for $i = 0, 1$ and $j = 0, 1$. We can write a system of four equations in the four unknown p_{ij},

$$\begin{aligned} r = \frac{\text{cov}(X_1, X_2)}{\sqrt{\text{var}(X_1)}\sqrt{\text{var}(X_2)}} &= \frac{p_{11} - (p_{10} + p_{11})(p_{01} + p_{11})}{\sqrt{(p_{10} + p_{11}) - (p_{10} + p_{11})^2}\sqrt{(p_{01} + p_{11}) - (p_{01} + p_{11})^2}}, \\ z &= p_{10} + p_{11}, \\ z &= p_{01} + p_{11}, \\ 1 &= p_{00} + p_{01} + p_{10} + p_{11}, \end{aligned}$$

and use this system to solve for the p_{ij} in terms of r and z.

	$X_2 = 0$	$X_2 = 1$
$X_1 = 0$	$1 - 2z + rz + (1-r)z^2$	$z - rz - (1-r)z^2$
$X_1 = 1$	$z - rz - (1-r)z^2$	$rz + (1-r)z^2$

This joint probability mass function completely describes the sample space of X_1 and X_2, and properly accounts for their correlation.

In our simple crash risk model, we focus on the probability of at least one crash. In terms of our correlated Bernoullis, we may write

$$\Pr(X_1 + X_2 > 0) = 1 - p_{00} = (2-r)z - (1-r)z^2. \tag{2.7}$$

Note that if $r = 0$, the expression reduces to $2z - z^2 = 1 - (1-z)^2$, matching the

z	r	$\Pr(X_1 + X_2 > 0)$ true	$\Pr(X_1 + X_2 > 0)$ approximate	difference
0.5	0.50	0.6250	0.7500	-0.1250
	0.10	0.7250	0.7500	-0.0250
	0.01	0.7475	0.7500	-0.0025
0.9	0.50	0.9450	0.9900	-0.0450
	0.10	0.9810	0.9900	-0.0090
	0.01	0.9891	0.9900	-0.0009

Table 2.4: The error that results from treating two dependent Bernoulli random variables as if they were independent. Following the notation of the text, $z = \Pr(X_1 = 1) = \Pr(X_2 = 1)$ and $r = \text{corr}(X_1, X_2)$. The "true" column uses Equation 2.7, and the "approximate" column uses Equation 2.6.

form of the simple crash risk model which assumed independence of the northbound and southbound trains (Equation 2.6). Equation 2.7 allows us to determine the error that accrues when we treat two identically distributed correlated Bernoulli random variables as if they were independent, as a function of their correlation, r, and their success probability, z.

$$\begin{aligned} \text{error} &= \Pr(X_1 + X_2 > 0, \text{ true, via Eqn 2.7}) \\ &\quad -\Pr(X_1 + X_2 > 0, \text{ approximate, via Eqn. 2.6}) \\ &= [(2-r)z - (1-r)z^2] - [1 - (1-z)^2] \\ &= -r(z - z^2) \end{aligned} \tag{2.8}$$

If X_1 and X_2 are positively correlated, treating them as if they were independent causes us to overestimate $\Pr(X_1 + X_2 > 0)$. The magnitude of this error increases linearly with r and decreases quadratically with $|z - 1/2|$. In all cases, the maximum absolute error of this approximation never exceeds $1/4$, obtained when $r = \pm 1$ and $z = 1/2$. Table 2.4 presents an illustrative numerical example.

We may apply these results to our simple crash model. We determined that, for our base-case parameter values, the probability of a northbound crash and the probability of a southbound crash were roughly one third ($z = 0.33$). Depending on the correlation between these two events, the error of Equation 2.6 could range from 0 (if $r = 0$) to about $2/9$ (if $r = 1$).

We next conducted 10,000 Monte Carlo experiments based on our crash risk model with our base-case parameters to determine the correlation between the crash status of the northbound train and the crash status of the southbound train (Table 2.5). In each simulation, we generated a Poisson number of track failures and placed them uniformly across a region of potential damage, represented by the interval $[-r, r]$. We placed the northbound train at position Y_n uniformly at random over the interval $[-(r+x), r]$, and the southbound train at position Y_s uniformly at random over the interval $[-r, (r+x)]$. (Thus, this simulation allows for trains to enter and exit the region of potential damage while braking; this has negligible effect.) If there were any failures in the interval $[Y_n, Y_n + x]$, our simulated northbound train crashes; if there were any track failures in the interval $[Y_s - x, Y_s]$, our simulated southbound train crashes.

From these Monte-Carlo experiments, we observe that the correlation between the crash status of the northbound and southbound trains is very small: $r = 0.012$. We believe this follows from the high expected number of track failures within the region of potential damage, so the northbound and southbound train seldom crash at the same failure. Equation 2.8 indicates that we incur an error of about $(0.012)(0.33 - 0.33^2) = 0.003$ by assuming independence of the northbound and southbound train. The observed difference between Pr(at least one crash) obtained from our simple crash risk model and Pr(at least one crash) obtained from this suite of Monte Carlo experiments is 0.007, clearly of the same order of magnitude as the predicted 0.003 error.

Due to the large distances between Shinkansen trains and the large number of expected track failures within the region of potential damage, our simulation experience suggests that the crash risk status of the northbound and southbound train are effectively independent, justifying this assumption in the simple crash risk model.

Results from Monte-Carlo simulation		
	$X_1 = 0$	$X_1 = 1$
$X_2 = 0$	0.44001	0.22146
$X_2 = 1$	0.22115	0.11738
Observed correlation $r = 0.012$		

Results from simple crash risk model		
	$X_1 = 0$	$X_1 = 1$
$X_2 = 0$	0.44735	0.22149
$X_2 = 1$	0.22149	0.10967
Model assumes correlation $r = 0$		

Table 2.5: Monte Carlo simulation results for the simple crash risk model. The table on the left provides the results of 10,000 Monte-Carlo experiments using the simple crash risk model and the base-case parameters (Table 2.1). The table on the right calculates these same probabilities analytically, assuming independence of the crash status of the northbound (X_1) and the southbound (X_2) trains.

2.3 Regional crash risk model

Though the simple crash risk provides some powerful insight, this model also ignores some issues. For example, the simple crash risk model does not account for differing train densities and speeds in different parts of a rail line. The simple model only considers moving trains, despite the fact that trains spend several hours each night out of service, and several minutes of each hour paused in stations. The simple model does not consider that track structures along the line may vary in their seismic vulnerability. The simple model ignores train length, and it does not consider that trains may enter or exit the region of potential damage while braking. We next weave these factors into our previous framework to form our "regional crash risk model." We will see that most of these refinements have small effect. In contrast, we will see that the regional variation in train density will have a significant effect on our results.

2.3.1 Formulation of regional crash risk model

Damage model

In our original model, we employ a homogeneous spatial Poisson process to model the distribution of track failures within the region of potential damage (Figure 2-2). However, an earthquake's force generally decreases with distance from the quake's

epicenter. (Local variations in local geological conditions may amplify or dampen local earth motion, obscuring this phenomenon to some degree.) As it seems reasonable to assume that the likelihood of track damage decreases as the track is subjected to less shaking, we would also expect that the likelihood of track damage generally decreases with increasing epicentral distance.

Seismologists have conducted statistical studies to determine how various measures of shaking attenuate with epicentral distance. One such measure is peak ground acceleration (PGA), which describes the maximum ground acceleration in the horizontal plane observed at a location over an earthquake. Another is peak spectral acceleration (PSA), which filters PGA through a simple model of a damped inverted pendulum to describe the expected maximum acceleration of a structure in response to the ground motion. After examining records from 400 off-shore Japanese earthquakes, Kawashima, Aizawa, and Takahasi [KAT84] propose that both median PGA and SA are log-normally distributed, with medians proportional to $(a + x)^{-b}$, where x represents epicentral distance, and a and b are fitted constants. For the Japanese quakes in their study, they suggest $a = 30$ km and $b \approx 1.2$. Most roughly, then, peak shaking drops off somewhat faster than the reciprocal of distance. Though such attenuation studies provide valuable insight into the relationship between distance and shaking, they do not address the relationship between shaking and damage.

Papadimitriou and Veneziano introduce a model for seismic damage to the Shinkansen track in their analysis of Shinkansen early warning systems. [PV95] The Shinkansen track runs on a concrete viaduct, supported every ten meters by concrete columns. Papadimitriou and Veneziano use these inter-column sections as the basic element of their analysis. They assume a track section will fail if the peak spectral acceleration at that location exceeds a fixed threshold. For each ten meter section, they use the Kawashima model to determine the probability that the PSA exceeds this threshold and causes the section to fail (Figures 2-6, 2-7, and 2-8). They employ a two-state discrete Markov model to describe the number of track sections within a run of damaged track and the number of track sections within a run of undamaged track.

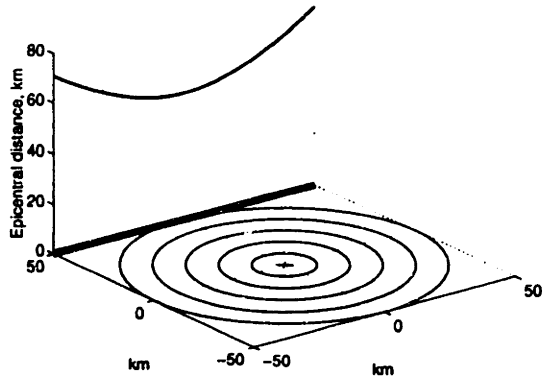

Figure 2-6: Epicentral distance along Shinkansen track, here represented by the bold black line. In this example, the track is 50 km from the epicenter of an earthquake, represented by an "x" at the origin. The concentric circles represent points 10, 20, 30, 40, and 50 km from the epicenter, respectively.

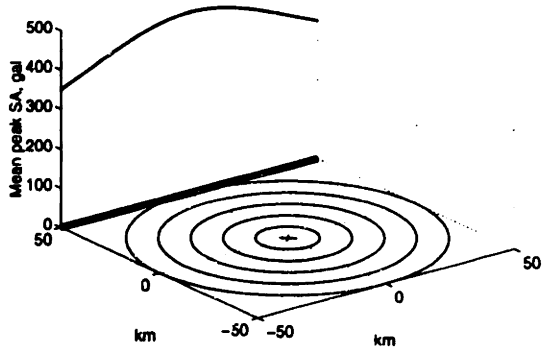

Figure 2-7: Peak spectral acceleration of track structure using the Kawashima model. [KAT84] In this example, the track is 50 km from a 7.2 M earthquake (similar to the situation in the Kobe quake); the soil type is "normal;" and the period of the viaduct columns is 0.4 seconds. [PV95]

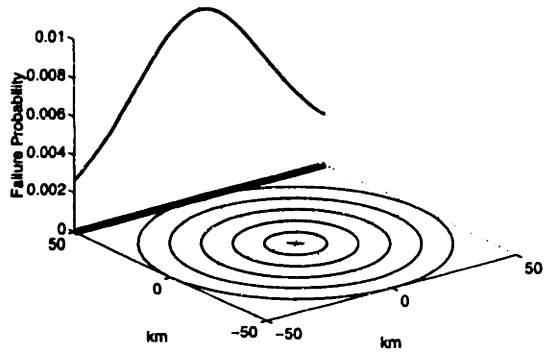

Figure 2-8: Probability that a ten meter inter-column track section fails, from the Papadimitriou and Veneziano track damage model. [PV95] In this example, the track is 50 km from a 7.2 M earthquake; the soil type is "normal;" the period of the viaduct columns is 0.4 seconds; and a section fails if PSA exceeds 1850 gals, or 1.85 times the force of gravity.

We choose to avoid the complexities and uncertainties inherent in such detailed engineering analyses. However, we do wish to include the effect of epicentral distance in our regional model. Observing the Papadimitriou and Veneziano track damage model, we note that the effective damage intensity is essentially triangular (Figure 2-8). This observation allows us to capture, albeit crudely, the effect of epicentral distance on intensity of the spatial track failure process. Instead of a constant λ, we allow this damage intensity rate to vary according to a triangular function $\lambda(s)$. We will denote the point on the track closest to the epicenter as the origin of our coordinate system (the point under the peak of the triangle), and maintain this convention throughout the entire model (Figure 2-9).

$$\lambda(s) = \begin{cases} \lambda_m \left(1 - \frac{|s|}{r}\right), & -r \leq s \leq r \\ 0, & \text{else.} \end{cases} \tag{2.9}$$

We wish to make the base-case earthquake of our regional model equivalent to the base-case earthquake of our simple crash risk model. This suggests we should match the length of the region of potential damage, and match the expected number of track failures in this region (that is, match the base of Figure 2-9 to the base of Figure 2-2,

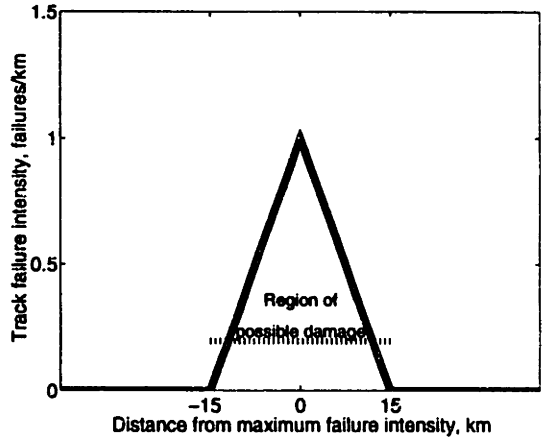

Figure 2-9: A triangular track failure intensity function captures the first-order effects of epicentral distance. The origin denotes the point along the track closest to the quake epicenter, which is also the point of maximum failure intensity.

and match the area of triangle to the area of the rectangle). In the simple crash risk model, the expected number of track failures in the region of potential damage is $2\lambda r$; in the regional model, the expected number is $\lambda_m r$. Thus, to match the expected number of failures, we set $\lambda_m = 2\lambda$.

As an aside, we note that while our base-case track failure intensity function (Figure 2-9) and the Papadimitriou and Veneziano track failure intensity function (Figure 2-8) share the same maximum intensity (1 failure/km vs. 0.01 failure/10m), we assume a much smaller region of potential damage (30 km, vs. about 100 km).

It will be helpful to define $m(s)$, the cumulative track failure intensity.

$$\begin{aligned} m(s) &= \int_{-\infty}^{s} \lambda(u)du \\ &= \begin{cases} 0, & s < -r \\ \frac{\lambda_m(s+r)^2}{2r}, & -r \leq s < 0 \\ \frac{\lambda_m r}{2} + \frac{\lambda_m(2r-s)(s)}{2r}, & 0 \leq s < r \\ \lambda_m r, & r < s \end{cases} \end{aligned} \tag{2.10}$$

With this definition, the expected number of track failures in the interval $[a, b]$ is $m(b) - m(a)$. Using the standard rescaling approach for non-homogeneous Poisson

processes [Gal96a, p. 43], the probability of at least one track failure in the interval $[a, b]$ is $1 - \exp(-[m(b) - m(a)])$.

Train speed

As Shinkansen trains spend most of their distance between stations traveling at cruising speed, we will treat a train's speed as constant, and use the mean speed $E(V)$ for this value. We next calculate the train's mean speed over an inter-station run. To do so, we assume that stations are spaced a constant D kilometers apart, and trains accelerate and decelerate with constant rate a_n m/s/s. Thus, a train spends $v_c^2/2a_n$ meters accelerating, $v_c^2/2a_n$ meters decelerating, and the remaining $D - v_c^2/a_n$ meters at full speed (Figure 2-10). With this model, at distance $d \in [0, D]$ from a station, a train has speed $V(d)$.

$$V(d) = \begin{cases} \sqrt{2a_n d}, & 0 \le d < v_c^2/2a_n \\ v_c, & v_c^2/2a_n \le d < D - v_c^2/2a_n \\ \sqrt{2a_n(D-d)}, & v_c^2/2a_n \le d \le D. \end{cases} \tag{2.11}$$

Let random variable V denote the speed of a train located at a random spot between two stations. We use the function $V(d)$ to determine the probability density function of this random variable (Figure 2-11),

$$f_V(v) = \begin{cases} \frac{2v}{a_n D}, & 0 \le v < v_c \\ \frac{D - v_c^2/a_n}{D}, & v = v_c \\ 0, & \text{else,} \end{cases} \tag{2.12}$$

and integrate to obtain v_m, the spatially-averaged mean speed.

$$p_{cruise} = \frac{D - v_c^2/a_n}{D} \tag{2.13}$$

$$v_m = E(V) = v_c \left[p_{cruise} + \frac{2}{3}(1 - p_{cruise}) \right] \tag{2.14}$$

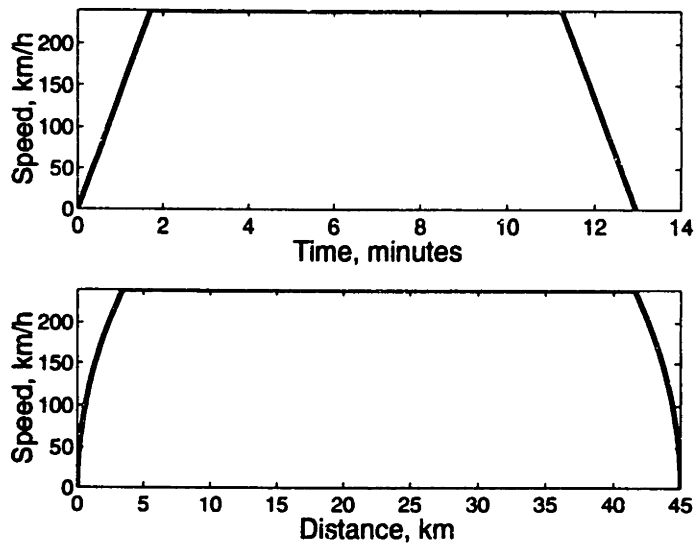

Figure 2-10: Speed versus time and distance for a typical Shinkansen inter-station run. This plot assumes constant linear acceleration and deceleration with rate a_n = 0.65 m/s/s, a cruising speed of v_c = 240 km/h, and an inter-station distance of 45 kilometers.

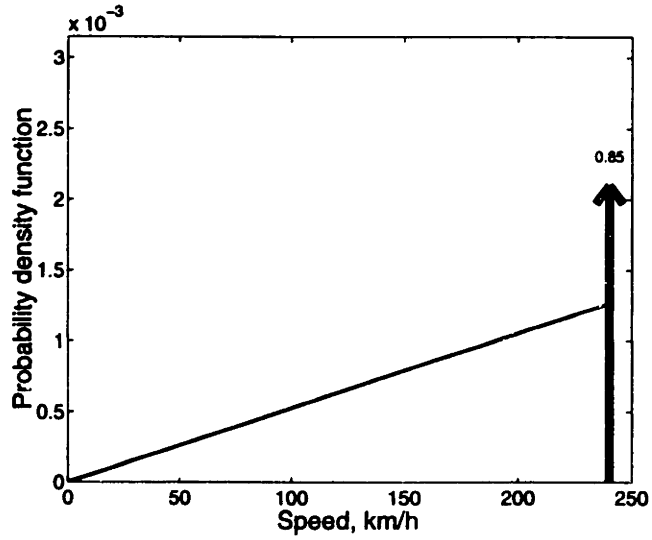

Figure 2-11: Mixed probability density function for the speed of a moving Shinkansen at a random location between two stations. The arrow represents a probability mass at cruising speed v_c. This plot assumes constant linear acceleration and deceleration with rate a_n = 0.65 m/s/s, a cruising speed of v_c = 240 km/h, and an inter-station distance of 45 kilometers.

Train length

The simple crash risk model treated a Shinkansen train as a point. In reality, trains are usually 250 meters long. In the regional crash risk model, we shall assume all Shinkansen trains are a constant positive length l_{train} meters. This small change ripples through the entire model, but has the most significant effect on the exposed stopping distance.

Consider the location of the rear of the last car of a moving train when an earthquake strikes, and consider the location of the front of the first car when this train brakes to a safe stop. Any failures between these two locations would endanger the train. Accordingly, we see that including the length of the train serves to increase the exposed stopping distance by l_{train}.

$$x = \begin{cases} l_{train} + v_m\tau + \frac{v_m^2}{2a}, & \text{if } \tau \geq 0 \\ l_{train} + \frac{[\max\{0, v_m + a\tau\}]^2}{2a}, & \text{if } \tau < 0 \end{cases} \tag{2.15}$$

This adjustment is a modeling trick to capture the first-order effects of train length on the exposed stopping distance. After this change, we continue to treat the train as a point for the rest of our model.

In the next section, we will argue that tunnels shield trains from some of their crash risk exposure. For precision, we will now refer to x from Equation 2.15 as the "potentially exposed stopping distance."

Tunnels

The simple crash risk model does not account for different seismic vulnerabilities of different track structures. The most significant of these differences involves tunnels. Our conversations with railroad engineers suggest that track in tunnels is far less likely to fail in an earthquake than track on viaduct. Evidence from Kobe supports this assumption, as the Kobe tunnel appears to have fortuitously protected 40 km of track from damage. (Section 2.3.3 specifically considers the Kobe damage.) While we suspect the probability of track failure in a tunnel is not zero, we do believe it is quite

small. We thus make the following simplifying assumption for our regional crash risk model: an earthquake causes no track failures inside tunnels. This assumption of zero-risk in tunnels (versus very-small-risk) has only a minor effect on the distribution of the number of trains that crash in the aftermath of a quake. This assumption could have a large effect on the distribution of the number of on-board passengers that perish in such crashes. Though rare, train crashes in tunnels may be exceptionally deadly for passengers and crew: fire may consume all available oxygen, and narrow tunnels limit access and hinder rescue attempts. Such concerns are very important, but estimating fatalities is beyond the scope of our models.

Tunnels of various lengths are scattered across the Shinkansen lines. Given the precision of our models, it makes little sense to consider each tunnel in detail. Instead, we approximate the tunnel situation with a stochastic proxy. We use an alternating renewal process to describe the alternating stretches of above-ground and in-tunnel track. Let random variable C_i denote the length of a randomly chosen stretch of tunnel, and let random variable Q_i denote the length of a randomly chosen stretch of above-ground track, with cumulative density functions $F_C(c)$ and $F_Q(q)$, respectively. We model the rail line by assuming it begins with an above-ground stretch of track of length Q_1, followed by a tunnel of length C_1, followed by an above-ground stretch of length Q_2, a tunnel of length C_2, and so on. We also assume trains travel at the same speed in tunnels as on above-ground track.

With this proxy model for tunnels, the chance a randomly chosen moving train is inside a tunnel at a randomly chosen instant is $E(C)/(E(C)+E(Q))$, and the chance such a train is on above-ground track is $E(Q)/(E(C)+E(Q))$. Now suppose the train is in a tunnel at this randomly chosen instant. We can use a random-incidence calculation [Dra88, p. 149] to determine the probability that the residual tunnel length (that is, the distance from the front of the train to the end of the tunnel ahead of it) exceeds x.

$$\Pr(\text{residual tunnel distance} > x) = 1 - \frac{1}{E(C)} \int_0^x (1 - F_C(s))ds$$

Similarly, if this train happens to be on an above-ground stretch at this random instant, we can determine the probability that the residual above-ground track length exceeds x.

$$\Pr(\text{residual above-ground distance} > x) = 1 - \frac{1}{E(Q)} \int_0^x (1 - F_Q(s))ds$$

Thus, at a random instant, a randomly chosen moving train finds itself in one of three situations:

"Tunnel ahead for x" The train is in a tunnel, and the tunnel extends in front of the train for x meters or more, or

"Above-ground ahead for x" The train is on above-ground track, and this stretch extends in front of the train for x meters or more, or

"Mixed ahead for x" The train spends some of the next x meters in tunnels and some of the next x meters on above-ground track.

We may compute the probabilities of these three mutually exclusive, collectively exhaustive states.

$$\Pr(\text{Tunnel ahead for } x) = \frac{E(C) - \int_0^x (1 - F_C(s))ds}{E(C) + E(Q)} \tag{2.16}$$

$$\Pr(\text{Above-ground ahead for } x) = \frac{E(Q) - \int_0^x (1 - F_Q(s))ds}{E(C) + E(Q)} \tag{2.17}$$

$$\Pr(\text{Mixed ahead for } x) = \frac{\int_0^x (1 - F_C(s))ds + \int_0^x (1 - F_Q(s))ds}{E(C) + E(Q)} \tag{2.18}$$

Let random variable Y denote the length of open track in the x meters ahead of a randomly chosen moving train at a random instant. Clearly, if the train is in a tunnel and the tunnel continues for the next x meters, then $Y = 0$. Similarly, if the train is on above-ground track, and the above-ground track continues for at least x meters, then $Y = x$. If neither of these two situations hold, then the train spends some of the next x meters in tunnels and some of the next x meters on above-ground track. By symmetry considerations, in this case Y is uniformly distributed on $[0, x]$. We thus obtain the mixed probability density function of Y.

$$f_Y(y) = \begin{cases} \Pr(\text{Tunnel ahead for } x), & y = 0 \\ \frac{1}{x} \cdot \Pr(\text{Mixed ahead over } x), & 0 < y < x \\ \Pr(\text{Above-ground ahead for } x), & y = x \end{cases}$$

The fate of an "at-risk" train

Assume an earthquake strikes, and consider a randomly-chosen moving train that will spend some or all of its potentially exposed stopped distance inside the region of potential damage. We use the description "at-risk" to describe such a train, as it is at risk to crash at a track failure. We now determine the probability that a randomly chosen at-risk train crashes at a track failure in the aftermath of an earthquake.

With the origin of our coordinate system still on the point of maximum track failure intensity, let random variable S denote the position of this at-risk train when it begins to brake. We assume the position of the earthquake and the position of the train at this moment are independent, so S is uniformly distributed on the interval $[-(r+x), r]$. We have extended this region by x to account for a train entering the region while braking.

We have defined random variable Y, which describes the amount of open track in front of a randomly chosen train within its potentially exposed braking distance at a random instant. We have defined random variable S, which describes the position of an at-risk moving train relative to the region of potential damage. Conditional on these two random variables, we can determine the crash probability of a randomly chosen at-risk moving train via the cumulative track failure intensity function, $m(s)$, from Equation 2.10.

$$P(s,y) = \Pr\left(\begin{matrix}\text{train}\\ \text{crashes}\end{matrix} \,\middle|\, \begin{matrix}\text{train}\\ \text{at-risk}\end{matrix}, \begin{matrix}\text{train}\\ \text{moving}\end{matrix}, S = s, Y = y \right) = 1 - e^{-[m(s+y)-m(s)]} \quad (2.19)$$

We integrate twice to remove the conditioning on Y and S.

$$\begin{aligned}
\Pr\left(\begin{array}{c}\text{train}\\ \text{crashes}\end{array}\middle|\begin{array}{l}\text{train}\\ \text{at-risk},\end{array}\begin{array}{l}\text{train}\\ \text{moving}\end{array}\right) &= \int_{y=0}^{x}\int_{s=-(r+x)}^{r} P(s,y) f_S(s) f_Y(y)\,ds\,dy \\
&= \Pr\left(\begin{array}{c}\text{tunnel}\\ \text{ahead for } x\end{array}\right)\cdot 0 \\
&+ \Pr\left(\begin{array}{c}\text{mixed}\\ \text{ahead for } x\end{array}\right)\cdot\left(\tfrac{1}{2r+x}\right)\left(\tfrac{1}{x}\right)\int_{y=0}^{x}\int_{s=-(r+x)}^{r} P(s,y)\,ds\,dy \\
&+ \Pr\left(\begin{array}{c}\text{above-ground}\\ \text{ahead for } x\end{array}\right)\cdot\left(\tfrac{1}{2r+x}\right)\int_{s=-(r+x)}^{r} P(s,x)\,ds
\end{aligned} \tag{2.20}$$

Station stops

Using a train schedule, we can determine $p_{station}$, the fraction of a train's operating day spent in station stops. Thus, at a random instant across the operating day, the probability that a randomly chosen train is in a station is $p_{station}$. As we restrict our model to moving trains, trains paused in stations are immune from crashing.

$$\begin{aligned}
p_{crash} &= \Pr\left(\begin{array}{c}\text{train}\\ \text{crashes}\end{array}\middle|\begin{array}{l}\text{train}\\ \text{at-risk},\end{array}\begin{array}{l}\text{train}\\ \text{moving}\end{array}\right)\cdot\Pr\left(\begin{array}{c}\text{train}\\ \text{moving}\end{array}\right) \\
&+ \Pr\left(\begin{array}{c}\text{train}\\ \text{crashes}\end{array}\middle|\begin{array}{l}\text{train}\\ \text{at-risk},\end{array}\begin{array}{l}\text{train in}\\ \text{station}\end{array}\right)\cdot\Pr\left(\begin{array}{c}\text{train in}\\ \text{station}\end{array}\right) \\
&= (1-p_{station})\cdot\Pr\left(\begin{array}{c}\text{train}\\ \text{crashes}\end{array}\middle|\begin{array}{l}\text{train}\\ \text{at-risk},\end{array}\begin{array}{l}\text{train}\\ \text{moving}\end{array}\right)
\end{aligned} \tag{2.21}$$

The number of at-risk trains

In our simple crash risk model, we only considered an earthquake that struck when the trains were running. We also assumed the headways between successive trains were larger than the region of potential damage ($\Delta > 2r$), so at most one moving northbound and one moving southbound train could be within the region of potential damage of an earthquake. We relax both of these assumptions in our regional crash risk model.

For a quake that strikes when the trains are in-service, let random variables N_1 and N_2 denote the number of northbound and southbound at-risk trains, respectively. Recall that an "at-risk" train is a train that spends some or all of its potentially exposed stopping distance inside the region of potential damage. We assume trains

have constant headways of Δ kilometers. (This is a modeling convenience; we realize that actual headways vary slightly due to congestion, station stops, and unforeseen circumstances.) The region of potential damage is $2r$ long, and we "extend" it by x kilometers to account for trains entering the region while braking. (This extension has a small effect, as $x \ll \Delta$.) We may obtain the distribution of the N_i by straightforward spatial considerations; there must be $\left\lfloor \frac{2r+x}{\Delta} \right\rfloor$ or $\left\lfloor \frac{2r+x}{\Delta} \right\rfloor + 1$ at-risk northbound or southbound trains in the region of potential damage. Let A denote the event, "the quake strikes when the trains are in service."

$$\Pr(N_i = k|A) = \begin{cases} 1 - \frac{2r+x}{\Delta} + \left\lfloor \frac{2r+x}{\Delta} \right\rfloor, & k = \left\lfloor \frac{2r+x}{\Delta} \right\rfloor \\ \frac{2r+x}{\Delta} - \left\lfloor \frac{2r+x}{\Delta} \right\rfloor, & k = \left\lfloor \frac{2r+x}{\Delta} \right\rfloor + 1 \\ 0, & \text{else} \end{cases} \tag{2.22}$$

We assume N_1 and N_2 are independent, and we obtain the distribution of their sum via convolution.

$$(N_{tot}|A) = (N_1|A) + (N_2|A)$$

$$\begin{aligned} \Pr(N_{tot} = k|A) &= \sum_{j=\left\lfloor \frac{2r+x}{\Delta} \right\rfloor}^{\left\lfloor \frac{2r+x}{\Delta} \right\rfloor + 1} \Pr(N_1 = k|A)\Pr(N_2 = k - j|A), \\ &\quad \text{for } k = 2\left\lfloor \frac{2r+x}{\Delta} \right\rfloor, 2\left\lfloor \frac{2r+x}{\Delta} \right\rfloor + 1, 2\left\lfloor \frac{2r+x}{\Delta} \right\rfloor + 2 \end{aligned} \tag{2.23}$$

The time at which the quake strikes

The JR East passenger rail system operates 18 hours each day, from six in the morning until midnight. Assuming an earthquake strikes at a random time of day, there is 6/24 chance it strikes in the early morning, when the system is not in service. If a quake strikes when the system is not in service, no moving trains face a risk of crashing. We scale the distribution of $(N_{tot}|A)$ downwards accordingly to account for

the time at which the quake strikes.

$$\Pr\left(N_{tot}=k \,\middle|\, \begin{matrix}\text{quake strikes}\\ \text{any time}\end{matrix}\right) = \begin{cases} \frac{6}{24}, & k=0 \\ \frac{18}{24}\Pr(N_{tot}=k\,|A) & k=2\left\lfloor\frac{2r+x}{\Delta}\right\rfloor, 2\left\lfloor\frac{2r+x}{\Delta}\right\rfloor+1, 2\left\lfloor\frac{2r+x}{\Delta}\right\rfloor+2 \\ 0, & \text{else} \end{cases} \tag{2.24}$$

Mean number of crashing trains

Equation 2.21 describes the probability that a randomly chosen at-risk train crashes at a track failure in the aftermath of an earthquake. Equation 2.24 describes the distribution of the number of at-risk trains given a quake which strikes at a random time of day. We can combine these two expressions to determine the mean number of trains that crash at track failures in the aftermath of an earthquake striking at a random time of day.

$$E\left(\begin{matrix}\text{train}\\ \text{crashes}\end{matrix}\,\middle|\,\begin{matrix}\text{quake strikes}\\ \text{any time}\end{matrix}\right) = p_{crash}\cdot E\left(N_{tot}\,\middle|\,\begin{matrix}\text{quake strikes}\\ \text{any time}\end{matrix}\right) \tag{2.25}$$

Equation 2.25 provides a simple expression for the expected number of crashing trains. It is not simple, however, to determine the distribution of the number of crashing trains. The crash status of each at-risk train is not independent of the crash status of the others, due to the triangular track failure intensity function and protective effect of tunnels. Accordingly, we cannot assume that the distribution of the number of crashing trains, given N_{tot}, is Binomial(N_{tot}, p_{crash}). Thus, the regional crash risk model provides only an estimate of the first moment of the number of moving trains which crash due to track damage in the aftermath of an earthquake.

Equations 2.9–2.25 comprise our regional crash risk model. In the next section, we apply this model to the Tohoku line as a case study.

2.3.2 Case study: the Tohoku Shinkansen

After considering the operations of the Tohoku Shinkansen, we opt to partition this line into three regions for modeling purposes: the 160 km stretch from Tokyo to Nasu, the 190 km stretch from Nasu to Sendai, and the 180 km stretch from Sendai to Morioka. We allow model parameters to vary between regions, but we assume constant parameters within each region.

Our model involves twelve parameters for each region (Table 2.6). Seven parameters deal with the trains, stations, and schedule: l_{train}, the train length; a_n, the non-emergency acceleration and deceleration rate of a Shinkansen; a_e, the emergency braking rate of a Shinkansen; $p_{station}$, the fraction of a train's operating day spent paused stations; D, the average inter-station distance; v_c, the cruising speed of a Shinkansen; and Δ, the headway between successive trains. For our base-case, we assume all trains in each region are $l_{train} = 250$ meters long, have acceleration and deceleration rates of $a_n = 0.65$ m/s/s, and have emergency braking rates of $a_e = 0.75$ m/s/s. The Shinkansen time-table provides the distance between stations, the transit time between stations, the number of stations in each region, and the number of trains per day. [Eas95] We assume a Shinkansen station stop lasts 45 seconds. Working from these data, we assume train cruising speeds of $\{200, 240, 240\}$ km/h, average inter-station distances of $\{40, 50, 50\}$ km, train headways of $\{60, 90, 170\}$ km, and $p_{station}$ values of $\{0.05, 0.06, 0.07\}$ for the Tokyo–Nasu, Nasu–Sendai, and Sendai–Morioka regions, respectively.

Two parameters for each region address the distribution of tunnels and above-ground track: $F_C(c)$, the length of a randomly chosen tunnel; and $F_Q(q)$, the length of a randomly chosen stretch of above-track (Figure 2-12). We calculate these distributions from a tunnel map of the Tohoku Shinkansen line.

Three parameters deal with the assumed earthquake itself: λ_m, the maximum track failure intensity rate; r, the radius of the region of potential damage; and τ, the delay time (or lead-time) between the time of the quake and the railroad's reaction to the quake. To match the base-case earthquake from the simple crash risk model, we assume $\lambda_m = 1.0$ track failures/km, $r = 15$ km, and $\tau = 0$ seconds. We intend

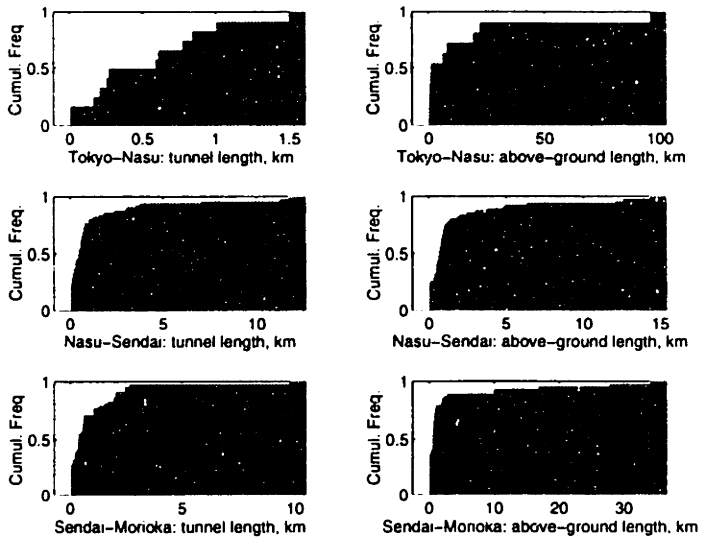

Figure 2-12: The cumulative distribution functions for tunnel length and above-ground length, by region of the Tohoku line.

these values to resemble the Kobe disaster, where a large quake struck close to the Shinkansen track without warning.

Crash risk per at-risk train, by region

With these values, we can use Equations 2.9–2.21 to calculate the probability that a randomly chosen at-risk train crashes at a quake-induced track failure, by region of the quake's occurrence (Figure 2-13). The similarity of p_{crash} across the three regions is striking. We may attribute most of the small inter-regional differences to differences in tunnel coverage and operating speeds. The difference between the Nasu–Sendai region and the Sendai–Morioka region is entirely due to the presence of more protective tunnels in the Nasu–Sendai region. The Tokyo–Nasu region has even fewer tunnels than the Nasu–Sendai region, but the slightly slower operating speed (and accompanying reduction in the potentially exposed braking distance) near Tokyo compensates somewhat for the increased risk exposure arising from less tunnel coverage.

Given the precision of our model, we do not believe such slight regional differences are meaningfully significant. We might coarsely summarize our results as follows:

Parameter		Tokyo–Nasu	Nasu–Sendai	Sendai–Morioka
Train length	l_{train}	250 m		
Accel/decel rate	a_n	0.65 m/s/s		
Emergency braking rate	a_e	0.75 m/s/s		
% time in stations	$p_{station}$	0.05	0.06	0.07
Train cruising speed	v_c	200 km/h	240 km/h	240 km/h
Train headway	Δ	60 km	90 km	170 km
Inter-station distance	D	40 km	50 km	50 km
Max failure intensity	λ_m	1 track failure/km		
Radius damage region	r	15 km		
Tunnel length	$F_C(c)$	(See Figure 2-12)		
Above-ground length	$F_Q(q)$	(See Figure 2-12)		

Table 2.6: Parameters for the regional crash risk model.

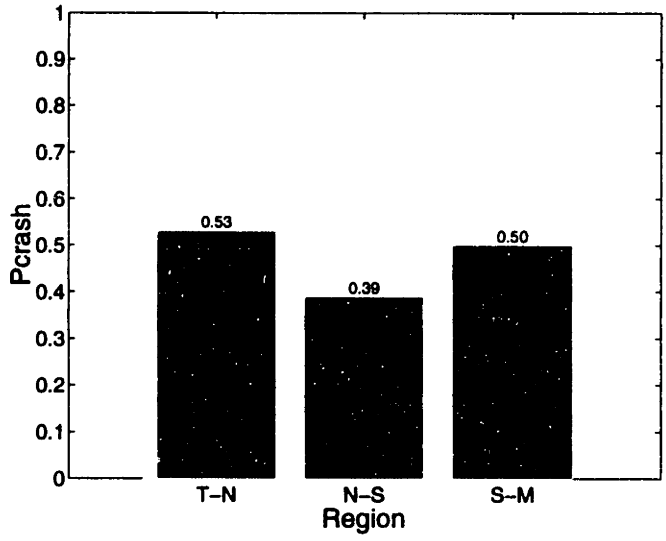

Figure 2-13: Crash risk per at-risk train, by region of the Tohoku line. This probability indicates the chance that a randomly chosen moving train that spends some or all its potentially exposed braking distance inside the region of potential damage crashes at a track failure in the aftermath of an earthquake.

a randomly chosen train that spends some or all of its potentially exposed braking distance inside the region of potential damage faces about a 0.5 probability of crashing at a track failure.

How does this result compare to the simple crash risk model? We next examine the differences between p_{crash} from these two models. To do this, we will start with the assumptions of the simple crash risk model and successively refine them until we reach the regional crash risk model, noting how p_{crash} changes as successive refinements are made (Figure 2-14). In the simple crash risk model, we assumed that a moving train braking to a stop within the region of potential damage faced a base-case probability of $1 - \exp(-\lambda x) = 0.77$ of crashing at a track failure. When we change the damage intensity function from the rectangular to the triangular form, this change tends to focus track failures towards the center of the region of potential damage. (Recall we selected the this triangle to match the base and the area of this rectangle.) This concentration of risk towards the center of the interval reduces the probability that a moving train strikes a track failure while braking to 0.63. When we next include the effect of station stops, these safe pauses in stations for 5% of the operating day reduce the per-train p_{crash} risk by 5%, to 0.59. Upon including tunnels, train length, and regional train speeds, we arrive at the regional risk model, with p_{crash} values of $\{0.53, 0.39, 0.50\}$ for the Tokyo–Nasu, Nasu–Sendai, and Sendai–Morioka regions, respectively. Taken together, we may conclude that the various refinements of the regional crash risk model reduce p_{crash} from about 0.8 to about 0.5. We also conclude that p_{crash} varies only slightly across these three regions.

Expected number of crashing trains, by region

We have seen that the per-train crash risk is approximately constant across the three regions. However, the distribution of the number of at-risk trains varies significantly across the three regions (Figure 2-15). There are more trains running near Tokyo; on average, then, a quake near Tokyo will place more trains at risk (Figure 2-16). If a quake near Tokyo places more trains at-risk, on average more trains will crash following a quake near Tokyo (Figure 2-17).

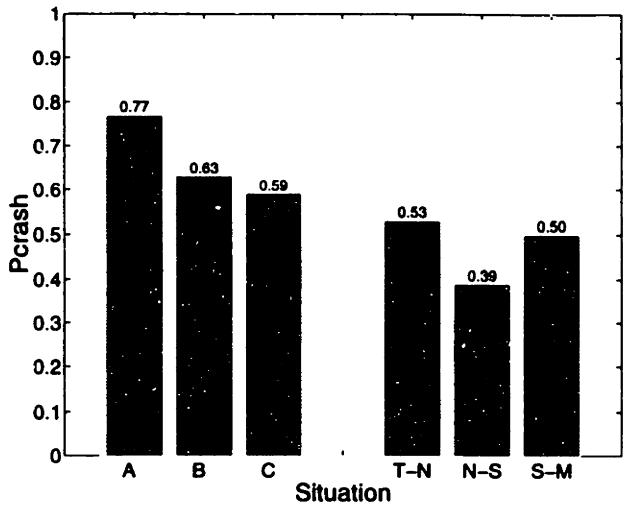

Figure 2-14: The contributions of various assumptions in the regional crash risk model. Each bar represents the crash risk faced a train inside the region of potential damage, by situation. Situation A corresponds to the simple crash risk model: no station stops, no tunnels, no train length, rectangular damage intensity function. Situation B modifies situation A by including the triangular damage intensity function. Situation C modifies Situation B by including station stops. Including tunnels, train length, and mean train speeds yields situations T-N, N-S, and S-M. These represent the results from the regional crash risk model for the Tokyo-Nasu, Nasu-Sendai, and Sendai-Morioka regions, respectively.

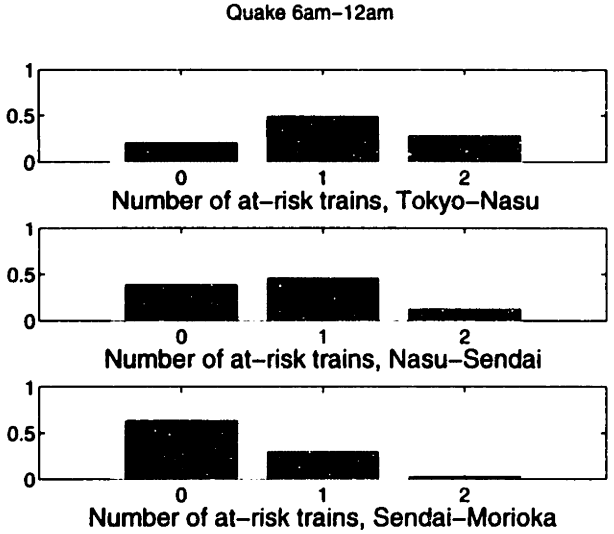

Figure 2-15: Probability mass functions for $(N_{tot}|A)$, the total number of at-risk trains by region for an earthquake that strikes between 6 am and 12 am, when the trains are in service. Base-case parameter values from Table 2.6.

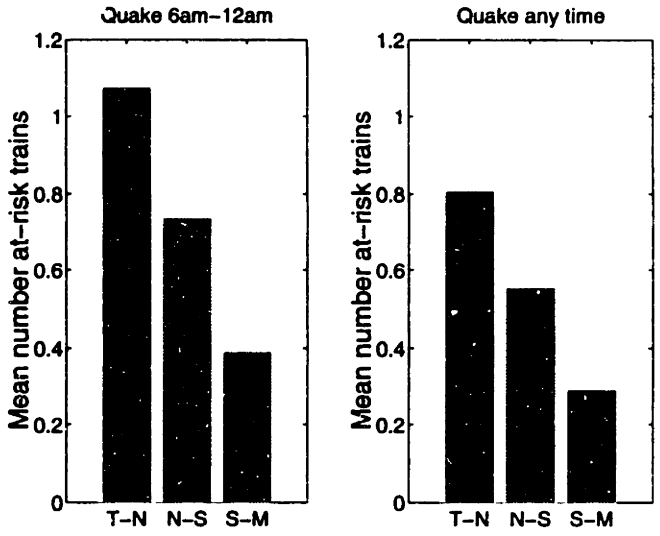

Figure 2-16: Mean number of at-risk trains, by region and time of day. T–N, N–S, and S–M denote the Tokyo–Nasu, Nasu–Sendai, and Sendai–Morioka regions, respectively. Base-case parameter values from Table 2.6.

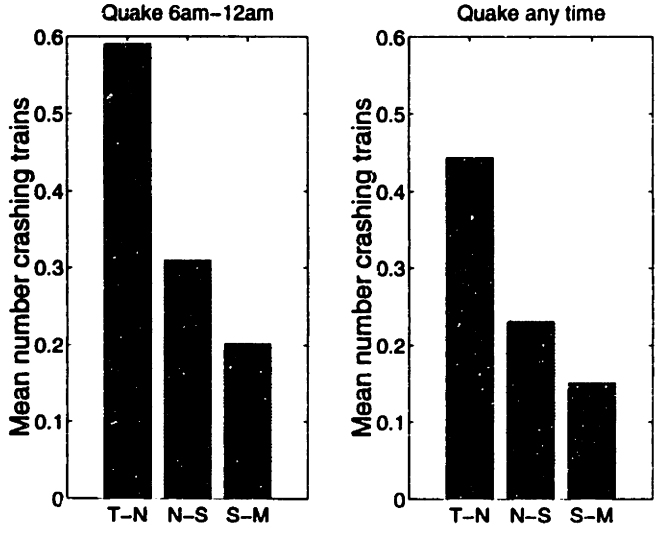

Figure 2-17: Mean number of crashing trains, by region and time of day, using the base-case parameter values from Table 2.6.

As per-train risk varies only slightly across the line, but the number of trains facing this risk varies greatly across the line, the railroad should focus their risk mitigation efforts in the areas with the greatest train traffic.

If we assume an earthquake is equally likely to strike anywhere along the Tohoku line, we may combine our regional crash estimates with a weighted average, weighting the expected number of crashing trains in each region by the regional fraction of total line length. With this assumption, our model suggests that if a major earthquake strikes somewhere along the Tohoku line when the trains are running, an expected 0.7 trains would crash at quake-induced track failures. If such a quake struck at a random time of day, an expected 0.5 trains would crash at quake-induced track failures.

2.3.3 Case study: the 1995 Kobe quake

Without warning, the magnitude 7.2 Kobe earthquake struck in the early morning of January 17, 1995 (Figure 2-18). The earthquake lasted less than twelve seconds and released more energy than 240 kilotons of TNT. [Ben95] The Kobe region was devastated: this disaster left 5,300 dead, 25,000 injured, and 50,000 buildings in ruins. The earthquake struck at 5:46 in the morning, fourteen minutes before the Sanyo Shinkansen was to have begun daily service. Thus, by a fortunate accident of timing, no in-service moving Shinkansen trains faced danger.

What would have happened to the Shinkansen system had this quake struck later in the morning, when the trains were running? We may use the framework of the regional crash risk model to answer this question. In this case, however, we may use the actual locations of failed track observed in the Kobe disaster. In effect, we will conduct a "thought experiment" in which we shift the Kobe disaster and its accompanying damage forward in time by several hours.

We assume the Sanyo Shinkansen trains cruise at $v_c = 240$ km/h and are $l_{train} = 250$ meters long. As this earthquake struck with no warning, we set $\tau = 0$, optimistically assuming the railroad would have reacted to the quake with no delay. We use Equation 2.1 to determine the potentially exposed braking distance for a Sanyo Shinkansen. The track failures observed in Kobe were more than $v_c^2/2a$ from any sta-

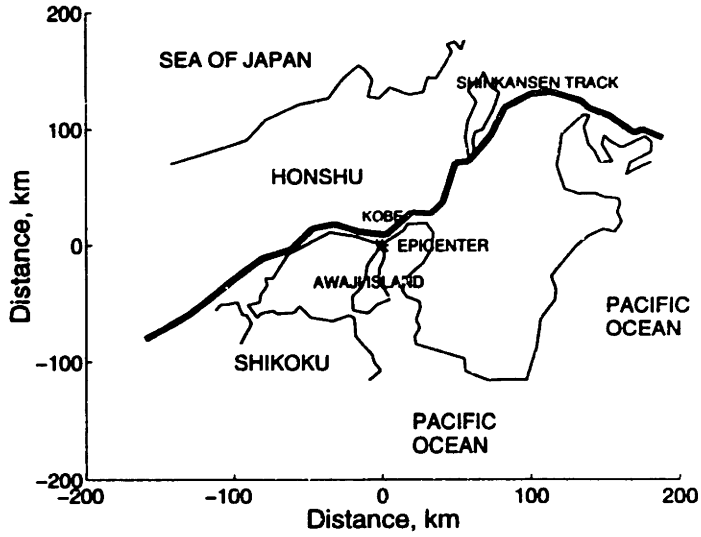

Figure 2-18: Location of 1995 Kobe earthquake.

tion; trains in the vicinity of these failures would have been traveling at full cruising speed.

$$x = l_{train} + \frac{v_c^2}{2a_e} = 0.25 + 2.96 = 3.2 \text{ km}$$

The Kobe quake shattered the Shinkansen viaduct in two areas, one at each end of the Kobe tunnel (Figure 2-19). We initially assume headways of $\Delta = 60$ km between successive trains in this region, and a $p_{station}$ value of 0.05. As the two failure areas are 40 km apart, at most one northbound and one southbound train could have crashed at these failures. Consider the closest northbound train approaching either one of these failure areas. If this train is within x kilometers before the southern track failure (modeled as a point), or within x kilometers before the northern cluster of track failures (modeled as a three kilometer interval), this train will be unable to stop before the damage and will crash. We assume the location of the train is independent of the time of the quake, and determine the probability the train is within x of either failure area by spatial considerations,

$$\begin{aligned} \Pr(\text{n'bound train crashes} \mid \text{n'bound train moving}) &= \frac{2x+3}{\Delta}, \quad \text{as} \quad \Delta > 2x+3 \\ &= (9.4/60) = 0.16 \end{aligned}$$

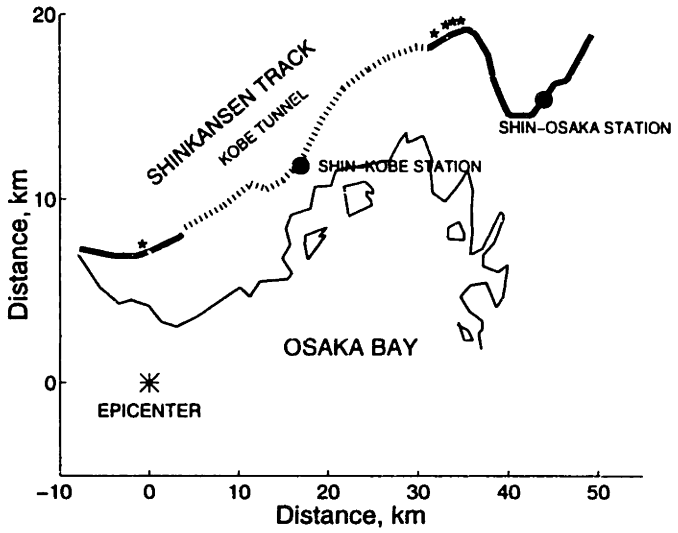

Figure 2-19: The damage to the Sanyo Shinkansen viaduct due to the Kobe quake. Small asterisks indicate the location of track failures, and the dashed line indicates the Kobe tunnel. Due to amplifications from local geological conditions, the point of peak ground acceleration lay 10 km northeast of the epicenter.

and account for station stops as before.

$$\begin{aligned} p_{kobe} = \Pr(\text{n'bound train crashes}) &= (1 - p_{station})\left(\frac{2x+3}{\Delta}\right) \\ &= (1 - 0.05)(0.16) = 0.15 \end{aligned}$$

We assume the positions of the northbound and southbound trains are independent. Thus, we may determine the distribution of the number of Sanyo Shinkansen trains that would have crashed in the Kobe disaster, had this earthquake struck several hours later in the day (Figure 2-20).

$$\Pr\left(k \text{ trains crash} \,\middle|\, \begin{matrix}\text{quake strikes when} \\ \text{trains in service}\end{matrix}\right) = \binom{2}{k}(p_{kobe})^k(1-p_{kobe})^{2-k}, \qquad k = 0, 1, 2.$$

This model suggests that had the Kobe quake struck later in the day, the most probable outcome would have been the loss of no trains (probability 0.71). The model estimates a 0.26 probability of one train crashing, and a 0.03 probability of two trains crashing. The expected number of crashing trains is 0.31 trains. We note the risk is

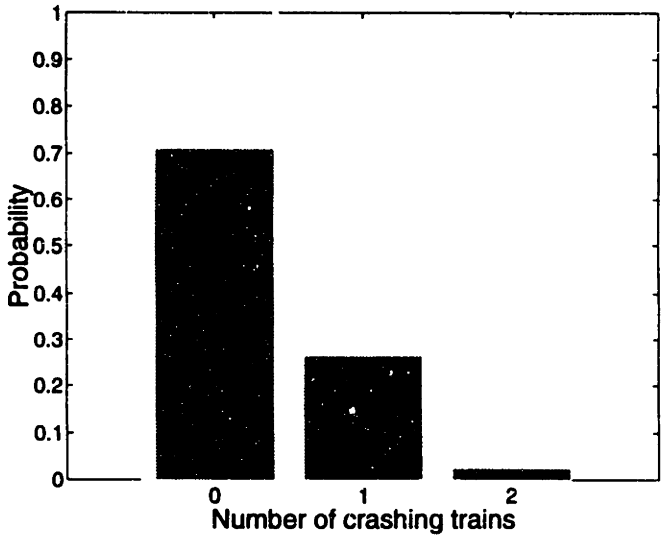

Figure 2-20: Distribution of the number of trains that would have crashed at the track failures caused by the Kobe quake, had this quake struck when the trains were running. This plot assumes trains cruise at $v_c = 240$ km/h; train headways of $\Delta = 60$ km; and neither early warning nor delay in responding to the quake ($\tau = 0$ seconds).

lower in this specific scenario than in the general regional crash risk model due to the presence of the Kobe tunnel. By fortunate luck, this tunnel shielded 40 kilometers of track from damage, cutting the average per-train risk by about one third.

Lacking good information on the actual headways on the Sanyo Shinkansen in this area, we use sensitivity analyses to determine how our results vary with Δ (Figure 2-21). We note that small uncertainties in Δ (say, between 50 and 70 km) lead to small shifts in the mean number of crashing trains (from 0.37 to 0.26, respectively).

2.4 Conventional Trains

We now shift our attention from Shinkansen trains to conventional trains. Compared to Shinkansen trains, conventional trains run more slowly, run closer together, run more frequently, and run shorter distances between stations. We modify the regional crash risk model to encompass these differences in Section 2.4.1. As a case-study, Section 2.4.2 considers the high-volume Yaminote line that circles downtown Tokyo.

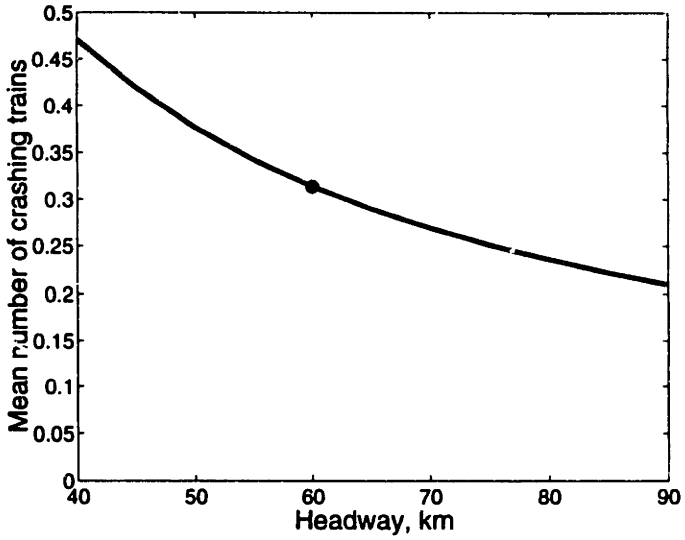

Figure 2-21: Sensitivity of the mean number of trains that would have crashed at the Kobe track failures, had this quake struck when the trains were running, as function of train headway, Δ. The black dot on the plot indicates the base-case parameter value.

2.4.1 Model formulation

Track damage model

As in our simple crash risk model, we employ a homogeneous spatial Poisson process to describe the locations of quake-induced track damage within the region of potential damage. As much of the conventional track rests upon earthen embankments, we believe that conventional track is somewhat less susceptible to outright failure than Shinkansen track, which rests on elevated viaduct. However, the urban environment presents additional hazards which could harm track in an earthquake. In urban areas, overpasses and bridges frequently cross above rail lines, and tall buildings often stand close to the tracks. A quake could cause such a bridge or building to collapse onto the track, destroying that area. These nearby structures serve to increase the rate of track failure. Our conventional model, then, uses the familiar Poisson damage model with constant rate λ over a region of length $2r$ from the simple crash risk model. The challenge is choosing sensible values for λ and r.

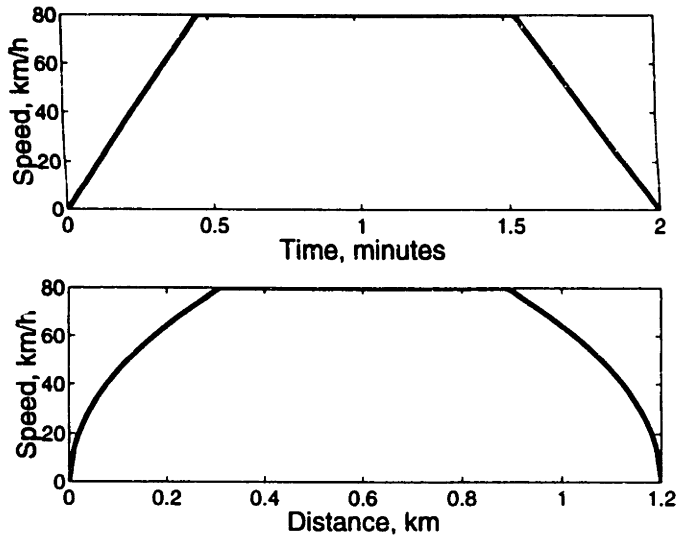

Figure 2-22: Speed versus time and distance for a typical inter-station run of a conventional train. This plot assumes constant linear acceleration and deceleration at rate $a_n = 0.8$ m/s/s, a cruising speed of $v_c = 80$ km/h, and an inter-station distance of $D = 1.2$ km.

Train speed, train length, tunnels, and station stops

Conventional trains spend almost half of their traveling distance between stations accelerating up to and decelerating down from their cruising speed (Figure 2-22). Following Equation 2.12, we let random variable V denote the speed of a moving train located at a random spot uniformly between two stations (Figure 2-23).

We treat the length of each train as a constant l_{train} meters. We assume the urban environment is free of bedrock tunnels that could protect a conventional train, and we assume a conventional train spends a fraction $p_{station}$ of its operating day paused in stations.

The fate of an "at-risk" train

We now consider a conventional train moving between two stations inside the region of potential damage when a quake strikes. For our base-case, we assume a train has no warning of the quake: $\tau = 0$. We also assume that when the quake strikes, the train is randomly located between two stations, and is traveling with speed V, where

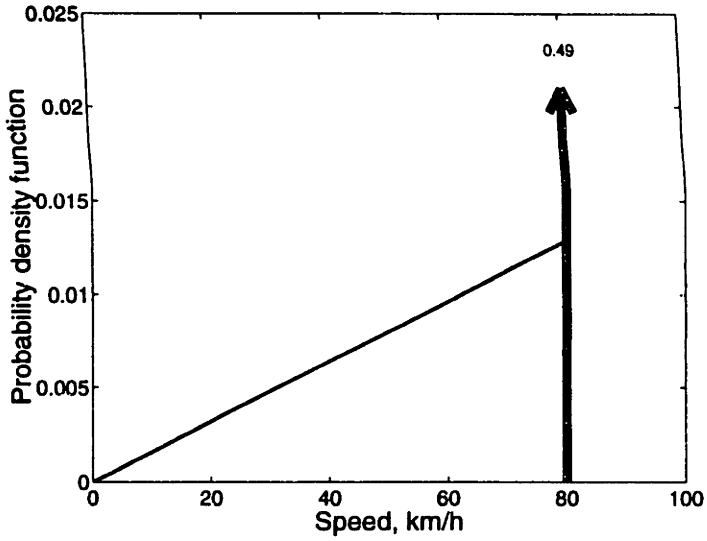

Figure 2-23: Mixed probability density function for the speed of a moving conventional train at a random location between two stations. The area represents a probability mass at cruising speed v_c. This plot assumes constant linear acceleration and deceleration at rate $a_n = 0.8$ m/s/s, a cruising speed of $v_c = 80$ km/h, and an inter-station distance of $D = 1.2$ km.

V is a random variable with density given in Equation 2.12. Assuming the train neither enters nor exits the region of potential damage while decelerating, we may determine p_{crash}, the probability the train encounters damaged track while braking.

$$\Pr\left(\begin{matrix}\text{train}\\ \text{crashes}\end{matrix}\,\middle|\,\begin{matrix}\text{train}\\ \text{at-risk}\end{matrix},\ \begin{matrix}\text{train}\\ \text{moving}\end{matrix}, V = v_0\right) = 1 - e^{-\lambda[v_0^2/2a_e + l_{train}]}$$

$$\begin{aligned} p_{crash} &= \int_{v=0}^{v_c} \Pr\left(\begin{matrix}\text{train}\\ \text{crashes}\end{matrix}\,\middle|\,\begin{matrix}\text{train}\\ \text{at-risk}\end{matrix},\ \begin{matrix}\text{train}\\ \text{moving}\end{matrix}, V = v\right) f_V(v)dv \qquad (2.26)\\ &= \left[\frac{2}{a_n D}\right] \int_{v=0}^{v_c} \left(1 - e^{-\lambda[v^2/2a_e + l_{train}]}\right) v dv \\ &+ \left[\frac{D - v_c^2/a_n}{D}\right] \left(1 - e^{-\lambda[v_c^2/2a_e + l_{train}]}\right) \qquad (2.27) \end{aligned}$$

At-risk trains, time of day, and mean crashes

Following the model and assumptions developed for the regional crash risk model, we use Equations 2.22 and 2.23 to describe the distribution of the number of at-risk

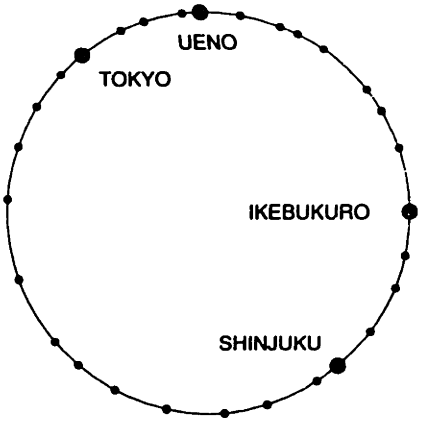

Figure 2-24: Schematic map of the Yaminote line. This 35 km line is not a perfect circle, but this schematic accurately presents the relative distances between the 29 stations, here indicated with black dots. Larger dots mark the largest stations.

trains in the region of potential damage; Equation 2.24 to adjust for the early-morning out-of-service interval; and Equation 2.25 to describe the mean number of crashing trains.

The regional Shinkansen crash risk model thus provides the backbone of our model for conventional trains. The conventional train model differs from the regional model in that it does not consider tunnels and that it does consider the acceleration-deceleration period.

2.4.2 Case study: the Yaminote line

We apply our crash risk model for conventional trains to the Yaminote line, the busiest JR line in metropolitan Tokyo. This rail line has 29 stations distributed along an irregular 35 km loop of track (Figure 2-24). For almost all of the operating day, "clockwise" and "counterclockwise" trains arrive at a station every two minutes. It takes a train a little over one hour to circumnavigate the entire loop, for a mean inter-station travel time of about two minutes. Thus, the spacing between adjacent trains is about the spacing between stations, and there are roughly 29 clockwise and 29 counterclockwise trains operating on the line during most of the day.

To assist in choosing an appropriate size for the region of potential damage for a

Parameter	Symbol	Value
Train length	l_{train}	50 m
Accel/decel rate	a_n	0.8 m/s/s
Emergency braking rate	a_e	0.9 m/s/s
% time in stations	$p_{station}$	0.15
Train cruising speed	v_c	80 km/h
Inter-station distance	D	1.2 km
Failure intensity	λ	0.25 failure/km
(Radius damage region)	(r)	(entire line at risk)

Table 2.7: Parameters for the crash risk model for conventional trains. These values are intended to model the effects on the Yaminote line of a major earthquake striking in metropolitan Tokyo.

hypothetical Tokyo earthquake, we consider the damage observed in the Kobe and Northridge quakes. In Kobe, most of the damage occurred in a 60 km by 20 km corridor, with the heaviest damage in a 40 km by 10 km corridor. In Northridge, most of the damage was concentrated within 10 to 20 km of the epicenter. [Ben95] The "radius" of the Yaminote irregular loop is between six and eight kilometers long. Thus, if a major earthquake struck in metropolitan Tokyo, a scenario similar to the Great Kanto Quake of 1923, the entire Yaminote line could fall within the region of potential damage.

We again turn to the Kobe quake to assist in choosing an appropriate rate of track failure. In the Kobe quake, conventional rail lines failed at six or seven locations between stations, and another five or six locations in stations. [IRJ95] (Recall we do not consider damage to trains in stations, as such trains are not moving, and passengers upon them incur little additional risk by virtue of being on a train.) With much uncertainty, we will take $\lambda = 0.25$ failures/km as our base case, and explore $\lambda \in [0.05, 2]$ failures/km.

Selecting values for the remaining parameters for the Yaminote line is simpler. We assume the emergency braking rate, a_e is 0.9 m/s/s; the length of a train is $l_{train} = 50$ m; and the standard acceleration and deceleration rates are $a_n = 0.8$ m/s/s. From the Yaminote time-table, we assume speed of $v_c = 80$ km/h. Assuming the average station stop is 20 seconds, and the average inter-station run lasts 2 minutes, we set $p_{station} = 0.15$ (Table 2.7).

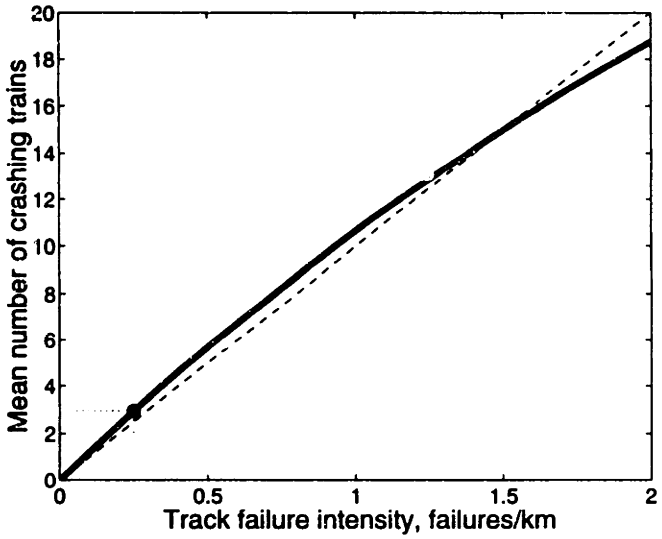

Figure 2-25: Mean number of crashing trains on the Yaminote line, given a major earthquake striking metropolitan Tokyo while the trains are running, as a function of track damage intensity. Base-case parameters from Table 2.7; a black dot indicates the base-case λ. The dashed line plots the function $y = 10x$.

With these values, we expect an average of $(35)(0.25) = 8.75$ failures along the Yaminote line. Trains need roughly $v_c^2/2a_e + l_{train} = 0.22 + 0.05 = 0.27$ kilometers to stop. From Equation 2.27, the probability a moving conventional train on the Yaminote line crashes at damaged track is about 0.06. Given our assumptions that trains spend 15% of their day in station, and that there are 58 at-risk trains on the line, we calculate that the mean number of crashing trains is $(58)(0.85)(0.06) = 3.0$.

We note that this result is strongly driven by our choice of λ, the track damage intensity (Figure 2-25). The relationship is nearly linear, with $E(\text{crashing trains}) \approx 10\lambda$. This indicates that our estimates of damage to the conventional rail system caused by a major Tokyo earthquake are quite sensitive to our assumption of the rate of track damage.

2.5 Comparing seismic risk: Shinkansen vs. conventional trains

Sections 2.2.1 and 2.3.1 estimated the mean number of Shinkansen trains that would crash at quake-induced track failures, given a major earthquake striking close to the Shinkansen track. Section 2.4 estimated the mean number of conventional trains that would crash at quake-induced track failures on a conventional line, given a major earthquake striking close to the line. The results of these two analyses are not easily compared as they are conditional on different events. In this section, we offer a first-order comparison between the seismic risk faced by the Shinkansen system and the seismic risk faced by the conventional system. We estimate N_s, the mean number of JR East Shinkansen trains that would crash if a catastrophic earthquake struck the island of Honshu during operating hours, and N_c, the mean number of JR East conventional trains that would crash if a serious earthquake struck the island of Honshu during operating hours.

For this calculation, we make the following simplifying assumptions.

- A serious earthquake is equally likely to strike at any spot in Japan.
- The damage region of a serious earthquake is a 20 km circle. (This is consistent with the Kobe and Northridge experience.) [Ben95]
- We assume (from Section 2.2.1) that

  $$\Pr(\text{Shinkansen train crashes} \mid \text{train within damage region of quake}) = 0.75$$

- We assume (from Section 2.4) that

  $$\Pr(\text{conventional train crashes} \mid \text{train within damage region of quake}) = 0.06$$

We first estimate N_s, the mean number of Shinkansen trains that would crash if a catastrophic earthquake struck anywhere on the island of Honshu during operating

Figure 2-26: The JR East Shinkansen network.

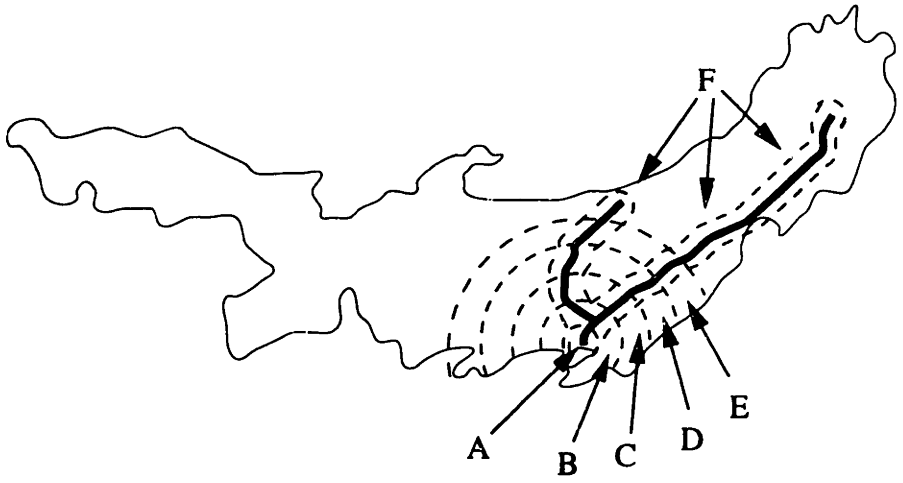

Figure 2-27: Partitioning the JR East network. Regions "A," "B," "C," "D," and "E" correspond to semicircular shells with radii 30, 50, 70, 100, and 130 km, centered on Tokyo Station, Tokyo; region "F" corresponds to a 40 km wide corridor around the two JR East Shinkansen lines.

hours. The JR East Shinkansen network consists of two lines: the 535 km Tohoku line, running from Tokyo to Morioka; and the 330 km Joetsu line, running from Tokyo to Niigata (Figure 2-26). As these lines run together for the first 50 kilometers from Tokyo, the total "unique" length of the JR East Shinkansen network is 815 km. If we assume the Shinkansen track is vulnerable to any serious earthquake which strikes within 20 km of the line, we estimate that the Shinkansen system faces risk from any quakes striking within a $(815)(20 + 20) = 32,600$ km^2 corridor containing the Shinkansen track (Figure 2-27). As the land area of the Honshu island of Japan is 2×10^5 km^2, we estimate that the probability an earthquake strikes within 20 km of the JR East Shinkansen network is roughly $32,600/(2 \times 10^5) = 0.17$.

JR East runs approximately 15 trains on these two lines at any given time during operating hours. [Eas95] Let us consider the 40 km wide corridor containing the Shinkansen track. The density of trains within this corridor is approximately $15/32,000 = 0.0005$ trains/km^2. With this density, a circle of radius 20 km contains an expected 0.57 trains. This calculation is valid only if the entire circle falls within the corridor; for these analyses, we adopt the simplifying assumption that this is always the case, ignoring "edge effects." (This leads to a slight overestimate in the mean number of trains in the danger region.) We may thus approximate the mean number of Shinkansen trains that would crash at quake-induced track failures, given a serious earthquake striking anywhere on the island of Honshu when the trains are running.

$$
\begin{aligned}
N_s &\equiv E(\text{number of crashing Shinkansen trains} \mid \text{Honshu quake, 6am-12am}) \\
&\approx \Pr(\text{train crashes} \mid \text{train in damage region}) \times \\
&\quad E(\text{trains in damage region} \mid \text{quake strikes in corridor}) \times \\
&\quad \Pr(\text{quake strikes in corridor}) \\
&= (0.75)(0.57)(0.17) = 0.07 \text{ trains}
\end{aligned}
$$

We next estimate N_c, the mean number of conventional trains that would crash if an earthquake struck anywhere in Honshu during operating hours. The geographic center of the JR East conventional rail system is Tokyo Station, in downtown Tokyo. [SRL96a] Most JR East train operations occur close to this center. We partition the JR East network into five regions, each a semi-circular shell centered on Tokyo station, with respective radii of 30, 50, 70, 100, and 130 km, respectively (Figure 2-28). We assume that these five regions collectively cover the entire JR East conventional rail system: this assumption is essentially correct. By dividing the number of conventional trains in each region by the area of each region, we estimate the average train density for each region (Figure 2-29). [SRL96a] Following the same procedure we used for the Shinkansen calculation, we determine the expected number of conventional trains that would crash in each region, given a serious earthquake striking the island of

Region	Area, km^2	Pr(strike)	Train density, trains/km^2	Mean trains, 20 km circle	Mean crash, if quake	% of total
A	1,400	0.01	0.156	196.6	0.074	57%
B	2,500	0.01	0.032	40.0	0.027	21%
C	3,750	0.02	0.011	13.3	0.013	10%
D	8,000	0.04	0.006	7.1	0.015	12 %
E	10,800	0.05	0.0001	0.1	0.0001	0%
Total	26,450	0.13	X	X	0.129	100%

Table 2.8: Seismic vulnerability calculation for the conventional network. The column labeled "Pr(strike)" indicates the probability that an earthquake strikes inside this region, given an earthquake strikes the island of Honshu. The column labeled "Mean trains, 20 km circle" indicates the expected number of trains inside a 20 km circle located entirely within the region. The column labeled "Mean crash, if quake" indicates the expected number of trains that would crash in this region if a serious earthquake struck the island of Honshu when the trains were operating.

Figure 2-28: Partitioning the JR East conventional network into five semi-circular shells. Figure 2-27 indicates the spatial location of these shells.

Honshu during operating hours, and sum these expectations to obtain $N_c = 0.129$ trains (Table 2.8).

From this analysis, we conclude that if one measures seismic risk via the expected number of crashing trains, the JR East conventional system faces roughly twice the seismic risk as the JR East Shinkansen system. We offer three additional observations.

- An earthquake is not equally likely to strike anywhere on the island of Honshu. The rate of earthquake activity varies across the island. Moreover, this rate is particularly large in the Tokyo region. [Uts74] Thus, it is likely that these calculations greatly underestimate the seismic risk faced by conventional trains.

- In terms of passenger death and injury, the consequences of each high-speed Shinkansen crash could be far more severe than the consequences of each conventional train crash. This analysis does not take this into account.

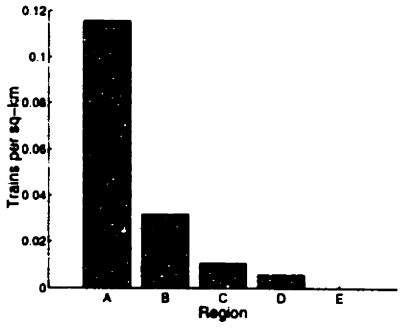

Figure 2-29: Density of conventional trains, by region.

- From the last column in Table 2.8, we observe that the large number of conventional trains operating in downtown Tokyo makes this region the primary contributor to total seismic risk faced by the conventional system. The heterogeneity of train density across the JR East conventional network increases the variance of the number of conventional trains that would crash in a serious earthquake. In contrast to the Shinkansen network, quake damage to the conventional network would more likely be "all-or-nothing," with severe damage if the quake strikes in metropolitan Tokyo, and little damage otherwise.

2.6 Crash speed model

We have analyzed the expected number of trains which would crash at track failures in the aftermath of an earthquake. We now address a different concern: given a train that crashes at a track failure, how fast would it be moving when it crashes?

Consider a train traveling at speed v_c when a quake strikes. Assume this train begins emergency braking, but crashes at a track failure before reaching a safe stop. Our intuition might suggest that the expected crash speed for such a train would be half the cruising speed, or $v_c/2$. Unfortunately, however, the actual expected crash speed is substantially higher. The shape of a train's braking curve drives this somewhat counterintuitive result (Figure 2-30). While braking with uniform deceleration, a train's speed decreases linearly with time, but with the square root of distance.

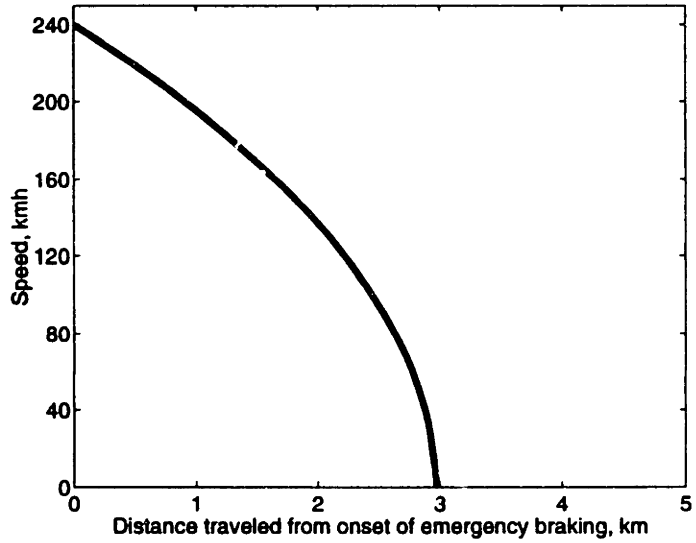

Figure 2-30: Sample emergency braking curve. This curve plots the speed of a braking Shinkansen against distance traveled while braking.

Thus, while braking, a train travels most of its braking distance at higher speeds.

2.6.1 Model formulation

We now determine the mean and median impact speed of crashing trains. We return to our earlier simplifying assumption that an earthquake and its accompanying damage take place in an instant, and we again denote this moment as time zero. For this model, we are only concerned with the distance traveled by the front of the train, and we recycle Equation 2.1 to describe the exposed stopping distance.

$$x = \begin{cases} v_c\tau + \frac{v_c^2}{2a} & \text{if } \tau \geq 0 \\ \frac{[\max\{0, v_c + a\tau\}]^2}{2a} & \text{if } \tau < 0 \end{cases}$$

Consider a train that will crash. At the moment just after the quake, there must be one (or more) track failures in the x km ahead of this train. Let random variable Y denote the distance from the train to the closest viaduct failure in front of it at time zero. For simplicity, we will use the constant track failure intensity function

(Figure 2-2). With this model, Y has a truncated exponential distribution.

$$f_Y(y) = \frac{\lambda e^{-\lambda y}}{1 - e^{-\lambda x}}, \qquad 0 \leq y \leq x \tag{2.28}$$

We next determine a function, $g(d)$, to describe a train's speed as a function of d, the distance traveled since the instant of the quake. From basic physics, we know that the velocity of a body experiencing constant deceleration drops off with the square root of distance traveled. [HR74, p. 34] For example, if the train begins emergency braking at the instant of the quake ($\tau = 0$), the function $g(d)$ takes a simple form.

$$g(d) = \sqrt{\max\{0, v_c^2 - 2ad\}} \tag{2.29}$$

Including early warnings and delayed braking is straightforward, but complicates the expression of $g(d)$ somewhat. Earlier warnings ($\tau < 0$) allow the train to begin braking before the earthquake arrives, shifting the braking curve to the left. Delayed braking ($\tau > 0$) means the train travels some additional distance at speed v_c, shifting the braking curve to the right. (Recall Figure 2-1). When $\tau > 0$, we must consider two cases: when the train is still at full speed ($d < v_c\tau$), and when the train has begun to slow ($d > v_c\tau$). As before, judicious use of maximum functions keeps all speeds nonnegative.

$$g(d) = \begin{cases} \sqrt{\max(0, (\max\{v_c + a\tau, 0\})^2 - 2ad)}, & \tau < 0 \\ v_c, & d \leq v_c\tau, \quad \tau \geq 0 \\ \sqrt{\max\{0, v_c^2 - 2a(d - v_c\tau)\}}, & d > v_c\tau, \quad \tau \geq 0 \end{cases} \tag{2.30}$$

Let random variable V denote the speed of the train when it strikes the closest viaduct failure in front of it. We use Equation 2.30 to express the relationship between random variables V and Y.

$$V = g(Y) \tag{2.31}$$

$$E(V) = \int_0^{v_c} g(y) f_Y(y) dy \tag{2.32}$$

$$\text{Med}(V) \;=\; w \text{ such that } \int_0^w f_Y(y)dy = 0.5 \tag{2.33}$$

Our crash speed model consists of Equations 2.28–2.33. Before providing numerical examples of these results, we pause to present interesting lower bounds for $E(V)$ and $\text{Med}(V)$. These bounds are obtained via a limiting argument on λ. As our model considers a train that will crash, there must be one (or more) track failures in the x kilometers ahead of the train. If λ is very small, there is most likely exactly one failure in $[0, x]$, and thus Y is essentially uniform over this interval.

$$\lim_{\lambda \to 0} f_Y(y) = \lim_{\lambda \to 0} \frac{\lambda e^{-\lambda x}}{1 - e^{-\lambda x}} = \frac{1}{x}, \qquad 0 \le y \le x$$

For simplicity, we consider the case when $\tau = 0$. The corresponding form for $g(d)$ from Equation 2.29 leads to a simple distribution for random variable V.

$$f_V(v) = \frac{v}{a} f_Y(\frac{v_c^2 - v^2}{2a}) = \frac{v}{a}\left(\frac{\lambda \exp\left[-\lambda \frac{v_c^2 - v^2}{2a}\right]}{1 - \exp[-\lambda x]}\right), \qquad 0 \le v \le v_c$$

As λ becomes very small, this distribution becomes essentially triangular on $[0, v_c]$.

$$\lim_{\lambda \to 0} f_V(v) = \left(\frac{v}{a}\right)\left(\frac{1}{x}\right) = \frac{2v}{v_c^2}, \qquad 0 \le v \le v_c$$

This limiting argument provides lower bounds for the mean and median crash speed.

$$E(V) \;\ge\; \lim_{\lambda \to 0} E(V) = \frac{2}{3} v_c \tag{2.34}$$

$$\text{Med}(V) \;\ge\; \lim_{\lambda \to 0} \text{Med}(V) = \frac{v_c}{\sqrt{2}} \tag{2.35}$$

2.6.2 Case-study: Shinkansen and conventional trains

Using our base-case parameters (Table 2.1), we calculate that the mean speed of a crashing Shinkansen is 180 km/h, and the median speed is 200 km/h. We can also use this model to investigate the sensitivity of these results to the values of a, v_c, λ, and τ (Figure 2-31). We observe that the mean crash speed varies almost linearly

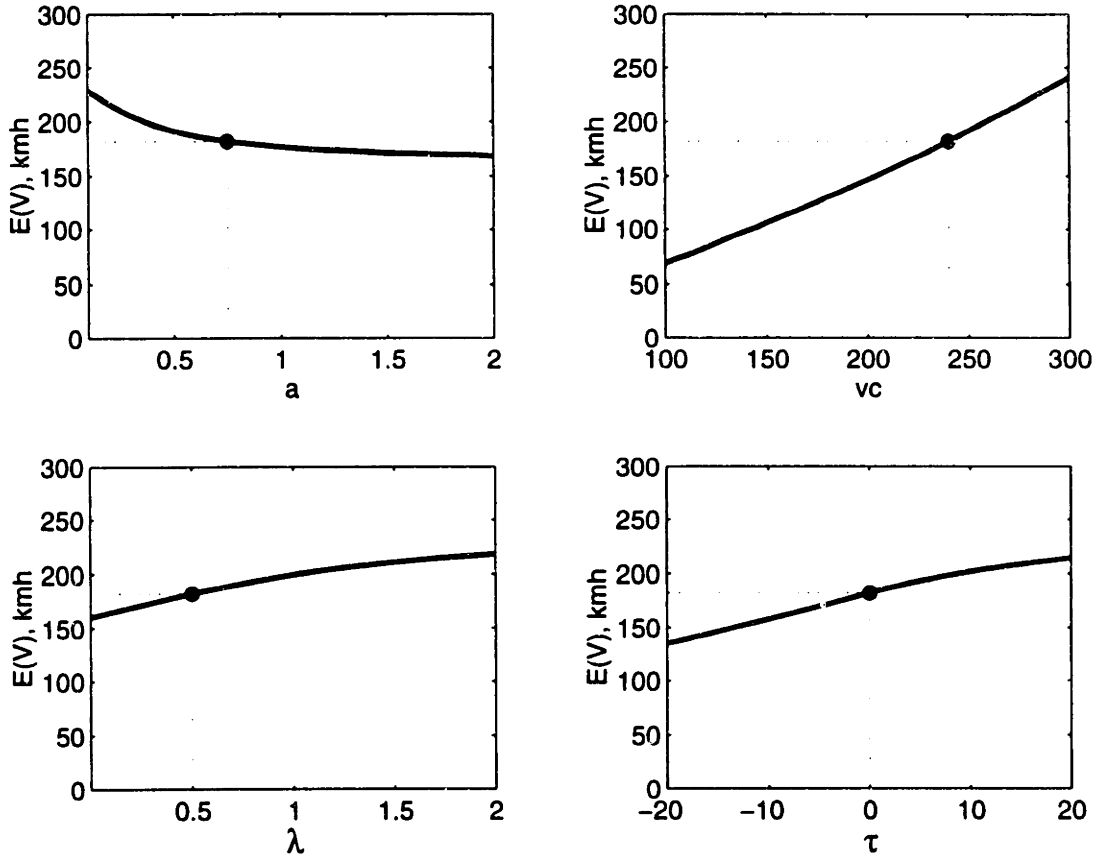

Figure 2-31: Sensitivity of $E(V)$, the mean impact speed of a crashing Shinkansen, to a, the emergency braking rate, in m/s/s; to v_c, the train cruising speed, in km/h; to λ, average rate of viaduct failure, in expected failures per kilometer; and to τ, the delay (or early warning) before braking, in seconds. The black dot on each plot indicates the base-case value.

with the train's cruising speed, v_c, with a 1 km/h change to v_c producing a 0.8 km/h to $E(V)$. The mean crash speed is less sensitive to a, the emergency braking rate; to λ, the mean viaduct failure rate; and to τ, the time braking begins relative to the instant of the quake. We note that each second of early warning reduces the mean crash speed by about 2 km/h.

Similarly, using our base-case parameters for conventional trains (Table 2.7), we can determine $E(V) = 55$ km/h and $\text{Med}(V) = 60$ km/h for trains which crash at track failures.

2.7 Conclusions

This chapter has presented a series of models that address the number and speed of conventional and high-speed Japanese trains that might crash at track failures in the aftermath of an earthquake. Our models are inherently simple; we do not intend these analyses to serve as detailed engineering studies (for they could not do so). Rather, we hope these calculations offer some valuable, though preliminary, insight into their subject.

We reiterate that these models do not include shaking risk. While we suspect this is a second-order phenomenon, it is not impossible that our models significantly underestimate the threat to trains.

Our most basic finding is that the earthquake risk to the Japanese rail system is very real, and numerically, somewhat large. That is, given a major earthquake, similar to the Kobe quake, striking when the trains are running, our models suggest a per-train crash risk in the range of 0.05 to 0.50, or $10^{-1.3}$ to $10^{-0.3}$. In no case is this probability 10^{-6} or 10^{-9}. The most basic conclusion of this chapter is also the most obvious: if an earthquake strikes near the trains when they are in service, there is a non-negligible probability that trains will crash.

In retrospect, we find the simple crash risk model to be surprisingly powerful. We find that most of the various refinements to this model (tunnels, station stops, acceleration and deceleration profiles, and so forth) change only the details, but not the substance, of our conclusions. The more complex models demonstrate that such refinements often yield only second-order corrections to the final result.

For Shinkansen trains on the Tohoku line in the region of potential damage, our models suggest that each moving trains faces a probability of crashing of about 0.5 if a major quake strikes close to the line. If a day-time quake strikes near Tokyo, where trains are more frequent, this implies that approximately one Shinkansen train is expected to crash due to track failures; if a day-time quake strikes in the northern region of the line, where trains are less frequent, this expectation drops to 0.4 trains.

In Kobe, the Sanyo Shinkansen was doubly fortunate. First, the quake struck

before the trains began service. Second, even if the quake had struck later in the day, the Kobe tunnel shielded 40 kilometers of track from damage; had the quake struck later, we estimate the probability of all trains stopping safely exceeds 0.7.

Our analysis of conventional trains suggest that they can brake to emergency stops in one tenth the distance required to stop Shinkansen trains. Thus, in contrast to a Shinkansen train, each conventional trains faces a much smaller chance of encountering a track failure; however, far more trains face this risk. Our models indicate that the expected number of conventional trains that would encounter track failures is highly sensitive to the intensity of track damage. On the Yaminote line, for example, this expectation could easily range from 3 to 20 crashing trains for plausible track failure intensities.

Our comparison of the conventional and high-speed networks suggests that the conventional system faces about twice the seismic risk of the Shinkansen system. The worst possible scenario—a quake striking downtown Tokyo during operating hours—accounts almost 60% of this risk. In contrast to the Shinkansen network, quake damage to the conventional network would more likely be "all-or-nothing," with severe damage if the quake strikes in downtown Tokyo, and little damage otherwise.

Our crash speed model suggests that if trains crash at track failures, they will tend to do so at higher speeds. This inescapable effect follows from physics of uniform deceleration: while braking, trains spend most of their braking distance traveling at high speed.

We offer a closing comment on these findings. It appears impossible for a railroad, and especially a high-speed railroad, to shield itself from earthquake risk altogether. If a large quake strikes close to the tracks when trains are running, it is very likely that trains will crash. That is not to say the railroad is powerless to reduce this risk. Indeed, we argue that stronger track, longer warnings before a quake, and faster brakes all serve to reduce this risk; hopefully, these models can assist a railroad determine the relative efficacy and cost-effectiveness of such improvements. Yet, we also feel that a railroad could benefit from explicitly acknowledging the inevitability of earthquake risk. This acknowledgment need not be advertised to the public. Internally,

however, this outlook could make safety planning more rational, perhaps leading to a greater focus on crash survivability of train cars, the railroad's emergency response procedures, and the issues of self-insurance.

Chapter 3

Slope failure risk to trains

3.1 Motivation and scope

Heavy rainfall can cause slopes to give way, can cause embankments to fail, and can cause rivers to flood over bridges. Near rail lines, such events can damage train tracks. The ensuing track damage, if not discovered in time, can lead to train derailments, and derailments may damage rolling stock and kill or injure passengers. To reduce the likelihood of such rain-induced derailment accidents, JR East monitors rainfall across its network using arrays of precipitation gauges. Using real-time precipitation data from these gauges, JR East slows and suspends train operations preventively during heavy rains in regions prone to rainfall-induced track disruption. This rainfall alert system reduces the probability that a train derails at track damage. To achieve this reduction in derailment risk, JR East pays the direct costs of maintaining and operating the alert system. The railroad also incurs the indirect costs of the disruption these train stoppages cause to regularly scheduled service.

Of the many ways heavy rainfall can lead to track damage, two are most common in the JR East system: slope failures, and embankment failures. A slope failure occurs when a hillside or cut becomes saturated with water and gives way. An embankment failure occurs when the earth-works supporting the track become saturated with water and collapse. We consider only these two scenarios, and uses the generic term "slope failure" to refer to both the failure of slopes and the failure of embankments.

This chapter investigates the performance of the current JR East rainfall alert rule, proposes and evaluates alternative alert rules, and quantifies the costs and benefits of such of systems.

3.2 The current rainfall alert system

This section reviews the current JR East rainfall alert system. We determine and discuss its performance; calculate the uncertainties associated with these performance estimates; introduce a model of driver awareness that permits a comparison of the different levels of JR East rainfall alerts; and illustrate the sensitivity of this alert system to slight modifications.

3.2.1 Description

We describe the current rainfall alert system using the Suigun line as a case-study. The Suigun line runs 100 km north of Tokyo. It is a single-track line, 150 km in length. The Suigun line runs through mountainous terrain on ledges cut into the hillside, and it travels through rice paddies and fields on raised earthen embankments. The Suigun line is a typical rural line, with infrequent train service, and several areas where the track bed is vulnerable to rain-induced failure.

To provide data for the rainfall alert system, JR East maintains an array of rain gauges along its network. On the Suigun line, gauges are spaced about every ten kilometers. Each gauge contains a tiny cup that holds one millimeter of rain. Upon accumulating one millimeter of rain, the cup tips over, empties itself, and sends a signal to the control center for conventional trains in Tokyo. At the control center in Tokyo, the natural hazards computer system receives these signals. These data streams provide the computer system with real-time precipitation information for the entire conventional network. The natural hazards computer system uses these rainfall data to decide when to issue operating regulations to trains. (This same computer system also monitors wind intensity, river height under bridges, and seismic acceleration, and issues alerts for these natural hazards as well.) Though it is the computer

system in Tokyo that actually issues these operating regulations, it is convenient to speak of each gauge as if it had local autonomy.

Using this convention, we say each gauge "controls" a fixed region of track. On the Suigun line, these gauge control regions overlap (Figure 3-1). At every instant, each gauge is in one of four states:

NO ALERT Trains operate normally, at full speed. For the Suigun line, this speed is 55 km/h.

WATCH Trains operate at full speed. Drivers are instructed to be extra-vigilant.

SLOW Trains operate at reduced speed. For the Suigun line, this speed is 35 km/h. Drivers are instructed to be extra-vigilant.

STOP Trains continue at reduced speed to the next station, then halt operation. Drivers are instructed to be extra-vigilant.

Intervals of track controlled by multiple gauges follow the most stringent alert level issued by those gauges. For example, if an interval of track is controlled by two gauges, and the first gauge issues a WATCH alert and the second issues a STOP alert, trains in that interval of track obey the STOP alert.

Each gauge uses two rainfall indices, CRAIN and HRAIN, to determine which level of alert to issue. HRAIN is defined as the total rainfall over the last hour. CRAIN is defined as the cumulative rainfall since the last period of twelve or more rain-free hours. JR East divides the positive HRAIN vs. CRAIN quadrant into four irregular regions, corresponding to the four alert levels (Figure 3-2). Each gauge issues the appropriate alert level based on the location of the current (CRAIN, HRAIN) point.

Let us demonstrate the operation of the current rainfall rule by considering the effect of a hypothetical storm on one gauge. Our example begins with an extended clear period, when CRAIN and HRAIN are both zero. A steady rain begins. For the first hour, CRAIN equals HRAIN, and both indices increase with additional rainfall, and the (CRAIN, HRAIN) point moves outward from the origin along the line $y = x$. After the first hour of steady rain, CRAIN continues to increase. As

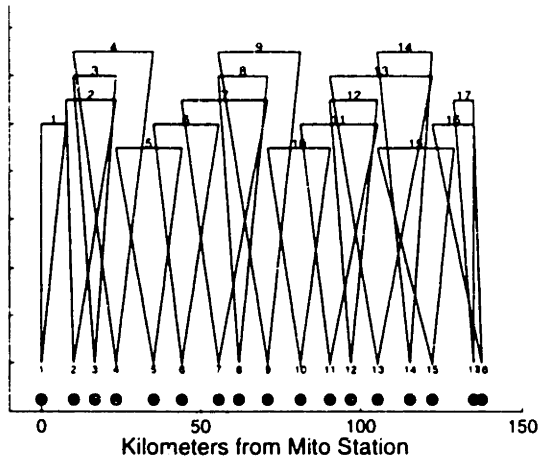

Figure 3-1: Rainfall gauge control regions on the Suigun line. The 17 dots along the x axis indicate the location of the 17 gauges along the line. The top of each triangle indicates the control region of each gauge. Triangles heights staggered for presentation only.

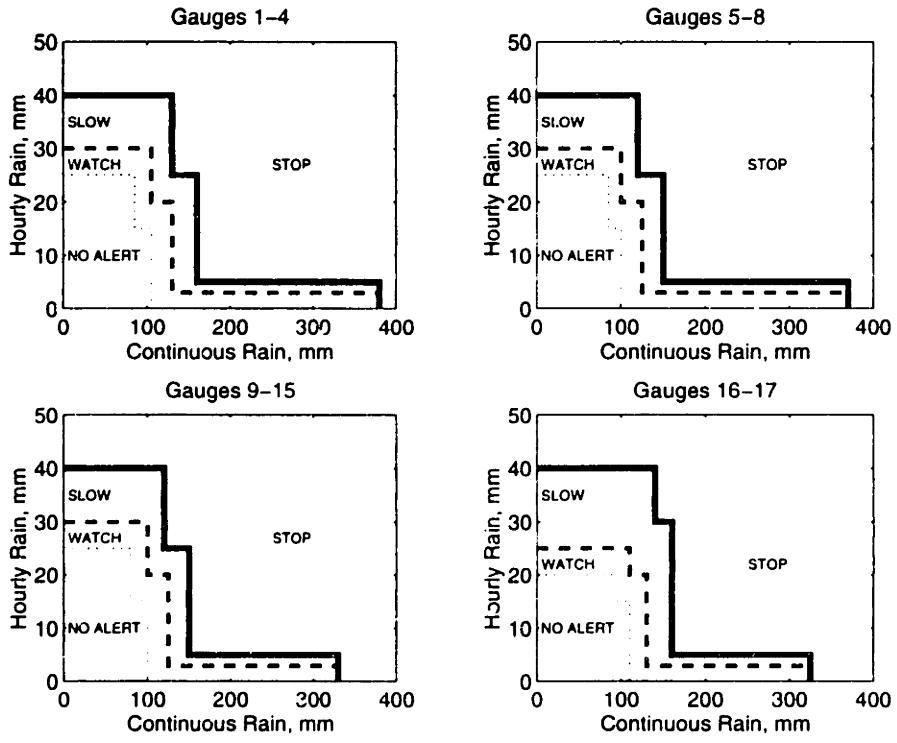

Figure 3-2: Rainfall gauge thresholds for the Suigun line, as functions of continuous rainfall (CRAIN) and hourly rainfall (HRAIN).

HRAIN accumulates rain only over the last hour, CRAIN now exceeds HRAIN. The rain continues. With additional rain, the (CRAIN, HRAIN) point moves to the left, further from the origin. When this point enters the WATCH region, the gauge issues a WATCH alert. The rain persists. When the point enters the SLOW region, the gauge issues a SLOW alert.

Suppose the rain stops abruptly, and the sun shines for an hour. Over this hour, HRAIN decreases to zero, while CRAIN maintains its value. Over this hour, as the (CRAIN, HRAIN) point moves vertically downward, the alert level drops from SLOW to WATCH to NO ALERT as the point traverses these regions. Let us suppose this clear spell lasts for the next eleven hours. Over this period, HRAIN equals zero, but CRAIN maintains its value. After a total of twelve rain-free hours, CRAIN also resets to zero.

Instead, let us suppose this clear spell lasts only two hours, and then the rain resumes. In this latter case, both HRAIN and CRAIN increase again with the additional rainfall. HRAIN starts from zero, while CRAIN starts from the accumulated rain from the previous storm. This aspect of the rule captures the effect of soil saturation: as CRAIN starts the second storm with some previous accumulated rainfall, it takes less additional rain to issue alerts than when CRAIN starts at zero.

3.2.2 Performance of the current alert system

We measure the raw performance of the current rainfall alert system using six quantities: four for safety, and two for delay. The four safety measures are:

1. the number of slope failures that occur during STOP alert periods,
2. the number of slope failures that occur during SLOW alert periods,
3. the number of slope failures that occur during WATCH alert periods, and
4. the number of slope failures that occur during NO ALERT periods.

Recall we use the generic term "slope failure" to refer to both collapses of the slope and collapses of the embankment. The two delay measures are:

1. the average annual hours each gauge is in the SLOW state, and

2. the average annual hours each gauge is in the STOP state.

In Section 3.2.4, we aggregate these six raw measures into a single measure of safety and a single measure of delay.

These six quantities may not be the natural performance measures of interest to the railroad. JR East is probably more concerned with hours of train disruption and service delay than with hours of gauge alerts. Likewise, the railroad is probably more concerned with the number of derailments or injuries than with the number of slope failures during STOP, SLOW, WATCH, and NO ALERT periods. However, we may approximate such measures from these six proxies. For example, we may estimate hours of train delay from gauge alert hours by multiplying by the average density of trains within a gauge's control region. Further, Section 3.2.4 offers a model-based procedure to aggregate these four safety measures into one, and Section 3.7 considers the consequences of derailment accidents.

To evaluate the performance of the current alert rule, we utilize three JR East datasets provided by the Safety Research Laboratory: a rainfall database, a rainfall alert database, and a slope failure database. [SRL95d] [SRL95f] [SRL95g]

Precipitation dataset

The JR East rainfall database provides precipitation data for each of the 17 gauges along the main Suigun line. These data cover the period from April 1, 1991 to December 31, 1994. The dataset provides precipitation data for each five minute period, rounded to the nearest millimeter. With seventeen gauges, twelve observations per gauge per hour, and 3.75 years of data, this dataset includes 6.7 million rainfall observations. As clear weather is far more common in this region than rainy weather, 98.9% of these observations are zero (Table 3.1). This does not imply the Suigun line enjoys sunny weather for 361 days (98.9%) each year; rather, it indicates that most rain storms are peppered with many five minute intervals with virtually no rain.

mm	0	1	2	3	4	5	6	7	8	9	10+
#	6,636,554	68,519	5,440	1,153	340	170	111	34	11	13	7
%	98.9	1.0	0.08	0.01	0	0	0	0	0	0	0

Table 3.1: Histogram of five-minute JR East precipitation data. The first row indicates millimeters of rain in a five minute period, the second indicates the number of such periods in the JR East dataset, and the third row provides the fraction of such periods.

Rainfall alert dataset

The rainfall alert database provides information on the schedule disruptions caused by rainfall alerts. For each alert, the database indicates the alert's date, starting time, and the track region affected. These data were collected to quantify the delay costs of rainfall alerts. They were not collected with the direct intention of evaluating the rainfall alert system. For this reason, the dataset lacks information on the duration and the level of each alert. Further, the dataset only includes those alerts which cancelled one or more trains or delayed a train for at least thirty minutes. As such, all WATCH alerts and some SLOW alerts are not included. Despite these limitations, the rainfall alert dataset is useful, for it provides data on the occurrence of rainfall alerts along the Suigun line between 1991 and 1994. The rainfall alert database contains 41 such alerts within this region and time period.

We apply the current rainfall alert rule to the JR East rainfall precipitation data to generate our own set of "inferred" alerts, so named as we infer them from the rain data using the alert rule. We thus obtain 53 inferred alerts. Thirty-six of these (68%) match actual alerts in the alert database. Seventeen of the inferred alerts do not have analogues in the actual alert data. As these alerts are typically WATCH alerts, we assume they are absent from the alert database as they did not cause train delay or cancelation. Five actual alerts do not have analogues in our set of inferred alerts. These discrepancies may well arise from small patches of missing data in the precipitation database. While no slope failures occurred during these patches of missing data, the fact that $5/58 = 9\%$ alerts are obscured by missing data call into question the reliability of the precipitation dataset for periods of heavy rain—those periods where accurate data are most critical.

	Inferred alerts	Non-inferred alerts
Actual alerts from database	36	5
Non-actual alerts	17	X

Table 3.2: Matching actual and inferred JR East alerts.

Matching actual and inferred alerts in this way, we can determine the total number of alerts issued, their spatial extent, their level, and their duration (Table 3.2). As WATCH alerts cause neither train delay nor cancelation, we focus our attention on SLOW and STOP alerts. We conclude that between 1991 and 1994 on the Suigun line, the JR East rainfall alert system issued a total of 58 WATCH and SLOW alerts. These involved 532 gauge hours of STOP alerts and 303 gauge hours of SLOW alerts. With seventeen gauges and 3.75 years of data, this implies that each rainfall gauge was in the SLOW state for an average of 4.8 hours per year, and each gauge was in the STOP state for an average of 8.2 hours per year.

Slope failure database

The slope failure database provides information on all rain-induced track failures near the Suigun line between 1938 and 1994. Twenty of these failures occured between 1991 and 1994, the period for which we have precipitation data.

The Safety Research Laboratory classifies slope failures as serious (in its terms, "affecting train service") or minor ("not affecting train service"). A serious slope failure is a failure that damages the track sufficiently to endanger passing trains. A minor slope failure is a failure that occurs on JR East property or right-of-way but does not harm the track. The Safety Research Laboratory classifies eight of the twenty slope failures between 1991 and 1994 as serious. Of these eight, five were collapses of the track embankment, two were collapses of the cut slope alongside the track, and one was a collapse of a culvert under the roadbed. This last failure was caused by gradual deterioration of the culvert, not by excessive precipitation. Following the recommendation of the Safety Research Laboratory, we omit this culvert collapse from our analyses. [SRL96a] We are left with seven serious slope failures along the Suigun main line between 1991 and 1994; this small sample constitutes our study

SAFETY
7 slope failures;
5 during STOP, 1 during SLOW, 1 during WATCH
DELAY
4.8 SLOW hours/gauge/year, 9.2 STOP hours/gauge/year

Table 3.3: Performance of JR East rainfall alert rule on the Suigun line, 1991–1994.

population of actual slope failures.

For each failure, the slope failure database provides the date and time at which the failure was discovered. We are less interested in the time the failure was discovered than we are in the time the failure actually occurred; unfortunately, the latter is unknown. As the railroad conducts on-foot inspections for slope failures every four hours during rain storms, each failure time must be within the four hours prior to its discovery. To resolve this ambiguity, we consider the hourly rainfall over the four hours prior to discovery (as measured at the gauge closest to the site of the failure), and we assign the hour with the highest rainfall as the time of failure.

Using the precipitation database, the rainfall alert database, the slope failure database, and the rainfall alert rule, we may determine how the current rainfall alert system performed in providing warnings for these seven slope failures. We find that the alert system issued STOP alerts for five of the seven, a SLOW alert for one of the seven, and a WATCH alert for one of the seven. Table 3.3 summarizes the performance of the JR East rainfall alert rule on the Suigun line between 1991 and 1994.

We note there were no rain-induced derailments on the Suigun line over this period. The slope failure that occurred during the WATCH period was discovered by an on-foot inspection team before it could threaten any trains. The slope failure that occurred during the SLOW period was discovered by an approaching train: a vigilant driver noticed the failure in time and brought his train to a safe stop.

3.2.3 Uncertainty in safety performance measures

We have estimated six performance measures for the JR East rainfall alert system. Two of these estimates, the average annual SLOW hours per gauge and the average annual STOP hours per gauge, are based on over 500,000 hours of precipitation data. As the sampling error associated with such averages drops off with the square root of the sample size, we have reasonable confidence that these point estimates capture the approximate magnitude of these delays.

The four safety performance measures are more uncertain. Faced with seven slope failure events, the current alert system provided STOP alerts for five, a SLOW alert for one, and a WATCH alert for one. Given an hour during which a slope failure occurs, let p_i denote the conditional probability that the current rainfall alert rule provides alert level i, where $i = 0$ for NO ALERT, $i = 1$ for WATCH, $i = 2$ for SLOW, and $i = 3$ for STOP. We may calculate MLE point estimates of the p_i directly,

$$\begin{aligned} \hat{p}_0 &= 0/7 = 0.00 \\ \hat{p}_1 &= 1/7 = 0.14 \\ \hat{p}_2 &= 1/7 = 0.14 \\ \hat{p}_3 &= 5/7 = 0.72, \end{aligned}$$

but as these estimates are based on only seven slope failures, the uncertainty of such estimates is of concern.

To assess this uncertainty, we take a Bayesian approach. [Box92] Let us first define a polytope $\mathcal{S}$ for convenience.

$$\mathcal{S} = \{p_1, p_2, p_3 \mid 0 \leq p_1 + p_2 + p_3 \leq 1, 0 \leq p_1 \leq 1, 0 \leq p_2 \leq 1, 0 \leq p_3 \leq 1\}$$

We begin with an uninformative prior on the vector $\vec{p} = [\ p_1 \quad p_2 \quad p_3\ ] \in \mathcal{S}$: that is, before considering the data, we favor all $\vec{p}$ values in $\mathcal{S}$ equally. Because the volume of polytope $\mathcal{S}$ is $1/6$, the probability density over this region is 6.

$$f(\vec{p}; \text{prior}) = 6, \quad \vec{p} \in \mathcal{S}. \tag{3.1}$$

Note this is a pessimistic assumption, for it implies we have little prior confidence in the efficacy of the alert system. Also note that we need not specify p_0, as $p_0 = 1 - p_1 - p_2 - p_3$. We write the likelihood function for the observed distribution of the seven failures across the four alert levels via a multinomial distribution.

$$\begin{aligned} L(\vec{p}) &= \frac{7!}{(0!)(1!)(1!)(5!)}(1 - p_1 - p_2 - p_3)^0(p_1^1)(p_2^1)(p_3^5) \\ &= 42(p_1^1)(p_2^1)(p_3^5), \quad \vec{p} \in \mathcal{S} \end{aligned}$$

We then use Bayes Law to determine the posterior distribution of $\vec{p}$.

$$\begin{aligned} f(\vec{p}; \text{posterior}) &= \frac{L(\vec{p})f(\vec{p}; \text{prior})}{\int_{\vec{r}\in\mathcal{S}} L(\vec{r})f(\vec{r}; \text{prior})d\vec{r}} \\ &= \frac{(p_1)(p_2)(p_3^5)}{\int_{\vec{r}\in\mathcal{S}}(r_1)(r_2)(r_3^5)d\vec{r}} \\ &= (30240)(p_1)(p_2)(p_3^5), \quad \vec{p} \in \mathcal{S} \end{aligned} \tag{3.2}$$

We next calculate the marginals of this joint posterior distribution.

$$\begin{aligned} f_{p_1}(p_1; \text{posterior}) &= \int_{p_2=0}^{1-p_1}\int_{p_3=0}^{1-p_1-p_2} f(\vec{p}; \text{posterior})dp_2dp_3 \\ &= 90p_1^8(p_1 - 1), \qquad p_1 \in [0,1] \\ f_{p_2}(p_2; \text{posterior}) &= \int_{p_1=0}^{1-p_2}\int_{p_3=0}^{1-p_1-p_2} f(\vec{p}; \text{posterior})dp_1dp_3 \\ &= 90p_2^8(p_2 - 1), \qquad p_2 \in [0,1] \\ f_{p_3}(p_3; \text{posterior}) &= \int_{p_2=0}^{1-p_3}\int_{p_1=0}^{1-p_1-p_3} f(\vec{p}; \text{posterior})dp_1dp_2 \\ &= 1260p_3^5(1 - p_3)^4, \qquad p_3 \in [0,1] \end{aligned}$$

We also recover the marginal for p_0 as well.

$$\begin{aligned} f_{p_0}(p_0; \text{posterior}) &= \frac{d}{dp_0}\Pr(P_0 \le p_0) \\ &= \frac{d}{dp_0}[1 - \Pr(P_1 + P_2 + P_3 \le 1 - p_0)] \end{aligned}$$

$$= \frac{d}{dp_0}\left[1 - \int_{p_1=0}^{1-p_0}\int_{p_2=0}^{1-p_0-p_1}\int_{p_3=0}^{1-p_0-p_1-p_2} f(\vec{p};\text{posterior})dp_1 dp_2 dp_3\right]$$

$$= 10(1-p_0)^9, \qquad p_0 \in [0,1]$$

We note that because our prior is a Dirichlet distribution, all of these posterior distributions are also Dirichlet. [Arn90, p. 538]

We use these posterior distributions to compute Bayesian 95% confidence intervals for p_0, p_1, p_2, and p_3. We determine the endpoints of these intervals by calculating the 2.5% and 97.5% quantiles of these posterior distributions (Figure 3-3). We thus obtain confidence intervals of $[0.0025, 0.31]$ for p_0, $[0.25, 0.445]$ for p_1 and p_2, and $[0.263, 0.812]$ for p_3. These wide intervals reflect the large uncertainties in the p_i arising from this small sample of seven slope failures. That is, while it is true that none of these seven failures occurred during the NO ALERT state, and $\hat{p}_0 = 0/7$, it is also true that p_0 could be significantly larger than zero. Indeed, with our assumption of a uniform prior, it is even plausible that p_0 is as large as 0.3, and this perfect record in seven trials is only a matter of good luck.

3.2.4 Derived performance measures

The six measures in Table 3.3 summarize the performance of the current rainfall alert rule on the Suigun line between 1991 and 1994, but such multivariate performance measures are unwieldy. To permit simple safety vs. delay tradeoff analysis, we shall aggregate the four measures of safety into a single univariate safety measure, and aggregate the two measures of delay into a single univariate delay measure.

Such aggregation, however, requires subjective judgment. For example, we might believe that trains near slope failures are safest when stopped, and if this is not possible, are safer when moving more slowly. In terms of reducing the probability of derailing at a slope failure, we might believe that STOP alerts are most effective, SLOW alerts are less effective, and WATCH alerts are no more effective than the NO ALERT state. Evaluating the performance of the current rainfall alert system from this viewpoint, we might say the alert system "caught" five of the seven failures (the

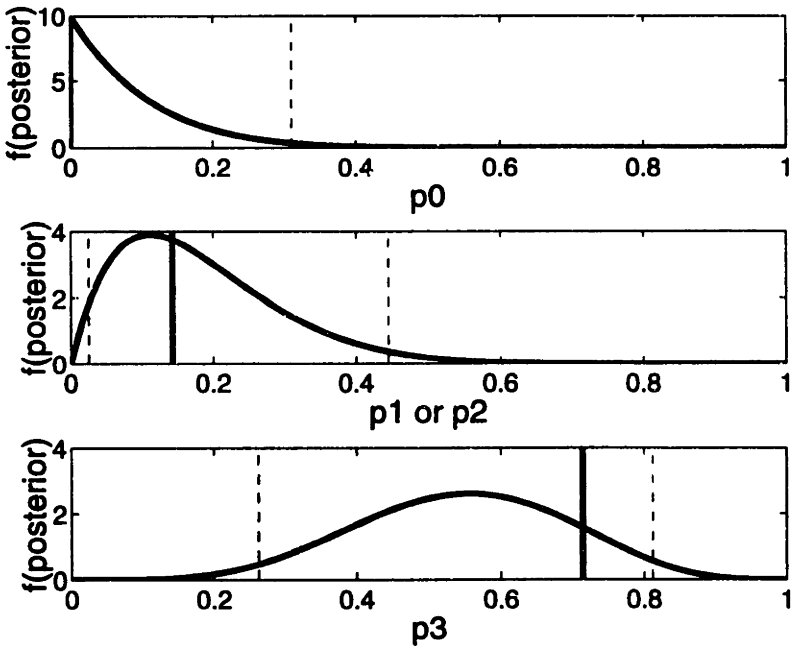

Figure 3-3: Bayesian confidence intervals on JR East alert rule performance, given a uniform prior. The upper figure plots $f_{p_0}(p_0; \text{posterior})$; the middle figure plots $f_{p_i}(p_i; \text{posterior})$ for $i = 1$ and $i = 2$; and the lower figure plots $f_{p_3}(p_3; \text{posterior})$. The dashed lines mark 2.5% and 97.5% quantiles; the interval between the dashed lines contains probability 0.95. The bold vertical lines indicate $\hat{p}_i$, the MLE point estimate of p_i.

five failures for which the system issued STOP alerts), "semi-caught" one of the seven failures (the failure for which the system issued a SLOW alert), and "missed" one of the seven failures (the failure for which the system issued a WATCH alert). Clearly, others could hold different beliefs. If one believed that each type of alert offered equal safety, one might evaluate the performance of the current system as follows: seven out of seven failures "caught" (that is, the system issued a WATCH or SLOW or STOP alert), and no failures "missed." With these beliefs, the current system performed perfectly.

Despite the inherent subjectivity, the benefits of aggregating these performance measures make such an effort worthwhile. We offer model-based procedures to generate a single safety measure and a single delay measure. We are mindful that others may prefer different aggregate measures based on different beliefs.

Aggregate measure of safety performance

For an aggregate measure of safety performance, we will assume that the primary objective of a rainfall alert system is to stop trains before slope failures. Thus, if the alert system is the STOP alert state when a slope failure occurs, we call this a success. Conversely, if the alert rule is in the NO ALERT state when a slope failure occurs, we call this an error, or a "false negative." We propose to use the average annual rate of false negatives for a univariate measure of safety performance. Of course, false negatives are not equivalent to derailments, for these slope failures could be detected by on-foot inspection teams or watchful drivers. (Indeed, this is exactly what happened for two of the seven slope failures in our study.) Nonetheless, as each false negative represents an occasion when the rainfall alert system did not properly provide a warning, we assert the false negative rate is a meaningful measure of safety performance.

Being in the STOP state when a slope failure occurs is clearly a success for the system. Being in the NO ALERT state when a slope failure occurs is clearly a false negative. We shall assume that train drivers are customarily vigilant during rain storms, and thus assume that WATCH alerts offer no additional benefit over the NO

ALERT state.

We need to determine the relative benefit of a SLOW alert relative to a WATCH or NO ALERT when a slope failure occurs, and we introduce a stochastic "driver detection" model for this purpose. We want to capture three effects with this model.

- Because of certain weather conditions (such as fog or snow) or because of certain track geometries (such as vision-obstructing curves), some track failures cannot be seen by an approaching train.

- For those failures which can be seen from an approaching train, the distance before the failure at which the driver can detect the track failure will vary by location, time of day, magnitude of the track damage, and so forth.

- A train traveling at lower speed requires less braking distance to stop. This gives a driver more opportunity to see a track failure early enough to stop his train safely.

To capture the first effect, we shall assume there is a probability p_s that a failure can be seen from an approaching train. We assume p_s reflects track and weather conditions, and thus does not depend on the speed of the approaching train.

To capture the second effect, let us consider the fraction p_s of failures than can be seen by the driver of an approaching train. For these failures, we use random variable D to model the furthest distance from the failure at which the driver can first see it. (For the $1-p_s$ fraction of failures that cannot be seen from an approaching train, random variable D is undefined.) Line-of-sight measurements could be made along a rail line to estimate a reasonable distribution for random variable D. Lacking such empirical data, we postulate that D involves many factors, such as the extent of the track damage, the track curvature, the ambient light level, and the driver's alertness. If we assume that are many of these contributing factors, that they have a multiplicative effect, and that they are independent, then a log-normal distribution is a sensible choice for random variable D. Two features of the log-normal distribution also support such a choice. First, log-normal random variables are non-negative, which is appropriate for D, as D represents a distance. Second, the hazard function for

the log-normal is unimodal; that is, $\Pr(D \in [x, x+dx] | D \geq x)$ as a function of x has a hump. The hazard function describes how the probability of first seeing the failure varies with distance from the failure. Moving towards the failure, this probability first increases, as closer failures are easier to see. After some point, however, this probability decreases, as those failures which still have yet been seen are increasingly difficult to see.

To capture the third effect, we recall from Chapter 2 that the distance required to stop a train experiencing uniform deceleration is $x = v^2/2a$, where x represents the stopping distance, v denotes the train's speed, and a denotes the deceleration rate.

We combine these three effects to produce a driver detection model. This model has five parameters: v, the train speed; a; the train deceleration rate; p_s, the fraction of all slope failures that can be seen from an approaching train; μ, the mean of $\ln D$; and σ, the standard deviation of $\ln D$. We shall use $v = 35$ km/h and $v = 55$ km/h to compare the benefit of the SLOW versus the NO ALERT states, and we shall assume $a = 0.9$ m/s/s, as in Chapter 2. With these values, it takes 17 seconds and 130 meters to stop a train traveling 55 km/h, and 11 seconds and 60 meters to stop a train traveling 35 km/h.

It is not so simple to choose values for p_s, μ, and σ. Lacking empirical data on visibility along the Suigun line, we shall employ values for these parameters that strike us as reasonable, then determine the sensitivity of our conclusions to these values. We caution that these values are speculative, and used only to demonstrate this model. With these caveats, we shall assume that 10% of all slope failures are undetectable from an approaching train, and thus $p_s = 0.90$. We shall assume the median value of D is 100 meters, and thus $\mu = \ln 100 = 4.6$, and we assume $\sigma = 0.5$. Figure 3-4 presents the density and hazard function for D using these values.

We estimate $p_{safe}(v)$, the probability that the driver of a train traveling at speed v sees an approaching slope failure in time to stop his train safely. With probability p_s, the approaching slope failure can be seen from a train. For such a visible failure, the driver must see the failure before he is $v^2/2a$ meters from it, or he will lack sufficient distance to stop the train safely. We use $\Phi(x)$ to denote the cumulative distribution

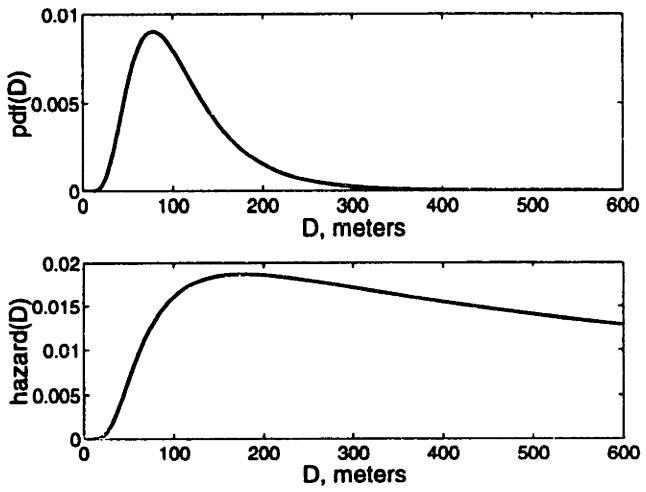

Figure 3-4: Density and hazard function for random variable D, the distance from a track failure at which the driver can first see the failure, for those failures that can be seen from an approaching train. These plots assume $\ln D$ is normally distributed with $\mu = 4.6$ and $\sigma = 0.5$.

of the standard unit normal distribution.

$$\begin{aligned} p_{safe}(v) &= p_s \Pr(D > v^2/2a) \\ &= p_s \left(1 - \Pr(\ln D \leq \ln(v^2/2a))\right) \\ &= p_s \left(1 - \Phi(\frac{\ln(v^2/2a) - \mu}{\sigma})\right) \end{aligned}$$

Figure 3-5 plots $p_{safe}(v)$ against v, for $p_n = 0.9$, $\mu = 4.6$, and $\sigma = 0.5$.

With these speculative parameter values, $p_{safe}(35 \text{ km/h}) = 0.81$, or train traveling at 35 kmh has a $1 - 0.81 = 0.19$ probability of being unable to stop before a track failure. In contrast, $p_{safe}(55 \text{ km/h}) = 0.27$, or train traveling at 55 kmh has a $1 - 0.27 = 0.73$ probability of being unable to stop before a track failure. That is, with this model and these values, running a train at 35 km/h instead of 55 km/h reduces the probability of hitting a track failure roughly four-fold, and thus it takes $0.74/0.19 \approx 4$ slope failures during SLOW alerts to accrue the risk of one slope failure during a WATCH or NO ALERT. Of course, this four-to-one ratio depends on the specific values of p_s, μ, and σ; in particular, $p_{safe}(v)$ appears most sensitive to p_s (Figure 3-6). For purposes of illustration, we assume for the remainder of this

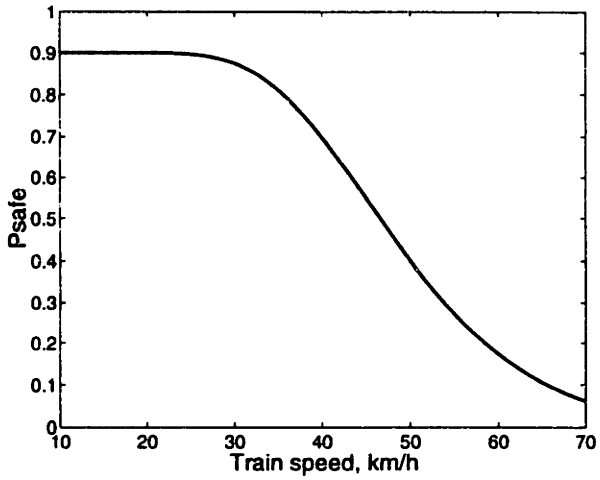

Figure 3-5: The probability the driver of a train traveling at speed v sees an approaching slope failure in time to stop his train safely. This plot assumes $\ln D$ is normally distributed with $\mu = 4.6$ and $\sigma = 0.5$, and $p_n = 0.9$.

Alert level at time of failure	False negative equivalents
NO ALERT	1
WATCH	1
SLOW	1/4
STOP	0

Table 3.4: Conversion factors from alerts to false negative equivalents.

chapter that 4 slope failures during SLOW alerts equal the train derailment risk of 1 slope failure during a WATCH or NO ALERT. Again cautioning that these values are speculative, we thus assign "equivalent false negative" values to each level of alert (Table 3.4).

Aggregate measure of delay performance

For our aggregate measure of delay performance, we define a "service disruption hour" as an hour when all regularly scheduled service is canceled. We also assume that a reduction in train speed leads to a proportional reduction in train service. That is, we assume that each hour the railroad operates in the SLOW state provides (35 kmh)/(55 kmh) = 65% of the customary service. With these assumptions, each operating hour in the SLOW state entails 0.35 hours of service disruption, and each operating in the

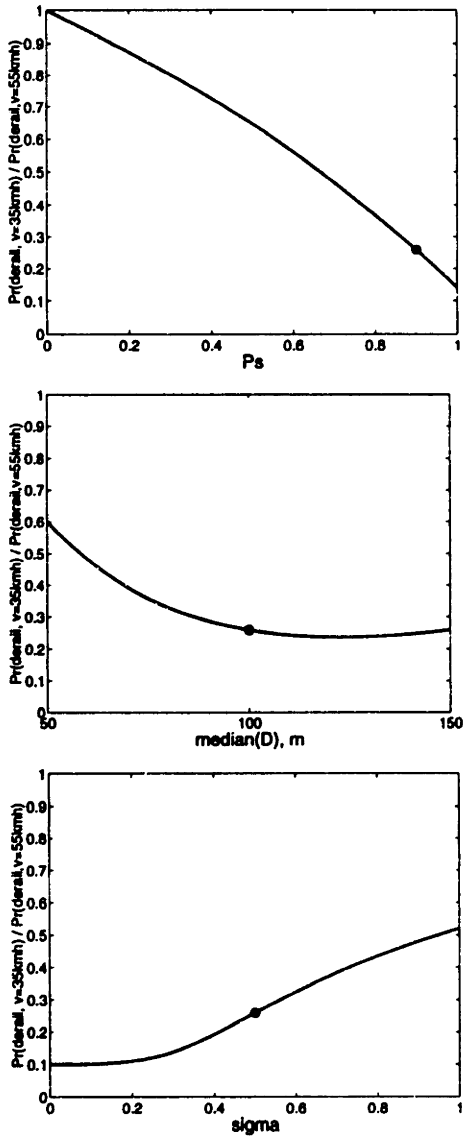

Figure 3-6: Sensitivity analysis of the driver detection model to p_s, to median$(D) = e^{\mu}$, and to σ; for $a = 0.9$ m/s/s, and $v = 35$ and $v = 55$ km/h. The dot on each plot indicates the base-case value.

Each hour in alert level	Equivalent service disruption hours
NO ALERT	0
WATCH	0
SLOW	0.35
STOP	1.0

Table 3.5: Conversion factors from delay hours to disruption hour equivalents

SAFETY
1/3 false negative equivalents/year
DELAY
11 service disruption hours/gauge/year

Table 3.6: Derived aggregate performance measures of the current rainfall alert rule on the Suigun line, 1991–1994.

STOP state represents 1.0 hours of service disruption (Table 3.5).

Aggregate performance measures for current alert system

Using Tables 3.3, 3.4, and 3.5, we apply the results of the last two sections to calculate univariate measures of safety and delay for the current rainfall alert rule (Table 3.6). We note these calculations are rather speculative, as the underlying models have not been properly calibrated on actual data.

$$
\begin{aligned}
FN &= \text{false negatives equivalents} \\
&= 5(0) + 1(\frac{1}{4}) + 1(1) + 0(1) = 1.25 \\
&= 1.25/3.75 \approx 1/3 \text{ false negative equivalents/year} \\
SDH &= \text{service disruption hours} \\
&= 4.8(0.35) + 9.2(1) \approx 11 \text{ service disruption hours/gauge/year}
\end{aligned}
$$

We also wish to assess the uncertainty in random variable FN. To do this, we hold the total number of slope failures fixed at seven, and explore the variation in FN due to our uncertainty in $\vec{p}$, the conditional probability the current rainfall alert rule

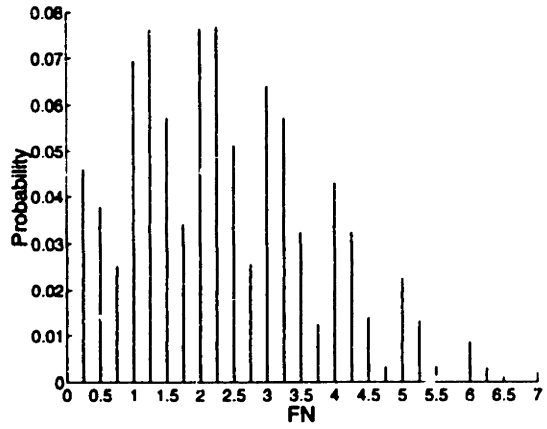

Figure 3-7: Probability mass function for FN, from Equation 3.3.

provides alert level i at the time of a slope failure. Given $\vec{p}$, let (N_0, N_1, N_2, N_3) follow a Multinomial$(7, p_0, p_1, p_2, p_3)$ distribution. Let $FN = N_0 + N_1 + (1/4)N_2$. Using our Bayesian posterior for $\vec{p}$ from Equation 3.2, we can determine the distribution of FN.

$$\Pr\left(\vec{N} = \vec{n}; \vec{p}\right) = \frac{(n_0 + n_1 + n_2 + n_3)!}{n_0! n_1! n_2! n_3!}(1 - p_1 - p_2 - p_3)^{n_0}(p_1)^{n_1}(p_2)^{n_2}(p_3)^{n_3}$$

$$\Pr(N_0 + N_1 = i, N_2 = j; \vec{p}) =$$

$$\sum_{n_0=0}^{i} \Pr(N_0 = n_0, N_1 = i - n_0, N_2 = j, N_3 = 7 - i - j; \vec{p})$$

$$\Pr\left(NF = i + \tfrac{j}{4}\right) = \int_{\vec{p} \in S} \Pr(N_0 + N_1 = i, N_2 = j; \vec{p}) f(\vec{p}; \text{posterior}) d\vec{p} \qquad (3.3)$$

As seen in Figure 3-7, there is considerable dispersion in this distribution for FN. This dispersion arises from both the uniform prior used in Equation 3.1 and the small sample of seven slope failures. That is, while the current record of no "missed" alerts is impressive, we lack sufficient data to prove that this is more than just a run of good luck.

3.2.5 Performance of modified versions of current rule

The current rainfall alert rule has several intricacies: the gauge control regions are complex (Figure 3-1); the gauge alert threshold functions are complex (Figure 3-2);

and the system employs fine temporal resolution precipitation data (Table 3.1). In Sections 3.4 and 3.5, we shall propose and evaluate a family of alternative rainfall alert rules. These rules will use simpler gauge control regions. In these alternative rules, each gauge will bear responsible for shutting down the interval of track which runs midway from its neighbor to the left to midway to its neighbor to the right. That is, these alternative rules will use non-overlapping gauge control regions, where each meter of track will fall under the responsibility of the gauge closest to it. We call these "closest gauge" control regions. Our motivation for considering such simplified alternative rules is to isolate the effect of varying the shape of the exponential filter. In other words, our goal is to determine whether the short-term or long-term rainfall trend is best suited to issuing slope failure alerts. By comparing the performance of these simplified alternative rules to the current system, we may also determine the value of current rule's overlapping gauge control regions.

We shall evaluate these alternative rules using two datasets: the JR East data described in Section 3.2.2, and an additional dataset from the Japanese national meteorological agency. This national dataset contains hourly precipitation data, a temporal resolution twelve times coarser than the five-minute JR East rainfall data.

Table 3.3 indicates the actual performance of the JR East rule. It will be helpful to know how the performance of the current rule changes upon slight modification.

Our first modification tests the sensitivity of the current rule to the temporal resolution of the precipitation data. Let us take the five-minute JR East precipitation data and aggregate these data hourly. Where the original dataset provided precipitation data for twelve five-minute intervals each hour (0 to 5 minutes, 5 to 10 minutes, 10 to 15 minutes, and so on), our derived dataset sums these to provide only one interval for each hour (0 to 60 minutes).

Why would such a change affect the performance of the current rule? The answer lies in the definition of the HRAIN index. HRAIN measures precipitation over the last hour. If the precipitation time-series varies rapidly, the hourly rain computed at five minute intervals could well exceed the hourly rain computed only on hour boundaries. Analysis by the Safety Research Laboratory suggests that this is indeed the case. Let

HRAIN5 denote the HRAIN index computed every five minutes, and MHRAIN5 denote the hourly maximum of the twelve HRAIN5 values. Let HRAIN60 denote the HRAIN index computed every hour. The Safety Research Laboratory suggests that, on average, MRAIN5 exceeds HRAIN60 by 50%. [SRL95i] Using hourly data degrades the current alert rule, as its design presumes MHRAIN5 values.

When we apply the current rule to our derived hourly dataset, we find the rule issues a STOP alert for one of the seven failures, WATCH alerts for two of the seven failures, and NO ALERT for four of the seven failures. We conclude that using the current alert rule with only hourly data would lead to a significant degradation in safety performance. (Of course, JR East receives instantaneous measures of rainfall from its network of rainfall gauges, so the railroad avoids any problems arising from data aggregation.)

Our second modification tests the sensitivity of the current rule to the geometry of the gauge control regions. We apply the current alert rule to the five-minute dataset, but we use the "closest gauge" control regions previously described. Upon this modification, we find the rule now issues STOP alerts for three of the seven failures, WATCH alerts for three of the seven failures, and NO ALERT for one of the seven failures. Again, we conclude using the current alert rule with "closest gauge" control regions results in a significant degradation in safety performance. Thus, the present system of overlapping gauges seems well-conceived.

For our third modification, we make both previous modifications simultaneously. Applying the rule to our derived hourly dataset using the "closest gauge" control regions, we find the rule now issues STOP alerts for one failure, WATCH alerts for one failure, and NO ALERT for five failures.

Table 3.7 summarizes these various modifications. We conclude that the current rainfall alert rule performs properly when it is used as intended; that is, when it is used with overlapping gauge control regions on frequent precipitation data. The performance of the current rule degrades significantly when these criteria are not met.

Distribution of seven failures
across alert levels

	Modifications to the current system	NO ALERT	WATCH	SLOW	STOP	Equivalent false negatives
0.	none	0	1	1	5	1.25
1.	hourly data	4	2	0	1	6
2.	"closest gauge"	1	3	0	3	4
3.	both 1 and 2	5	1	0	1	6

Table 3.7: Performance degrades upon modifying the current alert rule.

3.3 Alternative methodologies

Before proposing alternatives to the current system, we offer a brief review of common methodologies for issuing slope failure warnings from precipitation data.

Using mechanistic models to predict slope failure

Rainfall itself does not cause slope failures. Rather, heavy rainfall saturates the ground and increases soil pore pressure. These phenomena may destabilize a slope or embankment and cause failure. [Han84] One approach for predicting slope failure is to model these underlying physical processes directly. To avoid the complexity of a full hydrological analysis, the slope is often modeled as a simpler physical system, such as a series of tanks or leaky buckets.

For example, Kobashi and Suzuki employ a tank model to predict slope failure in mountainous regions in Japan. [KS87] Their model consists of a hypothetical series of linked tanks. Rainfall fills the uppermost tank, which then discharges at some rate into a cascade of lower tanks. Excessive rain may cause these tanks to back up and overflow, the model's equivalent of a slope failure. The difficulty with such tank models is they require estimating many site-specific parameters.

Wilson and Wieczorek propose a simpler mechanistic model to predict slope failures and debris flows in central California. [WW95] They model the reservoir effect of a hillside with the analogy of a leaky barrel. Their model has two parameters, one for soil drainage and one for soil porosity. After tuning these parameters for a particular hillside, Wilson and Wieczorek report that the model predicts the response of hill-

side pore pressure to rainfall reasonably well, as verified by actual site measurements. Wilson and Wieczorek then use this estimated hydrological measure to predict the onset of slope failures and debris flows.

Using a rainfall index to predict slope failures

Instead of modeling soil hydrodynamics directly, simpler methodologies attempt to capture such effects by applying an appropriate filter to rainfall data. This filtering process yields a "rainfall index." Sirangelo and Versace offer a general description of this approach. [SV94] Let $Y(t)$ denote the rainfall intensity at time t, $\psi(t)$ denote a nonnegative filter function, c denote a scaling constant, and $R(i)$ denote the filtered rainfall index at time t. The rainfall index is obtained by convolving the precipitation data with the filter.

$$R(t) = c \int_0^t Y(\tau)\psi(t-\tau)d\tau$$

This index is then used to predict slope failures. The filter function and an alert threshold are chosen empirically, after studying previous storms and failures. For such an approach to be successful, one must choose a function $\psi(t)$ that captures the typical pattern of antecedent rainfall before a failure while simultaneously minimizing the number of spurious alerts.

Using an exponential rainfall index to predict slope failures

One particular choice for $\psi(t)$ is an exponential filter.

$$\begin{aligned} \psi(t) &= \exp(-\lambda t), \qquad \lambda > 0, t > 0 \\ R(t) &= \int_0^t Y(\tau)\exp(-\lambda(t-\tau))d\tau \end{aligned}$$

This filter has one parameter, λ, which indicates how quickly the filter decays. As precipitation data are usually observed at discrete intervals evenly spaced in time, we replace this integral with a sum. This form leads to a convenient recursive definition

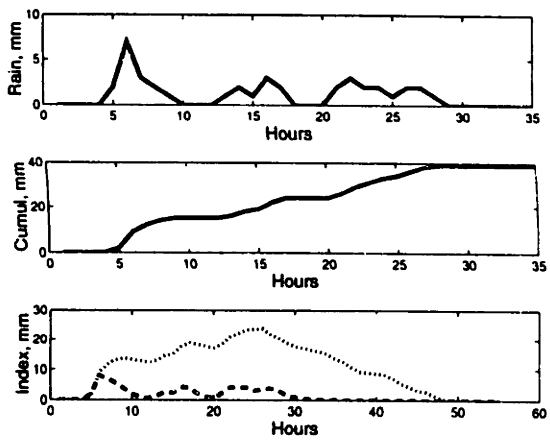

Figure 3-8: Demonstrating the exponential rainfall index. The upper plot indicates hypothetical hourly rainfall. The middle plot indicates cumulative rainfall. The lower plot indicates the response of two exponential rainfall indices: $h = 1$ hour (dashed line), and $h = 20$ hours (dotted line).

of the discrete-time rainfall index, R_t.

$$\begin{aligned} R_t &= \sum_{k=0}^{t} (Y_k) e^{-\lambda(t-k)} \\ &= Y_t + (e^{-\lambda}) R_{t-1} \end{aligned}$$

It is convenient to define a parameter $h = \ln 2/\lambda$, called the "half-life" of the index. The half-life indicates the amount of time it takes for the effect of the current period on the index to drop to one half its original value. A small value of h produces a large value of λ, yielding an index that resembles the pattern of recent rainfall. A large value of h produces a small value for λ, yielding an index that resembles the cumulative rainfall pattern (Figure 3-8).

To use an exponential rainfall index to issue slope failure alerts, one typically chooses a half-life and an alert threshold empirically, after studying data from previous storms and failures. We take this approach in Section 3.4, where we evaluate the performance of exponential rainfall indices for providing slope failure warnings for the Suigun line.

Using rainfall intensity and duration to predict slope failures

Perhaps the most common technique for predicting slope failures using precipitation data is the "intensity vs. duration" method. [CE95] This approach employs a collection of n indices, $\{I(D_1), I(D_2), \ldots, I(D_n)\}$, where $I(D)$ denotes the mean rainfall intensity over the last D time units. We fix our time origin at the onset of the current storm.

$$I(D) = \frac{1}{D} \int_{t=0}^{D} Y(t)dt$$

For index $I(D_j)$, a critical threshold $I^*(D_j)$ is calculated,

$$I^*(D) = \alpha D^{-\beta};$$

and if $I(D)$ exceeds $I^*(D)$, a slope failure alert is issued.

Often cited parameters for this model come from Caine, who suggests $\alpha = 14.8$ and $\beta = 0.39$, when D is measured in hours and $I(D)$ in millimeters of rain per hour. [Cai80] Caine derives these values by studying 73 rain-induced slope failures from around the world. He fit these parameters to place nearly all of the storms accompanying these failures above the critical threshold. However, Caine does not report how often a storm drives these indices above the threshold without causing an accompanying landslide. This conceptual blunder—estimating the distribution of a rainfall index at times when slope failures occurs instead of estimating the probability of a slope failure conditional on the value of a rainfall index—occurs several times throughout the literature. [OS94]

The US National Weather Service and US Geological Survey use such an $I(D)$ vs. D relationship to issue real-time landslide alerts for the San Francisco Bay region in California. [K+87] These agencies report success using this method. [USGS95]

3.4 The exponential rain index

The Safety Research Laboratory is considering replacing the current rainfall alert rule with a new system based on an exponential rainfall index. The motivation for

this change is the desire for a single index rule, and the hope that such a rule will offer superior performance. [SRL96a] Implementing such a change over the entire rail network would be a formidable task. The Safety Research Laboratory would need to determine appropriate half-lives, alert levels, alert thresholds, and gauge control regions for every rainfall gauge across the network.

In this section, we address a set of related but easier issues. Restricting our attention to a simple family of exponential rules, we investigate their performance at predicting slope failures along the Suigun line. To obtain a larger population of slope failures, we also study the performance of such rules using a precipitation dataset from the Japanese national meteorological agency. We also attempt to quantify the uncertainties in these performance estimates.

3.4.1 Performance of single exponential rules on JR East data

This section investigates the performance of slope failure alert rules based on a single exponential rainfall index. We restrict our attention to rules of the following simple form, where each rule has two parameters, r^* and h.

- Each gauge maintains an exponential rainfall index R_t with half-life h.
- Each gauge has a constant critical threshold, r^*.
- During all hours when $R_t \geq r^*$, the gauge issues a STOP alert. All other hours, the gauge is in the NO ALERT state.
- All gauges use the same value of h and r^*.
- The control region of each gauge runs from halfway to its neighboring gauge on the left to halfway to its neighboring gauge on the right.

We use the following discrete version of the single exponential index.

$$R_t = \sum_{k=0}^{t} (Y_k) e^{-\lambda(t-k)}$$

$$= Y_t + (e^{-\lambda})R_{t-1}$$

$$h = \ln 2/\lambda$$

Recall that h, the half-life of the index, indicates the amount of time it takes for the effect of the current period on the index to drop to one half its original value. When h is small, λ is small, and the index resembles the pattern of recent rainfall. When h is large, λ is small, and the index the pattern of cumulative rainfall.

Note that these simplified rules are more restrictive than the current rainfall alert rule. In contrast to these simplified rules, the current alert system allows gauges to issue three different types of alert (WATCH, SLOW, and STOP); the current system allows gauge control regions to vary by gauge; the current system allows these gauge control regions to overlap; and the current system allows critical thresholds to vary by gauge.

We assess these exponential rules using the same two aggregate performance measures discussed previously: the annual rate of false negatives, and the average annual number of service disruption hours per gauge. As these exponential rules only issue STOP alerts, there is a one-to-one correspondence between false negatives and the number of slope failures that occur in the NO ALERT state, and there is a one-to-one correspondence between service disruption hours and hours in the STOP alert state.

Retrospective performance assessment

For this first analysis, we take a "retrospective approach." That is, for each h and r^*, we determine how an exponential alert rule with these parameters would have performed had it been in use along the Suigun line between 1991 and 1994. Note that such rules, calibrated to fit the slope failures after these failures had occurred, would be expected to perform better than the current rainfall alert rule, which was designed before the slope failures occurred.

One method to address this "retrospective advantage" to divide one's sample into two subsamples. The first subsample is used to determine good parameters for a rule. The second subsample is used to evaluate how well these "good" parameters

generalize to fresh data. Reasonable positive correlation between the performance on both subsamples may suggest that these performance estimates are robust. Little or negative correlation between the performance on both subsamples may suggest over-fitting or data mining.

With only seven failures along the Suigun line between 1991 and 1994, it makes little sense to partition this sample into two subsamples. Instead, we take two approaches. First, aware of the limitations of this technique, we calibrate and test this collection of exponential rainfall alert rules retrospectively using the JR East precipitation data. Second, we employ a precipitation dataset from the Japanese national meteorological agency which includes more slope failures, allowing us to demonstrate the two subsample method.

Evaluation procedure

We now describe how we evaluate an exponential rule with half-life h and alert threshold r^*. We limit our attention to so-called "homogeneous rules," where each of the seventeen gauges employs the same values of h and r^*. This first analysis is retrospective: we determine how an exponential alert rule with these parameters would have performed between 1991 and 1994 along the Suigun line.

Let $Y(t,i)$ denote the raw precipitation recorded at gauge i over hour t. Let $t^*(j)$ denote the time of the jth slope failure, and $g^*(j)$ denote the closest gauge to the jth slope failure. Let $R(t,i,h)$ denote the filtered value of the exponential index with half-life h at gauge i at time t. For each value of h, we calculate this index for every hour for every gauge in the JR East dataset using the following recursive relationship.

$$R(t,i,h) = Y(t,i) + \exp(\frac{-\ln 2}{h})R(t-1,i,h)$$

We also compute the empirical complementary cumulative distribution of this index at each gauge.

$$G(r,i,h) \;=\; \Pr(R \geq r \text{ at gauge } i \text{ at a random hour; half-life}=h)$$

$$\approx \text{ (fraction of hours with } R \geq r \text{ at gauge } i\text{; half-life}=h)$$

We first examine the delay performance of an exponential rule with half-life h and alert threshold r^*. The probability that the rainfall index at gauge i at a random hour equals or exceed r^* is $G(r^*, i, h)$. With 6,500 operating hours each year, gauge i spends an average of $(6500)G(r^*, i, h)$ hours in the STOP state, and we count each of these STOP hours as one service disruption hour. We average these estimates by gauge to obtain the mean annual service disruption hours per gauge.

$$SDH(r^*, h) = \frac{6500}{17} \sum_{i=1}^{17} G(r^*, i, h)$$

With our notation, gauge $g^*(j)$ is the closest gauge to failure j; accordingly, it bears the responsibility of issuing an alert for this failure. Failure j occurs at hour $t^*(j)$. The value of the rainfall index at gauge $g^*(j)$ at hour $t^*(j)$ is $R(t^*(j), g^*(j), h)$. If this index is equal to or exceeds r^*, this gauge will issue an alert for this failure. If this does not occur, failure j will occur during the NO ALERT state, and we shall count this as a false negative.

$$FN(r^*, h) = \text{ \# of slope failures for which } R(t^*(j), g^*(j), h) < r^*$$

Due to the retrospective nature of this analysis, and given a sample of seven failures, FN takes on only eight values, ranging from zero (a rule which "catches" all seven slope failures) to seven (a rule which "misses" all seven slope failures). Accordingly, for a fixed value of h, the function $FN(r^*, h)$ resembles an increasing staircase function of r^*, with horizontal "steps" of varying width. Let $R^*_{(j)}(h)$ denote the jth order statistic of the set $\{R(t^*(j), g^*(j), h); j = 1 \ldots 7\}$; that is, $R^*_{(j)}(h)$ denotes the maximum value of the rainfall index with half-life h that catches $(7 - j + 1)$ of the seven failures. Recall that j ranges from to 7. When $j = 1$, the rule catches seven of the seven failures; when $j = 7$, the rule catches one of the seven failures; and when $r^* > R^*_{(7)}(h)$, the rule misses all seven failures.

Consider a rule that uses half-life h and alert threshold $r^* = R^*_{(j)}(h)$. Of all rules

with half-life h that catch $(7-j+1)$ failures, the rule with threshold $R^*_{(j)}(h)$ offers the minimum estimated disruption, and so we call this rule the "best" retrospective rule with half-life h for catching $(7-j+1)$ failures. As this assessment procedure is retrospective, it suffices to determine only these "best" rules for each half-life. Let us review this evaluation procedure in detail.

1. Determine the time of each failure, $t^*(j)$.
2. Determine the closest gauge to each failure, $g^*(j)$.
3. Calculate rainfall index at each hour at each gauge, $R(t,i,h)$.
4. Determine the complementary cumulative distribution at each gauge, $G(r,i,h)$.
5. Determine the set of "best" alert thresholds, $R^*_{(j)}(h)$, which achieve the lowest service disruption among all rules that catch $(7-j+1)$ of the 7 failures.
6. Determine the false negative performance of these thresholds, $FN(r^*,h)$.
7. Determine the service disruption hours of these thresholds, $SDH(r^*,h)$.

Performance analysis of single exponential rules

We apply this procedure to evaluate the retrospective performance of a family of exponential index rules on the JR East rainfall data for the Suigun line between 1991 and 1994. We investigate eighteen values for the half-life of the index.

$$h \in \{1,2,3,4,5,10,15,20,25,30,35,40,45,50,60,70,80,90\} \text{ hours.}$$

(The half-life of the index indicates the amount of time it takes for the effect of the current period on the index to drop to one half its original value.) Using the letter "A" to denote $h = 1$, "B" to denote $h = 2$, and so on, up to "R" for $h = 90$, we identify each exponential rule with a string consisting of this upper case letter prepended to an integer indicating r^*. For example, the exponential rule "A27" uses a $h = 1$ hour half-life and a $r^* = 27$ mm alert threshold. The rule "K71" uses a $h = 40$ hour half-life and a $r^* = 71$ mm alert threshold.

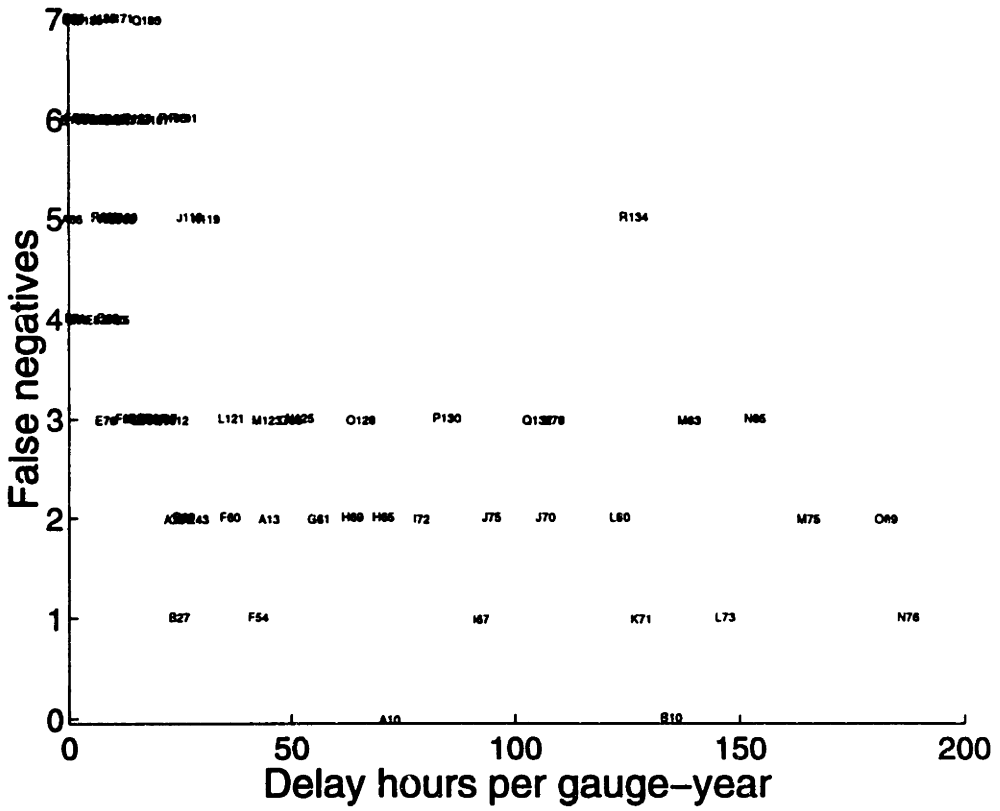

Figure 3-9: Performance of "best" exponential alert rules for the Suigun line between 1991 and 1994, using JR East precipitation data.

We determine the set of eight "best" rules for each value of h using the previous procedure. For each of these rules, we plot its safety vs. disruption performance using the string denoting h and r^* as a plotting symbol (Figure 3-9). We observe that most of the rules on the efficient horizon are of the "Axx" or "Bxx" type. This suggests that rules with short half-lives offer superior retrospective performance. This finding agrees with the design of the current JR East rainfall alert rule. In the current rule, alerts are more often triggered by HRAIN moving upwards than by CRAIN moving outwards (Figure 3-2). While the shape of the filters differ—HRAIN employs a one-hour rectangular filter, and an "Axx" rule employs a one-hour exponential filter—both rules place their primary emphasis on recent rain.

Though it is less than fair to compare the best of these retrospective rules to the performance of the current rule, we shall do so nonetheless. We find the current

rule outperforms all of these exponential rules, offering 1.25 false negative equivalents and an annual average of 11 service disruption hours per gauge (Table 3.6). This performance places the current alert rule well below the efficient horizon on Figure 3-9, indicating Pareto superiority. We suggest, however, that this may be an unfair comparison, given the restrictions (non-overlapping gauge control regions, same threshold value at each gauge, same half-life at each gauge, hourly data, etc.) we have imposed on these exponential rules. We recall that when we modify the current rule to employ hourly data and closest-gauge control regions, its performance degrades to 6 false negative equivalents and an annual average of 3 service disruption hours (Section 3.2.5). This performance places the modified current alert rule well above the efficient horizon, thick within a cloud of similarly-performing exponential rules. We again note that the present system of overlapping gauges seems well-conceived. Finally, we note that we lack sufficient data to conclude this estimate of 1.25 annual false negative equivalents is an accurate representation of the safety performance of the current system (recall Figure 3-7).

Complementary cumulative plots

A certain graphical presentation of the functions $G(y, i, h)$ and $R^*_{(j)}(h)$ offers considerable insight into the retrospective performance of these exponential alert rules. For each failure j, we graph $G(R^*_{(j)}(h), g^*(j), h)$ against h. We call this a "complementary cumulative plot" (Figure 3-10). Each line on this plot corresponds to one slope failure. These lines indicates where each failure falls in the distribution of rainfall indices as a function of half-life, thus summarizing the entire spectrum of retrospective safety vs. delay tradeoff in one graph.

For example, consider the line marked "2," corresponding to the second slope failure. Let us first employ a rainfall index with $h = 1$ hour. After computing this $h = 1$ index for every hour in the dataset, we compute this index at the hour of the second slope failure. We find the value of this index at the time of the second slope is very high relative to the value of the index of all other hours in the dataset. At the hour of the second failure, the value of the $h = 1$ hour index fell in 99.8th percentile of

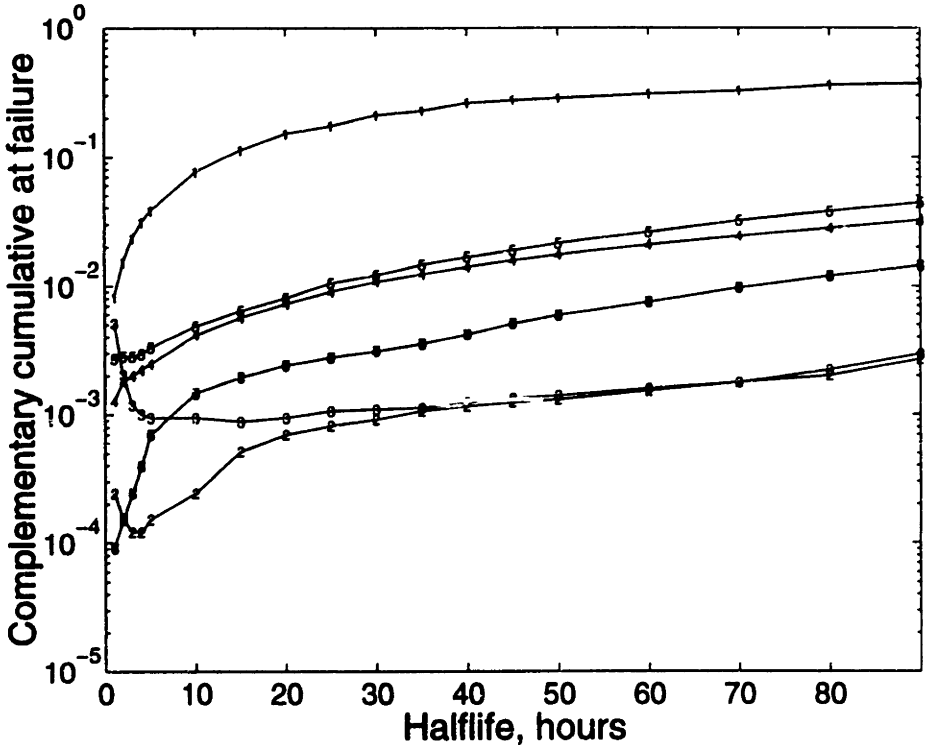

Figure 3-10: Complementary cumulative plot of exponential rainfall index rules.

its distribution. On Figure 3-10, we note the line marked "2" at $h = 1$ has the value $1 - 0.998 = 0.002$. Next, we employ a rainfall index with $h = 2$ hour half-life. We repeat the previous procedure for this index, and find that at the hour of the second failure, the value of the $h = 2$ hour index fell in 99.9th percentile of its distribution. As we begin to increase h, we find the value of the corresponding index at the hour of the second failure moves to right within its distribution (into the higher quantiles), indicated by the initial drop in the line 2. Once h exceeds 3 hours, we find the index moves leftward within its distribution, indicated by the subsequent rise in the line 2. This pattern means the rainfall before the second failure is best detected with a index using a two or three hour half-life, for such indices place this storm as far right as possible in the distribution of the corresponding exponential rainfall index.

We may use this complementary cumulative plot to determine the half-life of that rule which offers minimum delay and catches all seven failures. To determine rules that catch all seven failures, we choose $r^* \leq R^*_{(1)}(h)$. Graphically, we choose a point located on or above the highest curve on the complementary cumulative plot. To find that rule which catches all seven failures with minimum delay, we choose the lowest point above all seven curves. In Figure 3-10, this point is the left-most point on the highest curve. This procedure generalizes easily: to find the half-life of that rule which offers the minimum delay while catching k of the 7 failures, one finds the lowest point on the graph with $(7 - k)$ lines above it.

This complementary cumulative plot also offers insight into the rainfall pattern before each slope failure. Some lines on this plot are monotonically increasing, such as those marked "1," "4," and "5." These lines correspond to slope failures best detected with short half-life indices. This suggests that the storms preceding these failures were brief and intense rain bursts. In contrast, some lines on the plot decrease before increasing, such as those marked "2" and "3." These lines correspond to slope failures best detected with intermediate half-life indices. This suggests that the storms preceding these failures were extended periods of rain. Some of the lines on this plot cross one another, indicating that different half-lives offer different performance for different storms. This suggests that combining two indices with different half-lives

2c. By design, the rule using $r^* = R^*_{(j)}(h)$ and half-life h catches $(16-j+1)$ of the 16 failures in the first subsample; it thus has a false-negative frequency of $(j-1)/16$.

2d. Determine $SDH(r^*, h)$, the delay performance of this rule.

3. Without changing r^*, apply this rule to the second subsample. Determine k, the number of slope failures in the second subsample not caught by this rule. This rule has a false-negative frequency of $k/16$ on the second subsample.

4. As we do not change r^*, $SDH(r^*, h)$ does not change.

5. For half-life h, using the first subsample, we have found the minimum delay rule that offers a false-negative frequency of $(j-1)/16$. On the second sample, this same rule offers a false negative frequency of $k/16$ with the same average annual service disruption hours.

Performance analysis

After this procedure, we plot the safety vs. delay performance for each exponential rule (Figure 3-13). As each rule has two safety measures—its false-negative frequency on subsample one, and its false-negative frequency on subsample two—each rule is represented as line segment. Because delay is the same in both cases, this line segment is parallel to the y axis. If a rule performs better in the "tuning" subsample than in the "testing" subsample, the top of this line segment indicates its performance on subsample one and the bottom of this line segment indicates its performance on subsample two. If a rule performs better in the "testing" subsample than in the "tuning" subsample, the top of this line segment indicates its performance on subsample two and the bottom of this line segment indicates its performance on subsample one. If a rule performs equally well in both samples, this segment degenerates into a point. In any case, the length of the line indicates the discrepancy in the rule's safety performance between the two subsamples. The average of these two measures indicates a rule's safety performance on the full sample of 32 failures; these averages are plotted as black dots in Figure 3-13.

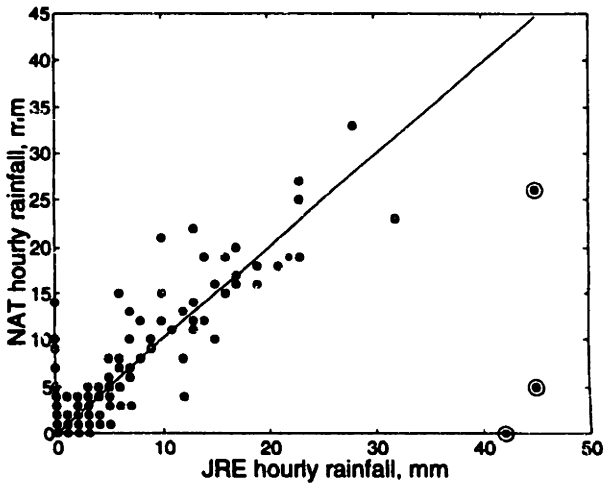

Figure 3-12: Similarity of pre-failure rainfall across the two precipitation datasets. This plot presents $Y_{NAT}(t^*(j) - \tau, g^*(j))$ vs. $Y_{JRE}(t^*(j) - \tau, g^*(j))$ for $j = 1 \ldots 7$ and $\tau = 0 \ldots 24$, or the twenty-four hours before each failure at the closest gauge to the failure. The line indicates the least-squares regression fit $y = 0.99x + 0.08$. The three circled outlier points were omitted from this regression.

catches $(16 - j + 1)$ of the sixteen slope failures in the first subsample and causes $SDH(R^*_{(j)}(h), h)$ average annual hours of service disruption. Instead of using the raw count of false negatives as a measure of safety performance, we divide this count by sixteen, obtaining the fraction of false negatives or the frequency of missed failures. We next determine how well these same rules perform providing warnings for the other sixteen slope failures in the second subsample. Instead of using the raw count of false negatives as a measure of safety performance, we again divide this count by sixteen to obtain the frequency of false negatives. This allows us to compare how a rule tuned on the first subsample performs on the second subsample. Let us describe this procedure formally.

1. Select a half-life h.

2. Select a safety performance level, j. We now determine the minimum delay rule with half-life h that detects $(16 - j + 1)$ of the failures in the first subsample.

   2a. Determine the complementary cumulative distributions, $G(y, i, h)$.

   2b. Using only the first subsample, determine $R^*_{(j)}(h)$.

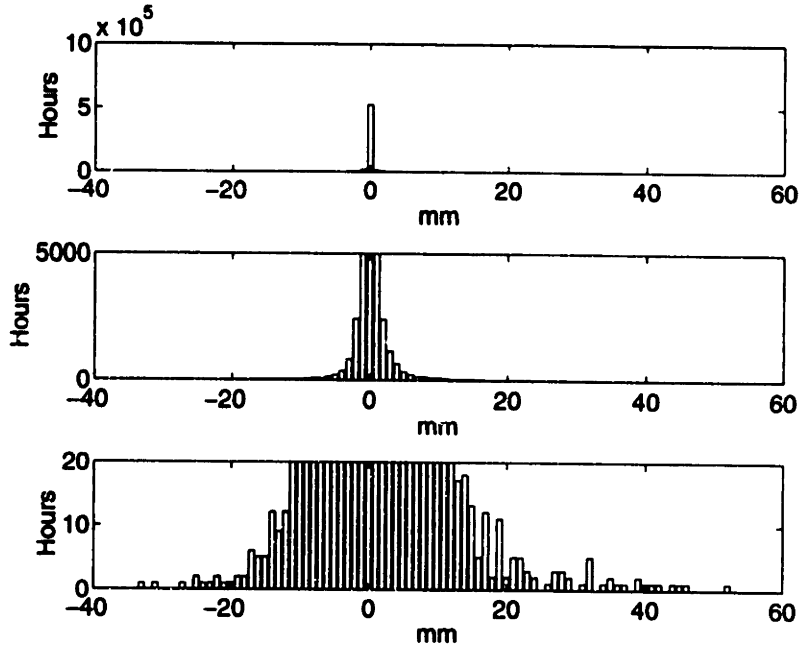

Figure 3-11: Histogram of $Y_{JRE}(t,i) - Y_{NAT}(t,i)$, the hourly discrepancies between the JR East and national precipitation datasets, 1991–1994. The three plots present the same data on different scales. This histogram is presented to demonstrate the lack of any gross overall discrepancies between the two datasets.

well (Figure 3-12).

Evaluation procedure

As the national dataset includes 32 slope failures, we can use the "two subsample" approach to evaluate the performance of this same family of exponential rules. We number the slope failures chronologically in time, then assign the odd-numbered failures to the first subsample and even-numbered failures to the second subsample. By alternating the assignments in this way, we hope to generate two representative and unbiased subsamples. We use the first subsample to choose parameters for these rules, and the second subsample to test these rules.

We begin with the first subsample of sixteen failures. For each value of h, we perform the same analysis as in the previous section, but we restrict our attention to only those slope failures in the first subsample. After this procedure, for a given value of h, we have a set of "best" r^* thresholds, where the rule that uses $r^* = R^*_{(j)}(h)$

into a single hybrid alert rule might offer superior performance; Section 3.5 considers this possibility.

3.4.2 Performance of exponential rules on national data

The JR East precipitation dataset for the Suigun line covers the 1991 to 1994 time period. Seven serious slope failures occurred along the Suigun line during these years. This small sample makes it difficult to reach statistically solid conclusions about the relative performance of various rainfall alert rules. To ameliorate this difficulty, we employ a larger precipitation dataset provided to us by the Safety Research Laboratory derived from data collected by the Japanese national meteorological agency. This agency maintains and monitors a network of rainfall gauges across Japan. The national rainfall gauges are spaced about every 30 kilometers. The locations of these national gauges do not necessarily coincide with the locations of the seventeen JR East gauges along the Suigun line, so the Safety Research Laboratory used a spatial interpolation scheme to estimate rainfall at their gauge locations using the national data. [SRL95i] This interpolation technique is a form of averaging, and averaging reduces the variance in the sample. In the context of precipitation data, we might expect that spatially interpolated rainfall data would show smaller extreme values.

This derived "national" dataset provides hourly rainfall observations for the seventeen gauges from the late 1970s to 1994. Starting dates differ by gauge. Beyond the seven failures between 1991 and 1994, the time period covered by this dataset includes an additional 25 serious slope failures on the Suigun line before 1991. Thus, this precipitation database allows us to increase our sample of slope failures from seven to thirty-two. Both the JR East and the national dataset contain precipitation data for 1991 to 1994. We use this overlapping period to check the consistency of the two datasets. If we compare the two datasets gauge by gauge, hour by hour, we find the hourly discrepancies are typically small, symmetric, and centered at zero (Figure 3-11). For our purposes, we are most concerned with the consistency of these two datasets over those critical hours preceding slope failures. Save one or two outlier hours, there is good agreement between the two rainfall datasets over these hours as

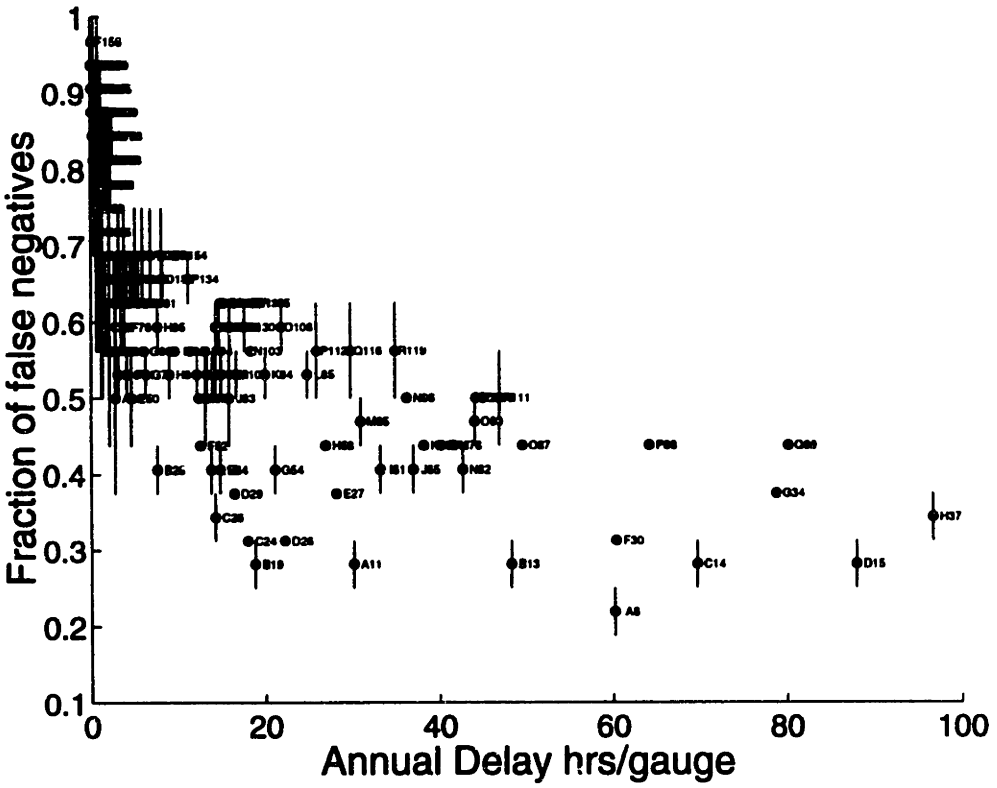

Figure 3-13: Uncertainty in performance of exponential alert rules. Each line represents one exponential rule; each end of the line indicates the rule's safety performance on one of the two subsamples. The black dot indicates the rule's safety performance on the full sample of 32 failures. Precipitation data from the national dataset.

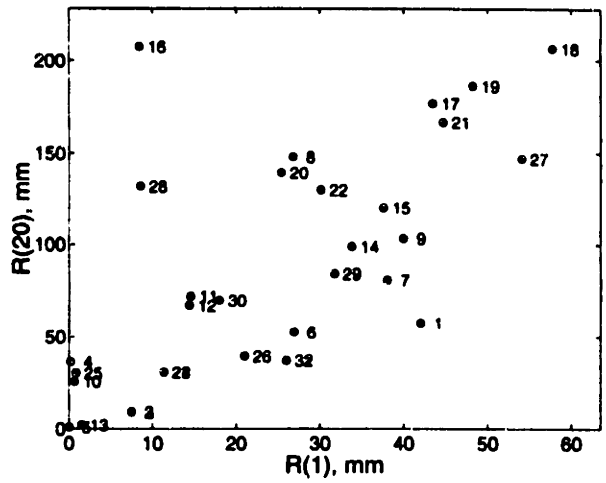

Figure 3-14: The $h = 1$ hour index vs. the $h = 20$ hour index at hours with failures. That these points do not entirely fall on a line suggests these two indices capture different phenomena.

We note that these performance discrepancies across the two subsamples are usually small. That is, a rule that catches k out of the sixteen failures in the first subsample catches k or $k - 1$ out of the sixteen failures in the second subsample. This experiment indicates that the general conclusions of the previous section are robust. As before, the current rainfall alert rule still outperforms these simplified exponential rules. After modifying the current rule to make it comparable to this family of exponential rules, or after accounting for the uncertainties associated with the performance of the current rule, this performance advantage evaporates.

3.5 The double exponential rain index

If different types of storms are best detected with different exponential filters, might not a hybrid rule based on two exponential indices offer superior performance? Initial support for this idea comes from the crossing lines in the complementary cumulative plot (Figure 3-10), and from plotting one exponential index against another at hours with failures (Figure 3-14).

However, such evidence is not as strong as it first seems. Many exponential rules perform admirably in providing warning for slope failures preceded by ample rain.

These, however, are the "easy" failures to catch. Exponential rules have more difficulty providing warnings for those failures preceded by little rain. No matter the half-life, such failures occur when the rainfall index is small. For example, in Figure 3-10, we see the hardest failure to catch (the top line) is always the hardest failure to catch, regardless of h. In Figure 3-14, we again see the two failures with the smallest $h = 1$ indices at failure also have the smallest $h = 20$ indices at failure. Failures preceded by light rain are difficult to catch simply because they are preceded by light rain, and a double index rule cannot solve this problem.

We have other concerns about combining two exponential indices into a single hybrid rule. One must choose two half-life values, h_1 and h_2, which specify the short-term and long-term components of the hybrid rainfall index. In addition, one must choose a function to map these two values into a one-dimensional index. Typically, this function will involve one or more parameters that will also require tuning. We can specify an exponential rule by specifying a single parameter, h; in contrast, choosing a double exponential rule requires specifying two values of h, one function, and possibly several parameters. The family of double exponential rules offers great flexibility, and we must exercise caution when approaching a small dataset with so many free parameters.

As an example, though, we shall consider two types of double exponential rainfall index rules. We do not intend these examples to be comprehensive, and we do not claim these two types are necessarily representative of the large class of possible double exponential rules. We present these two examples more as demonstration of the double index strategy than as proof of its efficacy.

Let us use "D_1" to denote our first type of double exponential rules. After examining the complementary cumulative plots, we choose to base the D_1 rules on the $h_1 = 1$ hour and the $h_2 = 20$ hour exponential indices. To select a function to map these two indices into a single value, we draw insight from the current rainfall alert system. The current rule uses irregular rectangles to partition the HRAIN vs. CRAIN plane into control regions (Figure 3-2). Using the $h = 1$ hour index as a proxy for HRAIN and the $h = 20$ hour index as a proxy for CRAIN, we opt to use convex

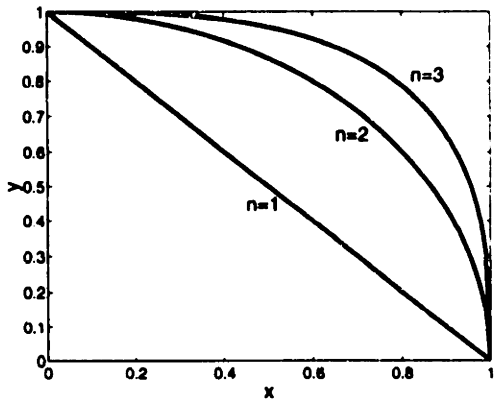

Figure 3-15: Ellipsoid-like control regions for double-index rule. This figure plots $(x^n + y^n)^{(1/n)} = 1$ over the positive quadrant for $n = 1, 2, 3$.

"ellipsoid-like" regions around the origin for the NO ALERT state. We use $R_{D_1}(t, i)$ to denote these hybrid D_1 indices at hour t at gauge i, and derive these indices from the $h = 1$ and $h = 20$ hour exponential indices using the following function.

$$R_{D_1}(t, i, c_1, c_2, c_3) = ([c_1 R(t, i, 1)]^{c_3} + [c_2 R(t, i, 20)]^{c_3})^{\frac{1}{c_3}}$$

This hybrid index involves three positive parameters: c_1 and c_2 scale the two exponential indices to adjust their relative magnitude, and $c_3 \geq 1$ controls the shape of the ellipsoid-like region. As c_3 increases, the control region moves from a triangular to a rectangular shape (Figure 3-15).

We observe a three-fold difference in the average magnitude of these indices at times of failures, $R^*(t, 1)$ and $R^*(t, 20)$, from Figure 3-14. We set the ratio of the two scaling parameters, c_1/c_2, to three, thus scaling both indexes to comparable size. As we round indices to the nearest integer when storing our complementary cumulative histograms, we opt to set $c_1 = 6$ and $c_2 = 2$ to produce $R_{D_1}(t, i)$ values which are neither too small nor too large. We thus obtain the following form for the D_1 double exponential rainfall indices.

$$R_{D_1}(t, i, c_3) = ([6R(t, i, 1)]^{c_3} + [2R(t, i, 20)]^{c_3})^{\frac{1}{c_3}}$$

We evaluate the retrospective performance of the D_1 rules using the national precipitation data and five values for the shape parameter:

$$c_3 \in \{1.0, 1.5, 2.0, 2.5, 3.0\}.$$

For this evaluation, we follow a procedure similar to the one used to evaluate the single exponential rules.

1. For each value of c_3, calculate the D_1 double exponential rainfall index at each hour at each gauge, $R_{D_1}(t, i, c_3)$.

2. Determine the complementary cumulative distributions, $G(r_{D_1}, i, c_3)$.

3. Determine the set of "best" alert thresholds, $(R_{D_1})^*_{(j)}(c_3)$.

4. Determine the false negative performance of these thresholds, $FN(r^*, c_3)$.

5. Determine the service disruption hours of these thresholds, $SDH(r^*, c_3)$.

After these calculations, we employ a complementary cumulative plot to summarize where each failure falls in the distribution of the D_1 indices as a function of the shape parameter (Figure 3-16). We observe that most of these lines are horizontal, and few of these lines cross. This indicates that varying c_3 may change the absolute size of the double index R_{D1}, but that varying c_3 does not change the relative location of each failure within the distribution of this index. In simpler terms, the shape parameter c_3 has little effect on the safety or delay performance. This occurs because—regardless of shape parameter—the most difficult slope failures to catch remain the most difficult slope failures to catch.

This experience with the D_1 rules suggests that the shape parameter contributes little value to these indices. For our second example of a class of double exponential indices, our so-called "D_2" rules, we set $c_3 = 1$ and instead vary the c_1/c_2 ratio. To this end, we set $c_2 = 1$ and vary c_1. This allows us to explore the relative importance of the short-term and the long-term components of this hybrid index. We use $R_{D_2}(t, i, c_1)$

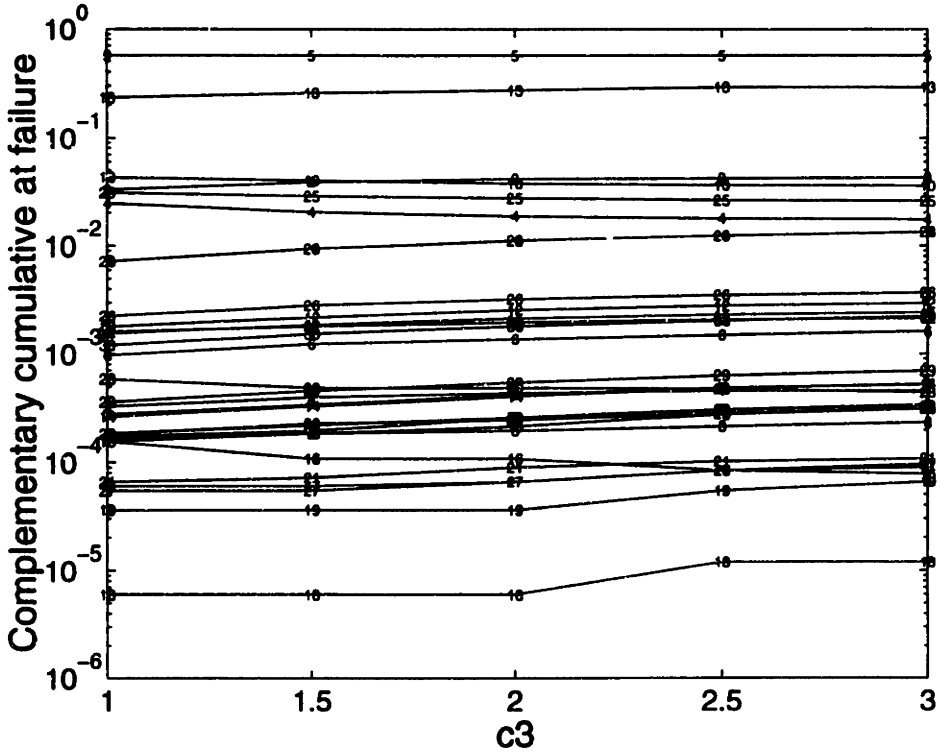

Figure 3-16: Complementary cumulative plot for the D_1 double exponential rules.

to denote these D_2 hybrid indices at hour t at gauge i.

$$R_{D_2}(t,i,c_1) = c_1 R(t,i,1) + R(t,i,20)$$

We evaluate the retrospective performance of double index rule D_2 using the national precipitation data and six values for the scaling parameter:

$$c_1 \in \{1,2,3,4,5,6\}.$$

For comparison, we plot the safety vs. delay performance of these rules along with the safety vs. delay performance of the single exponential rules (Figure 3-17). We note that these D_2 rules offer slight performance improvements over the single exponential rules. We also note the most effective of these rules employ the largest values of c_1. That is, the most effective rules place the greatest weight on the short-term $h = 1$ hour component of the index. This result has an analog in the design of the current rule, which usually issues alerts upon reaching a large value of HRAIN, but can also issue alerts when CRAIN becomes very large. As before, we note that the current rainfall alert rule outperforms even the best of double exponential index rules; we again attribute this difference to the simplifications we have imposed on our alternative rules.

3.6 Safety vs. disruption tradeoff

Each rainfall alert rule involves some tradeoff between expected safety and expected service disruption. In this section, we introduce a model that allows this tradeoff to be quantified. For this discussion, let us suppose a specific type alert rule has been selected, and our concern is choosing an appropriate alert threshold for that rule. For this example, we assume that the selected rule uses a single exponential index with half-life $h = 1$ hour, and the problem at hand is estimating the effect of the threshold r^* at which STOP alerts are issued. If the railroad chooses a small value for r^*, the alert system will be quick to call alerts, leading to a low false negative rate and a high

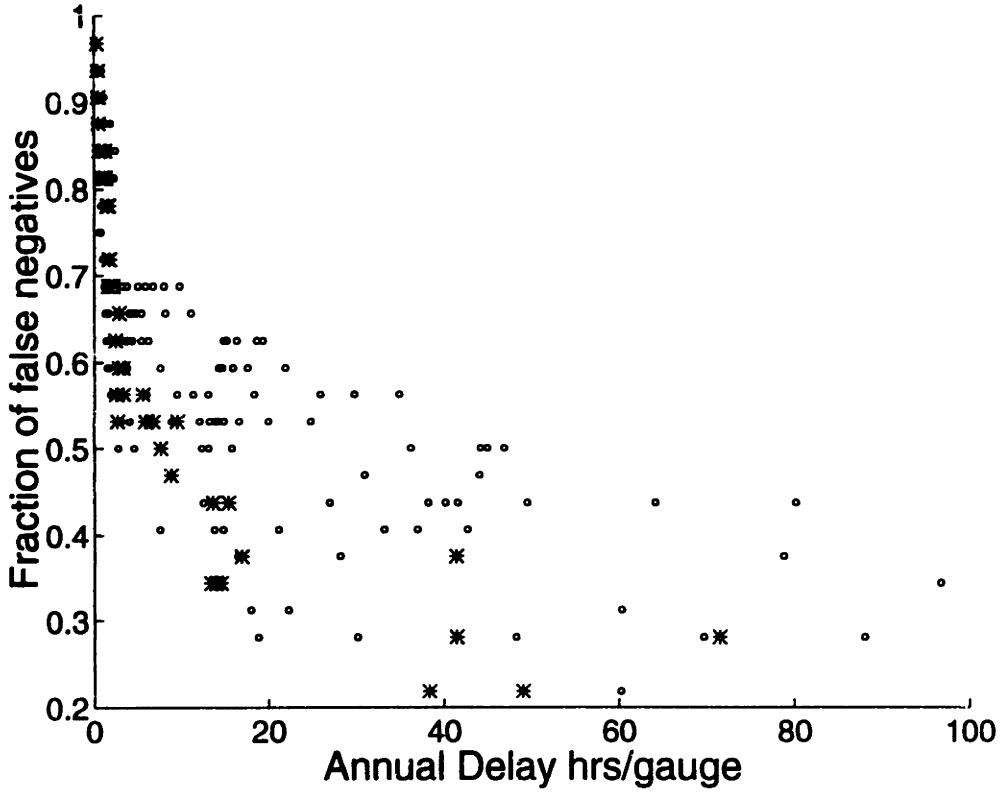

Figure 3-17: Performance of the D_2 double exponential rules, represented by stars, compared with the performance of the single exponential rules, denoted by small circles. Precipitation data from the national dataset.

service disruption rate. If the railroad chooses a large value for r^*, the alert system will be slow to call alerts, leading to a low service disruption rate and a high false negative rate.

To quantify this tradeoff between expected safety and expected disruption, we estimate the probability of a slope failure as a function of the rainfall index. Then, using this approximate relationship, we can determine the approximate expected safety and expected service disruption associated with each level of the rainfall alert threshold. We begin by modeling the probability of a slope failure the $h = 1$ hour half-life exponential rainfall index. We discretize this index by rounding it to the nearest integer. To model the occurrence of slope failures, we assume that when gauge i has rainfall index value r, slope failures occur within the operating region of gauge i according to a Poisson process with rate $\lambda(r)$ slope failures/hour. For this example, let us assume that $\lambda(r)$ does not depend on i, implying the slopes along the entire rail line respond similarly to heavy rainfall. We need a reasonable method to estimate the function $\lambda(r)$ for this exponential rainfall index. Let $n(r)$ denote the total number of gauge-hours in the dataset with a rainfall index value equal to r, and let $s(r)$ denote the total number of slope failures that occurred when the rainfall index had value r. The MLE estimate for $\lambda(r)$ is simply $s(r)/n(r)$; unfortunately, this point estimate yields zero for all index values which have yet to suffer a failure. We instead use a more sophisticated approach to estimate $\lambda(r)$, beginning with a qualitative specification of the properties we desire for this function.

- As $\lambda(r)$ is a rate of failures per unit time, we want $\lambda(r)$ to be positive for all r.
- As we believe the risk of slope failure increases with increasing rain, we want $\lambda(r)$ to be monotonically increasing.
- As slope failures may occur even when the rainfall index is equal to zero, we want $\lambda(0) > 0$. (As we have rounded the index to the nearest integer, $\lambda(0)$ denotes the hourly rate of slope failures when $R(t, i, 1) < 0.5$ mm.)
- As we believe that additional rain does little to increase the slope failure risk

once the slope is entirely saturated, we want $\lambda(r)$ to approaching a limiting asymptotic value as r gets very large.

Of the many functions fit these requirements, we have selected one.

$$\ln \lambda(r) = \alpha_1 + (\alpha_2 - \alpha_1)(1 - \exp(-\alpha_3 r)) \tag{3.4}$$

With this model, when $r = 0$, the hourly rate of slope failure is $\exp(\alpha_1)$, and when r is very large, the hourly rate of slopes failure approaches $\exp(\alpha_2)$. The parameter α_3 controls how quickly this rate increases between these extremes.

We use a numerical maximum likelihood approach to determine α_1, α_2, and α_3. In the following sum, we adopt the convention $0^0 = 0$.

$$\begin{aligned} \ln L(\alpha_1, \alpha_2, \alpha_3) &= \sum_{r=0}^{\infty} [s(r) \ln(n(r)\lambda(r, \alpha_1, \alpha_2, \alpha_3)) - n(r)\lambda(r, \alpha_1, \alpha_2, \alpha_3)] + c \\ (\hat{\alpha_1}, \hat{\alpha_2}, \hat{\alpha_3}) &= \arg \max_{\alpha_1 < \alpha_2 < 0, \alpha_3 > 0} \ln L(\alpha_1, \alpha_2, \alpha_3) \end{aligned}$$

For this $h = 1$ hour exponential index along the Suigun using the national precipitation dataset, we obtain MLE estimates of $\hat{\alpha_1} = -12.9$, $\hat{\alpha_2} = -0.8$, $\hat{\alpha_3} = 0.04$. Using these parameters, we plot $\lambda(r)$ and our point estimates $s(r)/n(r)$ on the same log plot (Figure 3-18). The fitted $\lambda(r)$ appears nearly always below these point estimates, suggesting a poor model. However, this graph deceives the eye, for most of the values of the rainfall index have yet to experience slope failures. These values have $s(r) = 0$ and thus $s(r)/n(r) = 0$; that is, the logarithmic scale hides the vast majority of point estimates lurking down below at zero. Presenting this same graph on a linear scale ameliorates this difficulty (Figure 3-19). The disadvantage of this second presentation is that one's eye is drawn to the right of the graph, where the function $\lambda(r)$ does not provide the most impressive fit of these data. This too is deceptive, for the most important region of the graph is on the left, when r is small; however, this region is obscured on the linear scale. For this discussion, we shall accept this functional form and these MLE parameter estimates for hourly rate of slope failure, $\lambda(r)$.

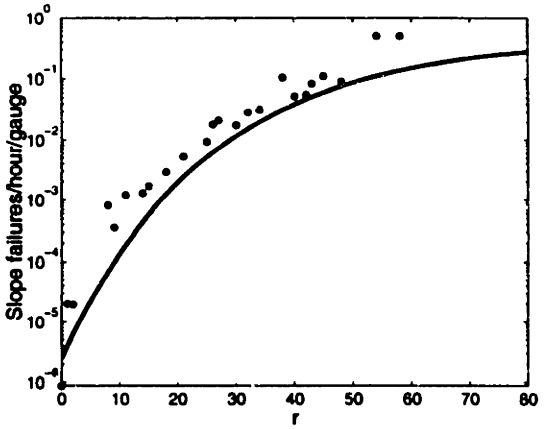

Figure 3-18: Best MLE fit for $\lambda(r)$, indicated by the black line, and the $s(r)/n(r)$ point estimates, indicated by black dots, based on the $h = 1$ hour exponential index.

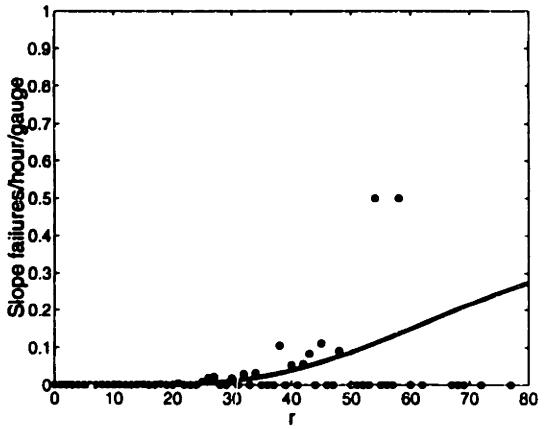

Figure 3-19: Best MLE fit for $\lambda(r)$, indicated by the black line, and the $s(r)/n(r)$ point estimates, indicated by black dots, based on the $h = 1$ hour exponential index.

With this model for $\lambda(r)$ in hand, we may determine the safety and delay performance of this $h = 1$ hour exponential rainfall rule as a function of its alert threshold, r^*. Let $f(r)$ denote the distribution of this discretized rainfall index averaged across all the gauges.

$$\begin{aligned} f(r) &= \Pr(\text{a random hour } t \text{ at a random gauge } i \text{ has } \lfloor R(t,i,1) + 0.5 \rfloor = r) \\ &\approx (\text{fraction of all hours in the dataset with } \lfloor R(t,i,1) + 0.5 \rfloor = r), \\ &\quad \text{for } r = 0, 1, 2, \ldots \end{aligned}$$

If we set the threshold at r^*, the rule will issue a STOP alert for every hour with an index value at or above this threshold. The probability that a randomly chosen hour has an index value at or above r^* is $G(r^*)$, which we now write as $\sum_{r=r^*}^{\infty} f(r)$. Further, this rule will successfully provide STOP alerts for all slope failures which occur at index value r^* or above. This rule will not catch slope failures which occur when the rainfall index is below this threshold; such events will be false negatives. From our Poisson model for slope failures, the chance of one or more slope failures during an hour with rainfall index r is $(1 - \exp(-\lambda(r))$. We can thus determine the probability of a false negative and the probability of a service disruption hour for a randomly chosen hour as a function of threshold r^*.

$$\begin{aligned} PFN(r^*) &= \Pr(\text{hour has false negative, threshold}=r^*) \\ &= \sum_{r=0}^{r^*-1} (1 - \exp(-\lambda(r)) f(r) \\ PSDH(r^*) &= \Pr(\text{hour has service disruption, threshold}=r^*) \\ &= \sum_{r=r^*}^{\infty} f(r) \end{aligned}$$

With 6,500 operating hours in the year, the expected annual number of service disruption hours per gauge is $(6500)(PSDH(r^*))$ hours. Similarly, the expected number of false negatives per gauge is $(6500)(PFN(r^*))$, and the expected number of false negatives across the system is $(17)(6500)(PFN(r^*))$ per year (Figure 3-20).

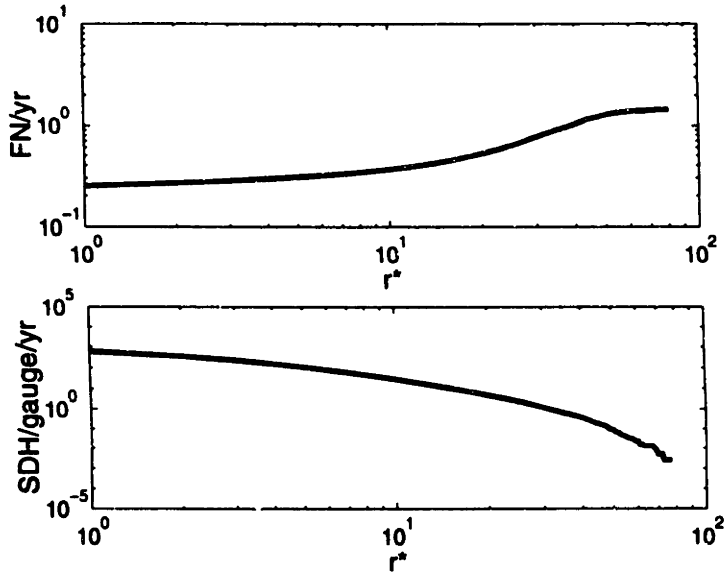

Figure 3-20: Annual expected number of false negatives, and annual expected number of service disruption hours per gauge, as a function of alert threshold r^*, using the exponential rainfall index with half-life $h = 1$ hour. Slope failure rate model from Equation 3.4; rainfall frequency data from national precipitation dataset.

3.7 Slope failure risk in context

We offer some back-of-the-envelope calculations intended to place slope failure risk within the larger context of total operating risk. Though these estimates are coarse, they serve to establish the magnitude of this hazard. We estimate the conditional probability of a derailment, given a slope failure; we also estimate the mean passenger injuries fatalities, given a derailment.

Probability of derailment, given a slope failure

Let random variable N_f denote the number of serious slope failures that have occurred across the JR East conventional network since privatization, and let N_d denote the number of derailments caused by these slope failures since privatization. We assume that each serious slope failure leads to a derailment with probability p, and next calculate a rough estimate of this probability.

JR East has suffered thirty-two derailment accidents since privatization; one of

these was caused by a slope failure; thus $N_d = 1$. [SRL95b] Lacking direct data on N_f, we cautiously estimate this number by extrapolating from the experience of the Suigun line. The Suigun line is 150 km long and has experienced thirty-two serious slope failures over the last twenty years. Assuming the rate of such failures is uniform in space and in time, these figures suggest a serious slope failure rate of 0.011 per kilometer per year. The total length of all lines in the conventional JR East network is 6,900 km; 2,600 of these kilometers are described as mountainous terrain. [SRL95k] If we make the conservative assumption that only mountainous track is vulnerable to slope failure, and if we make the additional assumption that the rate of serious failures along the Suigun line is typical of the rate along all mountainous lines, we obtain an estimate of $N_f = (2600 \text{ km})(19 \text{ years})(0.011 \text{ fails/km/year}) = 530$ serious slope failures across the JR East conventional network since privatization. With this value, the MLE estimate that a serious slope failure leads to a derailment under present policies is $\hat{p} = N_d/N_f = 1/530 = 0.002$.

Injury and death, given a derailment

Given that a train derails at a slope failure, what is the expected number of injuries and fatalities? What is the probability a passenger riding on such train is injured or killed? A derailment at a slope failure has occurred once since privatization, on September 16, 1988, on the Hachinohe line. There were 22 passengers on this train when it derailed; no passengers or crew were injured or killed in this accident. [SRL95b] Before privatization, over the region that would become JR East upon privatization, JNR experienced seven derailments caused by slope failures. None of these seven accidents caused fatalities; two of these accidents each injured one crew member; and one of these accidents injured one crew member and two passengers. [SRL95a] Over this combined sample of eight slope-failure induced derailments, we observe an average of 0.25 passenger injuries and 0 passenger fatalities per accident.

There is strong statistical evidence that the consequences of a derailment depend on the reason for that derailment, with derailments following collision accidents and derailments following level crossing accidents having the most severe consequences.

[RK95] Despite evidence to the contrary, one might argue that derailments at slope failures are just as severe as all other types of derailments, thus employ the entire derailment record to gain insight into the expected consequences of derailments at slope failures. Most of the 32 JR East derailments to date caused neither injuries nor fatalities. Six accidents injured passengers, and one accident killed one passenger. The average number of passenger injuries per derailment is 6.6; the average number of passenger fatalities per derailment is 0.032. These means are not entirely representative as these distributions are right skewed. Indeed, the median number of passenger injuries and the median number of passenger fatalities are both zero per derailment.

The JR East derailment database provides data on the number of passengers on riding trains involved in derailments. With a total 7400 passengers riding the trains involved in the thirty-two JR East derailments, and with 212 of these passengers sustaining injuries, we estimate the probability that a randomly-chosen passenger on a train that derails is injured is less than 0.03. Using a similar argument, the probability that a randomly-chosen passenger on a train that derails is killed is less than 0.00014. For a worst-case scenario, one might imagine a train derailing at a slope failure atop a steep cliff, with the train subsequently tumbling down the hillside. Such a disaster might be comparable to a train collision. If so, we note the worst JR East collision to date caused 116 passenger injuries and one passenger death. While there is always the possibility that heavy rainfall could lead to a train derailment responsible for many passenger fatalities, it is our belief that such events are exceedingly rare. and that overall rainfall-induced slope failures pose only a minor threat to JR East.

3.8 Conclusions

This chapter has presented a series of analyses addressing slope failure risk to the JR East conventional train network. The Suigun line has been used as case-study throughout. While the specific conclusions of this chapter apply to this particular rural line, the methodologies apply to any rail line in the network. Indeed, the tradeoff

between safety and disruption raised in this chapter is a more pressing problem on lines with greater economic importance. For example, JR East operates the Narita Express, the fastest available ground transportation from Tokyo to the Narita airport. This rail line runs through mountainous areas vulnerable to slope failures. On such a flagship line, the cost of disruption and the value of safety is likely many times higher than on the quiet Suigun line. Even if the ratio of safety to delay costs are similar for both high and low volume lines, the greater economic importance of larger lines calls for a finer calibration of this tradeoff.

We conclude this chapter by summarizing its major findings.

- The current rainfall alert system for the Suigun line performed well over the 1991 to 1994 time period, at least in comparison to the suite of alternative rules proposed in this chapter. One contributing factor to this strong performance may be the use of overlapping gauge control regions. Overlapping gauge control regions may offer valuable redundancy and help reduce the chance a failure is "missed" due to a local rainfall fluctuation. Alternatively, some or all of this fine performance could simply be the result of a run of good luck.

- When evaluating the performance of rainfall alert rules, one should attempt to determine the uncertainties in raw statistics given the small number of slope failures. Similarly, one should acknowledge and attempt to quantify the artificial performance inflation arising from retrospective analysis, in which one calibrates and tests the model on the same dataset.

- For any rainfall index, a complementary cumulative plot offers a comprehensive presentation of the full safety vs. delay spectrum.

- Our analyses suggest that single exponential rules using short one or two hour half-lives offer superior performance to rules using longer half-lives.

- Double exponential rules may offer slightly better performance than single exponential rules. The best of these double index rules place most of their emphasis on short-term rainfall.

- As slope failures preceded by ample rain are easily "caught" by any sensible rainfall index rule, it is those few failures preceded by light rain that most determine a rule's performance.

- Selecting a rainfall index and a rainfall alert threshold implies a relative cost for safety and delay. By modeling the hourly rate of slope failures as a function of a rainfall index, one may determine the optimal alert threshold for a given cost, or determine the implicit cost associated with each threshold. Our preliminary analysis in Section 3.6 does not account for SLOW alerts; this intermediate alert level between NONE and STOP could offer more efficiency, in the sense of providing the same level of safety with less delay or increased safety with less delay.

- Extrapolating from the Suigun line to the entire network, the probability of a slope failure causing a derailment is less than one in five hundred. Moreover, even the most serious derailments have typically small consequences, as over 95% of these accidents cause neither passenger injuries nor fatalities.

Chapter 4

Collision risk to trains

4.1 Motivation and scope

Train collisions are a railroad's nightmare. Collisions kill passengers and employees, destroy rolling stock and track, and tarnish a railroad's reputation. Unlike accidents caused by natural hazards, a railroad must take full blame for train collision accidents. There is simply no reason for two trains ever to crash head-on; there is simply no reason for one train to crash into another from behind.

Nonetheless, train collisions do occur, though rarely. That collisions are infrequent testifies to the success of the numerous safety systems used by modern railroads. Railroads worldwide have evolved safety systems in response to accidents. These safety systems pervade all aspects of a railroad. The track layout, the schedule, the signaling system, the dispatching system, and the training of drivers are all carefully planned to keep trains safely apart.

Such precautions are made necessary by the constraints of rail transport. Unlike cars or airplanes, trains move on a constrained right-of-way, with no ability to swerve to avoid a collision. Also, trains typically have considerable mass, and thus require long distances to stop. For both these reasons, railroads take particular care to keep trains safely apart.

This chapter investigates the collision risk faced by trains in the JR East system. We restrict our attention to the conventional train system, and do not consider col-

lisions between Shinkansen trains. Shinkansen trains run on double track with long headways between successive trains; the Shinkansen track is reserved for the exclusive use of Shinkansen trains; and the Shinkansen signaling and dispatching systems are near state-of-the-art. Furthermore, to date the Shinkansen lines all have perfect safety records: in thirty years of Shinkansen operation, no Shinkansen trains have collided. For these reasons, we limit our analysis to the conventional system.

Train collisions are extremely rare events, and assessing the likelihood of rare events is problematic. In this chapter, we apply two different approaches to this problem. First, in Section 4.3, we investigate collision risk empirically, applying statistical methods to historical data. Second, in Section 4.4, we propose a conceptual model to study collision risk more abstractly. We conclude this chapter discussing the implications of both approaches.

4.2 Collision summaries

JR East came into existence in April 1987, with the breakup and privatization of the Japanese National Railway. Since the privatization, the JR East conventional rail system has provided 2.4×10^9 train-km and 9.7×10^{11} passenger-km of service. [Nas95] Since privatization, the railroad has suffered three collisions.

JR East defines the term "accident" broadly, as the railroad counts a delayed train as an accident. Most JR East "accidents" are harmless cases of delay, and thus collisions comprise a negligible fraction of all JR East accidents. However, when we consider accident consequences, such as injuries and deaths, collision accidents comprise a more substantial fraction. Table 4.1 indicates the contribution of collisions to total accident consequences between 1987 and 1993. As there have been no additional collisions since 1993, the percentages in this table overestimate the relative contribution of collisions to total accident consequences between 1991 and 1996. Assuming constant annual growth in total consequences, we might multiply these six-year totals by 1.5 to obtain proxy nine-year totals; this adjustment would reduce the listed percentages by two-thirds.) We note that the majority of fatal accidents occur at

Number of accidents	3/57,440	0%
Suspended trains	193/47,811	0%
Total fatalities	1/605	0%
Total injuries	120/1,308	9%
Passenger fatalities	1/13	8%
Passenger injuries	116/367	32%

Table 4.1: Collision accidents as a percentage of total accident consequences, 1987–1993. Data from [Nas95].

level crossings, and most of these deaths are to third-parties, not passengers or crew. Accidents involving death or injury to JR East passengers or crew are very rare; two of the three collision accidents were of this type.

JR East has experienced three collisions to date. It is instructive to review these three accidents briefly. [SRL95e]

"Lost cars" collision. October 19, 1988.

This accident involved JR East freight trains and took place at 2 a.m. Train 3881, while passing oncoming train 4660 on an adjacent track, noticed flames coming from train 4660. Notified by radio, train 4660 stopped to investigate. The fire had gone out, but on-foot inspection revealed an additional problem: the last twelve of the train's thirty-two cars were missing. By radio, train 3881 confirmed that train 4660 was a full thirty-two cars long when they had passed, so the twelve lost cars could not be far behind. A search was begun, but the lost cars were not found in the darkness. The dispatchers warned all trains in the vicinity about the lost cars. Shortly thereafter, train 3781 approached the scene, traveling on an adjacent track. Train 3781 had been informed of the dangerous situation. Nonetheless, train 3781 struck two of the twelve lost cars, which had derailed and sprawled on its track. Fortunately, there were no crew injuries or fatalities.

"Station crash" collision. December 5, 1988.

This serious accident involved JR East passenger trains and took place during the morning rush hour on the busy Chuo line in downtown Tokyo. Train 801C

was paused in Higashinakano station, taking on and letting off passengers. Train 835B was approaching the same platform on the same track, behind schedule and speeding to catch up. Due to the presence of train 801C ahead, the ATS signaling system rang a warning bell in the cockpit of train 835B to alert the driver that the next signal was red. For some reason, the driver of train 835B pushed the ATS manual override button, silencing the alarm and disengaging the automatic train braking system. This override allowed train 835B to crash into train 801C. The crash killed the driver and one passenger on train 835B. In addition, 116 other passengers on trains 835B and 801C were injured. Because of this accident, JR East created the Safety Research Laboratory to monitor and improve system safety.

"Short train" collision. June 28, 1992.

This accident involved an out-of-service JR East passenger train and a JR East freight train and took place just before the morning rush hour on the Tokaido line before the Kinomiya station. Train 6726 was waiting to begin service for the day. The driver had stopped train 6726 so as to line up the last car of his train with the last car of the train on the adjacent track, to permit the conductors in the rear of both trains to chat. For some reason, train 6726 was unusually short that day, placing the front of train 6726 much further back from the usual stopping point. One of many signals in front of the parked trains turned from red to green. The driver of train 6726 mistook this green light for his own and started his train towards the station. Because his train was unusually short, and because he had stopped his train unusually far back, the driver achieved a higher speed than usual when he reached his red signal. The signaling system correctly instructed the train to brake to emergency stop, but train 6726 had sufficient speed (and thus momentum) to overrun the signal by 45 meters while braking. Within this distance, the train encountered a switch which routed the front of the train onto the main line. The driver radioed the dispatch control center of his plight. Freight train 106 was approaching the site and began

braking in an attempt to avoid colliding with train 6726. However, the freight train lacked sufficient time and distance to stop, and crashed head-on into train 6726. There were no passengers on either train. There were no crew fatalities. The driver of train 6726 sustained light injuries.

We offer three comments on these accidents. First, we observe that each of these accidents appears unique. Though it is likely that JR East will suffer additional collisions at some point in the future, it is plausible that none of these three specific scenarios will occur again. We suggest that this catalog of collisions contains examples of what is possible, not necessarily what is prototypical.

Second, we believe that these three accidents do not suggest the presence of glaring flaws in the safety systems. One might argue that train-detection circuitry that "loses" several cars represents a glaring flaw, but this circuitry usually operates perfectly. One might argue that a driver's ability to override an ATS emergency brake warning represents a glaring flaw, but such overrides are typically benign; in fact, such overrides are both necessary and common in the Tokyo area due to train congestion. One might argue that a signaling system which cannot protect over-speed trains represents a glaring flaw, but this system usually operates perfectly, and protecting over-speed trains would require larger block sizes and thus reduce the effective capacity of each line. After a collision, finding some fault or another to blame for "causing" the accident is not difficult. The meaningful issue is whether this fault is likely to cause additional collisions.

Third, we observe that the chain of events leading up to each accident is typically complex. Certainly in the case of the "lost cars" and "short train" accidents, it appears the interaction of several seemingly-innocuous events led to a collision. While the "station crash" accident may have had a single primary cause—driver error—system delay may have played some secondary role.

Qualitatively, the three JR East collisions to date appear unique, accidental, and complex. We next examine the trend of JR East collision risk quantitatively.

4.3 Statistical approach

In this section, we estimate the average JR East collision rate since privatization. As this rate is based on only three accidents, we cannot establish a time trend for this rate directly. Rather, we argue that signal overrun rates may serve as a proxy measure of overall probability of a collision. Arguing from a significant downward trend in the signal overrun rate, we suggest the collision rate may have also fallen since privatization. We estimate the number of collisions "averted" due to this "invisible" safety improvement. We also estimate the collision rate for the various signaling systems.

4.3.1 Collision data

Average collision rate

We now estimate the average rate of collisions in the JR East system over the nine years since privatization, from April, 1987 to April, 1996. We adopt a simple model: every time a train travels one kilometer, we assume there is a very small probability p it collides with a second train. We assume each of these Bernoulli trials are independent. We use v to represent the total train-km of service since privatization. We use a to represent the total number of collision accidents since privatization. Assuming v is very large, and p is very small, a is well-modeled by a Poisson random variable with mean λv. We use λ to denote the average rate of collisions per train-kilometer.

The obvious estimator of λ is $\hat{\lambda} = a/v$, the observed ratio of collisions to train kilometers. This estimator has pleasing properties: it is unbiased, and it has the lowest variance among all unbiased estimators. [Ing94b]

$$\hat{\lambda} = \frac{a}{v} = \frac{3}{2.4 \times 10^9} = 1.3 \times 10^{-9} \text{collisions/train-km}$$

With one collision every 770 million train-kilometers, a random passenger could expect to cross the island of Honshu over one half million times between collisions.

We note that over the same geographic region, between 1976 and April, 1987.

Japan National Railways provided 2.9×10^9 train-km of service and suffered thirteen collisions, for a collision rate of 4.5×10^{-9} per train-kilometer. [SRL95a] To determine if this difference in observed collision rates is statistically significant, let us test the following null hypothesis: the true JNR collision rate is equal to the true JR East collision rate. Given this null hypothesis, the number of the sixteen total collisions occuring in the JR East period should follow a Binomial$(16, p)$ distribution, where $p = (2.4 \times 10^9)/(2.4 \times 10^9 + 2.9 \times 10^9) = 0.45$. The probability that a Binomial$(16, 0.45)$ random variable is three or smaller is 0.006. We thus reject the null hypothesis and assert that over the JR East region, JR East collision rate is significantly smaller than the JNR collision rate.

Time trend of collision rate

The $\hat{\lambda}$ estimator does not address whether the collision rate has changed since privatization. For example, if the time since the last collision appears unusually long, this large "gap" might suggest that the risk has dropped since the last accident, perhaps reflecting successful improvements to various safety systems. (The annual volume of conventional train-kilometers has remained essentially constant since privatization, so we would credit an exceptionally long collision-free period to a reduction in the collision rate, not in the traffic volume.)

For a null hypothesis, let us assume the collision rate has remained stable from privatization to the present. Given this assumption, the instants at which collisions occurred should be uniformly distributed across this period. One test of the stable-rate assumption is whether all the collisions occurred "too early" in the period. We define the first day of JR East operations as our time origin, $T_0 = 0$; we use T_i denote the number of days from this date to the ith collision accident; and we let T_{now} denote the number of days from privatization until the present (Table 4.2). The probability that all three accidents occurred before time T, given our stable-rate null hypothesis, is $p = (T/T_{now})^3$. If this probability is sufficiently small, we would be forced to reject our null hypothesis. We observe that all three JR East collisions occurred by day 1,915; that is, within the first 58% of the period. The probability of this event, given

T_0	=	0 days	April 1, 1987	privatization
T_1	=	567 days	October 19, 1988	"lost cars" collision
T_2	=	614 days	December 5, 1988	"station crash" collision
T_3	=	1,915 days	June 6, 1992	"short train" collision
T_{now}	=	3,288 days	April 1, 1996	present

Table 4.2: Time-line of JR East collisions.

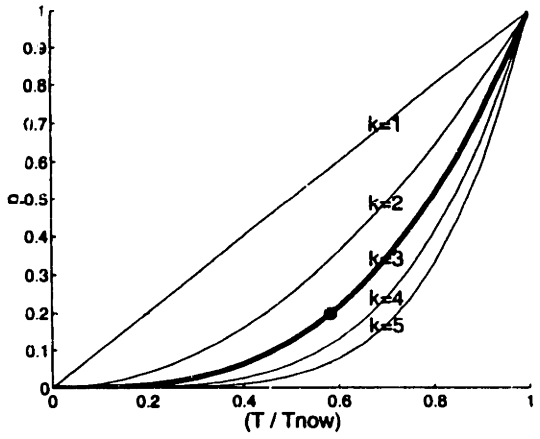

Figure 4-1: Testing the stable-rate hypothesis. This plot indicates the probability that k out of k observations from a uniform $[0, T_{now}]$ distribution fall in interval $[0, T]$. As indicated by the black dot, the probability that the $k = 3$ JR East collisions fell in the first 58% of the interval, given the null hypothesis, is $p = 0.2$.

our constant-rate null hypothesis, is 0.20 (Figure 4-1). We conclude that based on the time since the last collision, the current JR East collision record is consistent with a stable--rather than a falling—collision rate.

We note that if the current collision-free interval continued for the next five years, this would place the three collisions in the first 37% of the period, and lead us to reject the stable-rate hypothesis at the $(0.37)^3 = 5\%$ level.

4.3.2 Signal overrun data

A railroad strives to keep trains safely apart when scheduling and dispatching trains. Each train driver strives to keep his train a safe distance from the next. Sometimes these efforts fail, and moving trains become too close. When this occurs, the signaling system keeps trains from colliding.

At the heart of the signaling system is the concept of a "block." A block is a short interval of track, typically from one-tenth to two kilometers in length. Each rail line is partitioned into non-overlapping blocks. The signaling system is designed to ensure that at most one train occupies a block at a time. "Signals," or traffic lights for trains, use colored lamps to inform train drivers of the status of the next two blocks. For simplicity, we describe the three-color system: green, yellow, and red. On a typical double-track line, a red signal indicates that a train occupies the next block. A yellow signal indicates that the next block is clear, but a train occupies the subsequent block. A green signal indicates the next two blocks are clear (Figure 4-2). The situation is similar for single-track lines (Figure 4-3). All JR East lines have some form of "automatic train stopping" (ATS) or "automatic train control" (ATC); these systems intercede and apply the emergency brakes should a train attempt to run a red light. These systems consist of many interconnected devices. Track circuitry detects trains; inter-locking logic links signals and switches; signals convey information to drivers; and ATS or ATC equipment orders emergency braking when required. These devices work collectively to ensure that at most one train occupies each block at a time.

On rare occasions, something fails, and a train overruns a red signal. JR East calls this event a "signal overrun accident." Usually, a signal overrun accident has no consequences, for the train already occupying the block is typically at the far end, safely distant from the train entering the block. As defined, a signal overrun accident is not a necessary precondition for a collision. For example, in the 10/19/88 collision, the train-detection circuitry failed, and train 3781 passed green lights all the way to its crash.

While collisions are not always preceded by signal overruns, we suggest that signal overruns can serve as a proxy measure for the overall safety of a signaling system. Fewer signal overruns are clearly preferable to more, as fewer overruns suggest a smoothly-operating rail system. An increase in signal overrun accidents suggests a system under increasing stress. Let us assume that a signal overrun leads to a collision with a fixed probability q. Thus, a change in the signal overrun rate leads

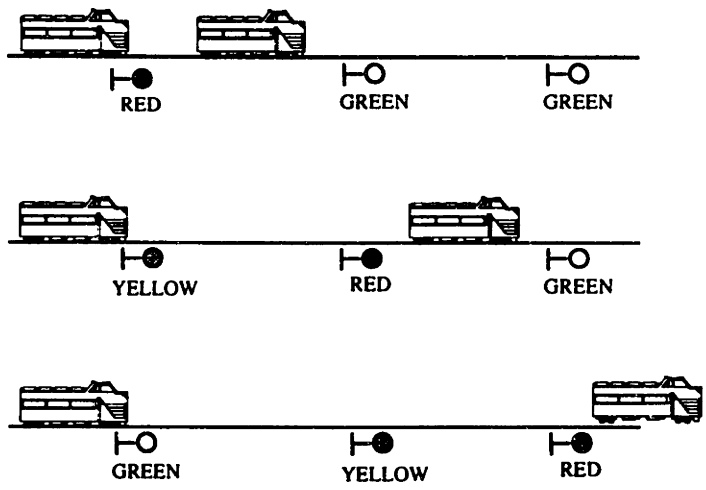

Figure 4-2: The double-track signaling system. Signals mark the end of blocks. In the upper diagram, the left train is approaching a red signal and must stop immediately. In the middle diagram, the left train is approaching a yellow signal and must proceed with caution. In the bottom figure, the left train is approaching a green signal and may proceed normally. Block lengths are not to scale.

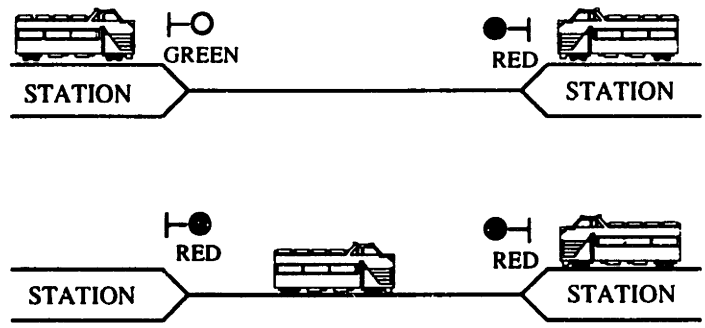

Figure 4-3: The single-track signaling system, used on lower-volume lines. Each inter-station interval is a block, protected on both sides by signals. In the upper diagram, the left train is about to enter the block, and the right train must wait. In the lower diagram, the left train occupies the block, and flanking red signals prevent other trains from entering. Inter-station distance is not to scale.

to a proportional change in the collision rate. Adopting such a model, we propose to use the trend in the signal rate to gain some insight into the trend in the collision rate.

Average signal overrun rate

We analyzed the JR East signal overrun database, which contains data on all signal overrun accidents between April, 1987, and March, 1995. [SRL95c] We use v to denote the total train-km of service, b to denote the total number of signal overruns, and δ to denote the average rate of signal overrun accidents per train-kilometer over this 95-month period. We estimate δ with $\hat{\delta} = b/v$.

$$\hat{\delta} = \frac{b}{v} = \frac{199}{2.05 \times 10^9} = 9.7 \times 10^{-8} \text{ signal overrun accidents/train-km}$$

We note that signal overruns were approximately 65 times more frequent than collisions. Accordingly, we estimate that the probability a signal overrun progresses to a collision is approximately $q = 1/65 = 0.015$.

Time trend of signal overrun rate

We investigated a time trend in the collision data by examining the interval since the last collision. Signal overruns, however, occur far more frequently than collisions (Figure 4-4). This permits us to take a more powerful approach and directly estimate an annual signal overrun rate.

Let v_t denote the total train-km and b_t denote the number of signal overrun accidents in year t, where $t = 0$ represents 1987 and $t = 8$ represents 1995 (Table 4.3). Let δ_t denote the signal overrun rate per train-kilometer in year t. We assume the number of signal overrun accidents in year t follows a Poisson distribution with mean $\delta_t v_t$. We impose the following structure on δ_t:

$$\delta_t = \exp(c_0 + c_1 t).$$

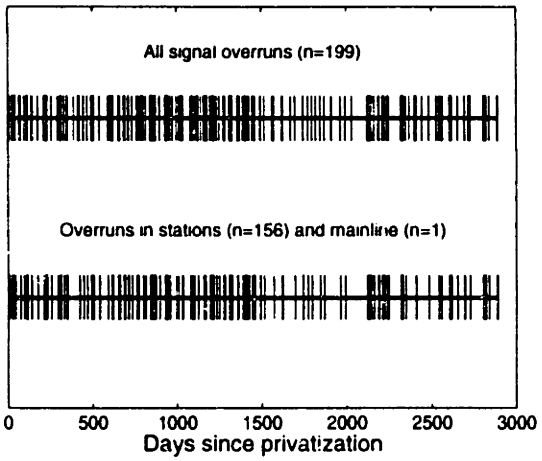

Figure 4-4: Time-line of JR East signal overruns, each represented by a vertical line. The upper plot displays all 199 signal overruns; the lower plot does not include the 42 overruns in depots.

Year	Train-km	Signal overruns
4/1/87–12/31/87	1.63×10^8	26
1988	2.42×10^8	28
1989	2.47×10^8	29
1990	2.70×10^8	32
1991	2.73×10^8	20
1992	2.72×10^8	14
1993	2.67×10^8	27
1994	2.67×10^8	20
1/1/95–2/28/95	1.45×10^7	3

Table 4.3: Annual signal overrun data, compiled from [SRL95c] and [Nas95]. Annual 1987 and 1995 train-km prorated accordingly.

That is, we assume the base signal overrun rate in 1987 is e^{c_0}, and we assume this rate changes by a constant $100(e^{c_1} - 1)\%$ each year. If $c_1 = 0$, the collision rate is constant over time. We estimate c_0 and c_1 by a maximum likelihood approach.

$$\begin{aligned} \{\hat{c_0}, \hat{c_1}\} &= \arg\max_{c_0, c_1} L(c_0, c_1), \quad \text{where} \\ L(c_0, c_1) &= \prod_{t=0}^{8} (\delta_t v_t)^{a_t} \exp[-\delta_t v_t]/a_t! \\ \delta_t &= \exp(c_0 + c_1 t) \end{aligned}$$

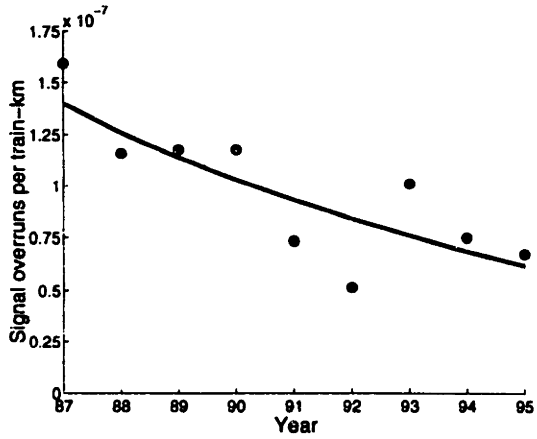

Figure 4-5: Annual signal overrun rate. Dots indicate the empirical annual rate, b_t/v_t. The line indicates the fitted rate $\hat{\delta}_t = \exp(\hat{c}_0 + \hat{c}_1 t)$.

We use the statistical package **STATA** to perform this maximization. [Ham93, p. 139] We find our model fits these data sufficiently well (model $\chi^2_7 = 8.7, df = 7, p = 0.28$), with MLE estimates $\hat{c}_0 = -15.8$ and $\hat{c}_1 = -0.1$. This implies that the base signal overrun rate is 1.4×10^{-7} in 1987, and this rate drops by $100|e^{-0.1} - 1| = 10\%$ each year (Figure 4-5). The 95% confidence interval for $\hat{c}_1$ is $[-0.16, -0.04]$. This interval does not include zero, so we conclude that the annual drop in this rate is significant.

Forty-two of these overruns occurred in depots. Such overruns occur at lower speeds and involve out-of-service trains carrying no passengers. One could argue that accidents in depots should not be counted against the safety of the system. If we omit these depot accidents and solve the maximum likelihood estimation again, we find the model still fits sufficiently well (model $\chi^2_7 = 9.1, p = 0.25$). The MLE estimate $\hat{c}_1 = -0.1$ remains the same, while $\hat{c}_0$ drops to -16.0. This indicates that the non-depot signal overrun rate also falls by $100|e^{-0.1} - 1| = 10\%$ each year. As expected, the 1987 base rate falls 20%; this is because we excluded 20% of the signal overruns.

Collisions "avoided"

Since privatization, JR East has suffered 199 signal overruns and three collisions, or roughly one collision for every 65 signal overruns. Let us assume this 65:1 ratio holds

constant over the period. The signal overrun rate in 1987 was $e^{-15.8} = 1.4 \times 10^{-7}$ overruns per train-km. (This is a conservative estimate, as the point estimate of this quantity, b_0/v_0, is 1.6×10^{-7} overruns per train-km.) Had this rate remained constant over the period, we would have expected approximately

$$\hat{\delta_0} \sum_{t=0}^{8} v_t = (1.4 \times 10^{-7})(2.06 \times 10^9) = 290$$

signal overruns over the period; instead, we observe 199 signal overruns. We argue that safety improvements of some sort 'averted" an expected 91 signal overrun accidents. With a 65:1 ratio of overruns to collisions, we propose that these improvements "averted" an expected 1.4 collisions since privatization, or about 50% of the collisions that actually occurred.

Several factors could have contributed to this safety improvement. In the next section, we consider one such factor: JR East's ongoing program of upgrading the signaling systems across its network.

Signaling systems

JR East employs four different signaling technologies in its network: automatic train control (ATC), and three forms of automatic train stopping (ATS-S/B, ATS-SN, ATS-P). [JICA93] Since privatization, JR East has invested significant resources to upgrade its signaling systems. We next consider if the safety improvements discussed in the previous section may be attributed to improved signaling technology.

Of these four technologies, the ATC system provides the greatest degree of safety at the greatest cost. The ATC system continuously transmits predetermined speeds and signals to each train based on the status of approaching switches and trains. The ATC system automatically brakes trains to avoid signal overruns. The system also stops trains in an earthquake. All Shinkansen lines employ ATC signaling. JR East also uses this technology on its busiest conventional lines: the Yaminote, the Keihin-Tohoku, and the Akabane.

In contrast to ATC, ATS systems do not continuously transmit information to the

trains. Rather, the role of an ATS system is to stop trains before they overrun red signals. ATS-S/B is the oldest, least expensive, and least effective of ATS systems. ATS-S/B rings a bell to warn the driver of an approaching red signal. If the driver does nothing, the system stops the train. The driver may push an override button; this disengages the ATS system and permits the train to continue. This override feature led to the 12/5/88 collision in Higashinakano station.

The ATS-SN system is a slightly improved version of ATS-S/B. This system places an additional transponder just before each signal. If a driver overrides an ATS stop request when approaching a red signal, and if the signal is still red when his train reaches the signal, this additional transponder nullifies the driver's override and orders the train begin emergency braking. If ATS-SN nullifies a driver's override, braking begins near the signal, and the train may travel a short distance into the occupied block while stopping.

The ATS-P system is the most advanced, most effective, and most costly of the ATS systems. The ATS-P system transmits speed vs. distance "patterns" to each train. These patterns indicate the maximum allowable train speed at each point in the block, including deceleration to avoid signal overruns as needed. When a train reaches an ATS-P signal, the signal system examines the state of the next two blocks. If there is a train in either block (that is, if the signal is yellow or red), the ATS-P system provides the train with a pattern designed to stop the train safely before the next train. If there are no trains in the next two blocks (that is, if the signal is green), the ATS-P system assumes there is a train in the third block—a worst-case assumption—and transmits a pattern designed to apply the brakes near the end of the second block to stop the train before the third block. A new pattern is supplied by each signal the train passes. Thus, as a train travels down empty track, each time it passes a green ATS-P signal it receives permission to proceed another block at full speed.

Since privatization, JR East has worked substantially and quickly to upgrade its signaling systems. In recent years, the railroad has allocated one third of its annual safety budget for this purpose; this represents annual expenditures of more than ten

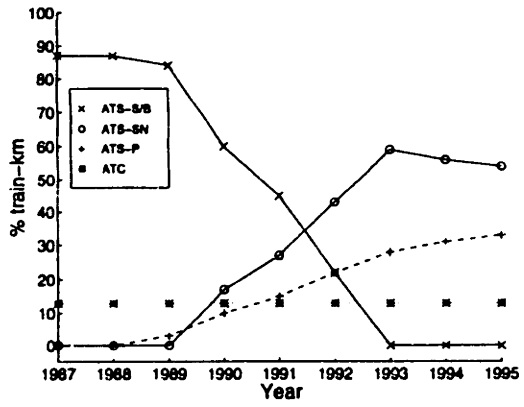

Figure 4-6: Fraction of annual JR East conventional train-km by signaling system.

billion yen (one hundred million dollars), or 0.5% of total annual railroad operating expenses. [Eas93, p. 30] [SRL95j] Over the short term, ATS-S/B lines have been upgraded to ATS-SN (Figure 4-6). Over the longer term, JR East plans to upgrade most ATS lines to ATS-P

There are two motivations for these upgrades: increasing revenue, and increasing safety. [Eas93, p. 12] Better signaling systems allow JR East to run more trains on the same amount of track, increasing capacity and thus increasing revenue. At the same time, the railroad hopes that improved signaling also increases the safety of the rail system. Everything else being equal, it seems better signals would offer improved system safety. However, it does not seem obvious *a priori* that the JR East signal upgrades would necessarily improve system safety: one could envision increases in train volume and density "canceling" out reductions in the per-train risk achieved through better signals. However, in the case of JR East, we do note that the signal overrun rate systemwide appears to have fallen since privatization.

Signal overrun rate by signaling system

There is clear evidence that the signal overrun rate has declined since privatization, possibly accompanied by an "invisible" reduction in the collision rate. Some of these safety gains could have resulted from the signaling system upgrades. Some of these gains could have resulted from other causes, such as better driving and dispatching.

We wish to assess the relative safety of the four signaling systems, and we wish to determine what role signaling upgrades may have played in the reduction of the signal overrun rate.

We will use the subscript j to denote the four signaling systems: $j = 1$ for ATS-S/B, $j = 2$ for ATS-SN, $j = 3$ for ATS-P, and $j = 4$ for ATC. Let f_{tj} denote the fraction of total train-km in year t that occured on lines using signaling system j (Figure 4-6).

$$\sum_{j=1}^{4} f_{tj} = 1, \ \forall t$$

As before, let v_t denote the total train-km in year t. Thus, $v_t f_{tj}$ indicates the total train-km in year t on lines using signaling system j. Let a_{tj} denote the number of signal overrun accidents in year t on lines using system j. Let δ_{tj} denote the annual rate of signal overrun rates in year t for lines using j. We assume that the distribution of a_{tj} is Poisson with mean $v_t f_{tj} \delta_{tj}$. If we had sufficient data, we would estimate the overrun rate by signaling system as a base-rate and an annual trend:

$$\delta_{tj} = \exp(c_{0j} + c_{1j} t), \ \ t = 0 \ldots 8, \ \ j = 1 \ldots 4.$$

This would allow us to determine what role upgrading the signaling systems has played in the reduction of the signal overrun rate. Such a model explicitly parcels out improvements due to signal upgrades. If $\hat{c_1} = 0$, there is no exogenous safety time-trend, and the reduction in the signal overrun rate can be attributed solely to better signals. If $\hat{c_1} \neq 0$, this vector allows us to compare the relative size of the exogenous safety trend to the safety gains of better signals.

Unfortunately, the complete a_{tj} data are not available, as the JR East signal overrun database does not provide signaling system information for accidents before 1992. Thus, we take a simpler approach, and calculate the historical signal overrun rate for each signaling system since 1992. For each system between 1992 and 1995, we consider the portion of the rail network protected by that system, and take the ratio of signal overrun accidents to train traffic, $\sum_t a_{tj} / \sum_t v_t f_{tj}$, over that portion. This approach does not parcel out any exogenous safety-trends, if any exist, but does

	ATS S/B	ATS SN	ATS P	ATC
ATS-S/B	1			
ATS-SN	3	1		
ATS-P	11	4	1	
ATC	11	4	1	1

Table 4.4: The relative signal overrun rates between the four signaling systems. Cell (i,j) indicates the ratio of $\hat{\delta}_j/\hat{\delta}_i$, or how many times safer the row is in comparison to the column.

allow a relative comparison of the four systems (Table 4.4).

$$\text{ATS-S/B:}\quad \hat{\delta}_1 = \frac{\sum_t a_{t1}}{\sum_t v_t f_{t1}} = \frac{12}{6.0 \times 10^7} = 2.0 \times 10^{-7} \text{ overruns/train-km}$$

$$\text{ATS-SN:}\quad \hat{\delta}_2 = \frac{\sum_t a_{t2}}{\sum_t v_t f_{t2}} = \frac{33}{4.5 \times 10^8} = 7.4 \times 10^{-8} \text{ overruns/train-km}$$

$$\text{ATS-P:}\quad \hat{\delta}_3 = \frac{\sum_t a_{t3}}{\sum_t v_t f_{t3}} = \frac{4}{2.3 \times 10^8} = 1.7 \times 10^{-8} \text{ overruns/train-km}$$

$$\text{ATC:}\quad \hat{\delta}_4 = \frac{\sum_t a_{t4}}{\sum_t v_t f_{t4}} = \frac{2}{1.1 \times 10^8} = 1.8 \times 10^{-8} \text{ overruns/train-km}$$

Let us assume these point estimates of the δ_i are reasonably accurate, and let us assume that the entire reduction in the aggregate signal overrun rate may be attributed to signal system upgrades. We may use the $\hat{\delta}_i$ and the f_{ti} to calculate the annual signal overrun rates "consistent" with these two assumptions.

$$\tilde{\delta}_t = \sum_i f_{ti}\hat{\delta}_i$$

Given the uncertainties in the $\hat{\delta}_i$, and given that exogenous safety trends may also explain some of the reduction in the signal overrun rate, we would not expect a perfect match between the modeled $\hat{\delta}_t$ and these "consistent" $\tilde{\delta}_t$. Nonetheless, there is substantial qualitative agreement between these two sets of rates (Figure 4-7). This provides support for the hypothesis that signal upgrades alone (Figure 4-6) can explain the improvement in the overrun rate (Figure 4-5).

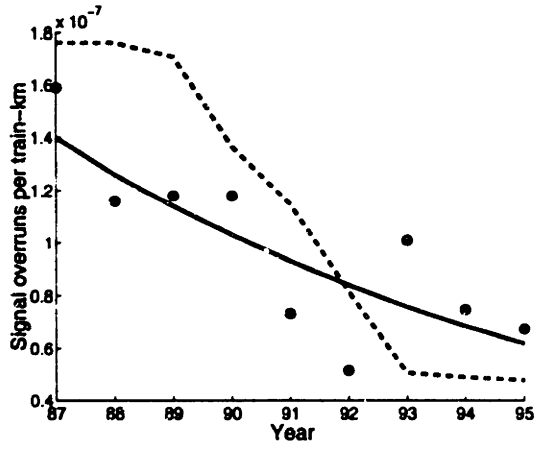

Figure 4-7: Annual signal overrun rate. Dots indicate the empirical annual rate, b_t/v_t. The solid line indicates the fitted $\hat{\delta}_t$ rate, and the dashed line indicated the "consistent" $\tilde{\delta}_t$ rate.

Summary

Let us summarize the results of our statistical analysis of JR East collisions and signal overruns. Since privatization, the JR East collision rate has been 1.3×10^{-9} collisions per train-kilometer. This rate is based on three collisions. Based on the dates of these three collisions, we cannot detect any trend in the collision rate. We note that the JR East collision rate between 1987 and 1996 is 3.5 times smaller than the Japan National Railways collision rate over the same geographical region between 1976 and 1987. Since privatization, JR East has experienced 199 signal overrun accidents; the signal overrun rate appears to have dropped by ten percent each year since privatization. Assuming a constant ratio of collisions to overruns, this improvement in the signal overrun rate may have "averted" an expected 1.4 collision accidents since privatization. We note that the signal overrun rate for the ATS-SN system is three times lower than the rate for the ATS-S/B system, and the signal overrun rate for the ATS-P and ATC systems are eleven times lower than the ATS-S/B system.

This concludes our statistical analysis of JR East collisions. We next approach this same problem from an entirely different angle.

4.4 Conceptual approach

In this section, we propose a conceptual model for train collisions. This model is inspired by the work of sociologist Charles Perrow and others on high-risk technologies. [Gla96] Our resulting model is conceptually simple and highly abstract, which has both advantages and disadvantages. We present the mathematics needed to evaluate our model, apply the model to the JR East system, and offer some conclusions.

4.4.1 "Normal Accidents"

In his book *Normal Accidents*, Charles Perrow presents a convincing framework for understanding accidents in high-risk technologies. [Per84] After studying accidents in many fields, including the nuclear industry, the airline industry, and the chemical industry, Perrow argues that highly infrequent, highly catastrophic accidents are an inevitable consequence of certain modern technologies.

Perrow lays out a hierarchy of failures. Perrow defines an "incident" as the failure of a single component that may bring a piece of a larger system to a halt. Incidents are neither unanticipated nor uncommon, and typically have safety systems associated with them to facilitate recovery. An example of an incident in the JR East system would be a jammed routing switch on the track. The railroad has special circuits to detect such failures, and these circuits ring alarms in the dispatching control center. Should a switch jam, train dispatchers would immediately reroute trains to avoid the switch if possible, and they would stop trains if not. A repair crew would be dispatched to fix the switch, and the entire system would recover smoothly and efficiently.

More severe than the incident is the "accident." An accident involves the failure of subsystems, or multiple components, that affects the entire system. An example of an accident in the JR East system would be a computer failure at the centralized train control center for ATC lines. Such a failure could bring an entire rail line to a temporary halt. While the computers were being repaired, service might resume at reduced speed. Dispatchers would have to dispatch trains "by hand," by maintaining

extensive radio communication with trains and stations, and by monitoring train positions using paper maps and chalkboards. Recovery from such an "unplanned" failure would demand creativity, cooperation, and quick-thinking on the part of many railroad employees.

Most severe is the "system accident," characterized by the unanticipated interaction of multiple failures. The "short train" Kinomiya station collision is one example of a system accident in the JR East system. First, the driver of train 6726 responded incorrectly to a green light on an adjacent track. Second, the ATS system was not prepared to stop a train moving unusually fast. Third, a subsequent switch happened to route the train into the path on an oncoming freight train. Together, these factors led to a collision. The railroad has no specific safety procedures to deal with this unusual scenario. In fact, one might argue that an element of the safety system itself bears some responsibility for this collision: a dense cluster of signal lights confused the driver.

Perrow argues that system accidents are most likely to occur in systems that are complex and tightly-coupled. The characteristics of a "complex" system include tight spacing of equipment, limited isolation of failed components, unfamiliar or unintended feedback loops, indirect or inferential information sources, and limited understanding of some of the subprocesses involved. [Per84, p. 88] A nuclear power plant is the quintessential complex system. Rail systems, in contrast, are far less complex. Rail equipment is spread out over great distances. Failures tend to be localized, and few unintended feedback loops exist. However, many of the information flows in rail systems are inferential or indirect. Dispatchers know which blocks contain trains, but may not know which train is in the block, or its position within the block. Drivers know the status of the next two approaching blocks by watching the signals, but they cannot see any farther ahead or any distance behind. Maintenance teams know scheduled train movements, but they must post watchmen ahead and behind the work site to warn of unexpected trains.

Perrow defines a "tightly-coupled" system as one where delays are "not possible," the sequence of events is invariant, alternate paths are not available, there is little

opportunity for substitutions or slack, and redundancies are "designed-in" and deliberate. [Per84, p. 96] JR East clearly fits this definition. The railroad runs on a very tight schedule. Trains can only pass one another at sidings and stations. There is typically only a single path between any two stations. It is difficult to swap trains between lines. The slack in the schedule is carefully planned and very small, as are the redundancies in track and trains.

The core of Perrow's argument is this: in a system with mature technology, good management, and strong incentives to avoid accidents, the occurrence of system accidents suggests that such catastrophes are an inevitable consequence of complex, tightly-coupled systems. In Perrow's words, such unexpected system accidents are "normal." The JR East collision record supports this contention. The three JR East collisions do not suggest flawed technology, poor management, or corporate recklessness. These three collisions seem, most simply put, accidental.

Perrow argues that complex, tightly-coupled systems naturally give rise to system accidents. When we apply this insight to the railroad, we conclude that unexpected collision accidents are inevitable. This pessimistic conclusion does not imply a railroad should abandon safety programs and invest its safety resources elsewhere; indeed, if a railroad adopted this foolhardy strategy, it would likely find itself plagued with no-longer-rare catastrophes. This conclusion simply restates the sensible notion that "zero-risk" is unattainable.

When we construct a collision model, we must address two issues: that the future occurrence of collisions is inevitable, and that we cannot anticipate all of the specific scenarios that may occur. This suggests that approaching this problem from a detailed "fault tree" framework could prove difficult. [Soc74] We instead offer a alternative framework based on Perrow's ideas.

4.4.2 Collisions as normal accidents

Most of the time, railroad safety systems operate properly. Moreover, these systems are designed redundantly to provide extra safety. For something to go seriously wrong, we argue that several failures must occur in the same place at the same time.

Suppose a train derails and comes to a sudden, unplanned stop. What must occur for the following train to crash into this disabled train from behind? First, the trains must be close, otherwise the dispatching control center would have time to learn of the failure from the first train and inform the following train to stop. Second, the driver of the following train must not notice the stopped train soon enough, otherwise he would apply the brakes and stop his train. Third, the signaling system must fail, otherwise it would not allow a second train to enter the block. We adopt a broad definition of "signaling system," including the information systems that link the track, the dispatching center, and the trains.

On the abstract level, then, three distinct types of events must occur in the same place at the same time for a collision to ensue. As "failure" is too strong a word for these precursor events, we shall instead refer to them as "intervals of stress." Most of the time, these stress intervals pass unnoticed, without consequences. We assume that there are three types of stress intervals:

- "Driver stress," denoting any time that the driver cannot respond to trouble ahead, perhaps due to poor visibility caused by weather or darkness;
- "Schedule stress," denoting any time that trains are abnormally close, usually due to congestion; and
- "Signal stress," denoting any time a signal cannot keep a second train from entering an occupied block, usually because wayside or on-board hardware have failed or been turned off.

We base our model on the assumption that a necessary precondition for a collision is that all these three stresses occur at the same place at the same time. A criticism of this assumption is the observation that the failure of a single critical component, such as train brakes, might suffice to cause a collision. If a train suffers a complete brake failure, it is likely to collide with something, even if the signals and the driver and the schedule all function perfectly. (Perrow calls these "final accidents," noting that once "the wing comes off an airplane in flight, or an earthquake shatters a dam, we hardly need further analysis.") [Per84, p. 71] However, designers are particularly skilled

at identifying such critical components, and they strive to make them as reliable as possible. We will argue from inspection and failure rates that the chance of total brake failure on JR East trains is negligible. We also note that two (and possibly all three) of the JR East collisions did not arise from the failure of a single component.

The recent Amtrak crash in Silver Spring, Maryland on February 17, 1996 provides anecdotal support for our multiple-stress framework. In this accident, an Amtrak passenger train crashed head on into a local commuter train near Kensington station. The collision killed eleven passengers on the commuter train. [Gal96b] [Heb96] [TAP96] Initial accident analyses suggest that the simultaneous occurrence of multiple "stress" factors created the dangerous situation that led to this tragedy. [Wal96]

"Signal stress" As part of a signal improvement program in 1993, a signal which had been located after the Kensington station was moved further down the track, to a location beyond the station. Because of this move, there was no signal light to inform the commuter train driver of the status of the track ahead when he departed the station. The closest previous signal was more than three miles behind the crash site.

"Driver stress" The accident occurred during a snow storm, limiting the visibility of both drivers. The driver of the Amtrak train did not apply his train brakes before the collision. The driver of the commuter train began applying his brakes as soon as he saw the Amtrak train, about ten seconds before impact.

"Schedule stress" Just before the accident, a dispatcher in the CSX Florida control center stopped a late freight train to allow the faster Amtrak train to pass. The Amtrak train was routed onto the left-hand track to pass the freight. Just before the Amtrak train was to have switched back to the right-hand track, it collided with the oncoming commuter train.

"Mechanical stress" The engine of the commuter train was in the rear of the train, pushing the passenger cars from behind. This arrangement subjected the commuter passengers to substantially greater impact force. The engine at

the front of the Amtrak train spewed diesel fuel upon impact, and this fuel ignited. The emergency exits jammed on the commuter train, trapping the passengers in the flaming wreck.

We see that the Silver Springs collision had not a single cause, but rather occurred when several systems were simultaneously stressed.

4.4.3 The triple-stress model

Viewing collisions as system accidents provides a conceptual framework for studying collision risk, and this framework leads us to a mathematical model. We call this the "triple-stress model," for at its core is the assumption that the occurrence of three stresses in the same place at the same time is a necessary precondition for a train collision accident. This model is coarse and somewhat hypothetical. Nonetheless, the framework can provide preliminary insight into the relationships between the driver system, the schedule system, and the signal system, and their collective effect on collision risk.

The scope of our triple-stress model is one signal. When all three stresses occur at the same signal at the same time, we say that this signal is in the "triple-stress state." When this occurs, we assume there is a possibility of a collision in the adjoining blocks. The triple-stress model allows us to address approximately the following questions.

- What is the probability a signal is in the triple-stress state at a random instant?
- On average, how many times each year is a signal in the triple-stress state?
- Given a signal is in the triple-stress state, what is the conditional probability of a collision in the adjoining blocks?
- What are the railroad's options to reduce the expected number of triple-stress intervals?
- How effective are these options?
- How cost-effective are these options?

Model Notation

We retain none of the notation introduced in the statistical analyses of Section 4.3, and the symbols a, b, λ, and μ take on new meanings in our triple-stress model. Let $X_1(t)$, $X_2(t)$, and $X_3(t)$ denote the state of the driver system, the schedule system, and the signal system, respectively, at time t.

$$X_i(t) = \begin{cases} 0, & \text{system } i \text{ is in the "normal" state at time } t \\ 1, & \text{system } i \text{ is in the "stressed" state at time } t \end{cases}$$

For each system i, we assume the mean duration of a stressed interval is a_i hours. We assume the mean duration of a normal interval is b_i hours. We assume the durations of stressed intervals are independent identically distributed random variables, and we make the same assumption for the duration of normal intervals. We assume the durations of normal intervals are independent of the durations of stressed intervals. We also assume the three alternating processes are independent of one another. On top of these three processes, we define a super-process $\tilde{X}(t) = \prod_i X_i(t)$; this process is zero except when the three types of stress occur simultaneously.

$$\tilde{X}(t) = \begin{cases} 1, & \text{if } \prod_{i=1}^{3} X_i(t) = 1 \quad \text{(triple-stress state)} \\ 0, & \text{otherwise} \quad \text{(non-triple-stress state)} \end{cases}$$

Probability that a signal is in the triple-stress state

From elementary renewal theory [Ros89, p. 320], we know the probability that $X_i(t) = 1$ at a random instant is $a_i/(a_i + b_i)$ (Figure 4-8). As we assume the occurrence of driver stress, schedule stress, and signal stress are independent, the probability that all three stresses occur at the same signal at the same random instant is the product of the probabilities that each occurs at a random instant.

$$\tilde{p} = \Pr\left(\tilde{X}(t) = 1\right) = \prod_{i=1}^{3} \Pr(X_i(t) = 1) = \prod_{i=1}^{3} \frac{a_i}{a_i + b_i} \tag{4.1}$$

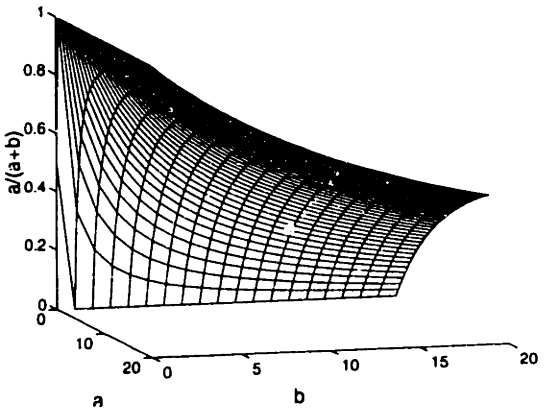

Figure 4-8: The function $f(a, b) = a/(a + b)$.

Annual occurrences of the triple-stress state

The railroad operates eighteen hours each day, or approximately 6,500 operating hours each year. The railroad is out-of-service for six hours each day, from midnight until 6am. This daily "rest period" serves to insulate each day from its predecessors, preventing schedule stress and driver stress from spilling from one day into the next. (In contrast, inoperative signals do not automatically become operative each morning.) In the context of our model, we might propose that the driver stress $X_1(t)$ and the schedule stress $X_2(t)$ processes should reset to the normal state at the beginning of each day. We do not include this feature in our model. Instead, we assume the four processes $X_1(t)$, $X_2(t)$, $X_3(t)$, and $\tilde{X}(t)$ occur in operating time, where t measures cumulative hours of operation. This assumption erases the daily six-hour rest periods. As a result, our model allows driver and schedule stress to spill over from one calendar day into the next. This assumption introduces some additional inaccuracy into an already coarse model. However, we believe this assumption does not affect our general conclusions.

For each system i, the average duration of a stressed interval is a_i hours and the average duration of a normal interval is b_i hours. Thus, on average, the time between the onset of successive stress intervals is $(a_i + b_i)$ hours. On average, then, there are $6,500/(a_i + b_i)$ stress intervals each year. In a similar fashion, we can determine

the expected number of triple-stress intervals each year. Let $\tilde{a}$ denote the average duration of a triple-stress interval in hours. Let $\tilde{b}$ denote the average duration of a non-triple-stress interval in hours (that is, the average time between the end of one triple-stress interval to the beginning of the next). Analogously, we expect an average of $6,500/(\tilde{a} + \tilde{b})$ triple-stress intervals to occur at each signal each year.

However, we do not know the value of $\tilde{a}$ or $\tilde{b}$. With one, we could determine the other using the relationship between the average duration of a triple-stress interval, the average duration of a non-triple-stress interval, and the probability of the triple-stress state at a random instant. Let us assume we know the value of $\tilde{a}$, and use $\tilde{p}$ to determine $\tilde{b}$:

$$\tilde{p} = \frac{\tilde{a}}{\tilde{a} + \tilde{b}}, \quad \text{thus} \quad \tilde{b} = \tilde{a}(\frac{1}{\tilde{p}} - 1). \tag{4.2}$$

Average duration of a triple-stress interval

We now determine $\tilde{a}$, the average duration of a triple-stress interval. For tractability, we impose an additional assumption on these three systems: we use the Erlang family of distributions to model the duration of stressed and normal intervals. With an Erlang distribution, we can model many types of system behavior, such as an exponentially distributed duration, an (approximately) normally distributed duration, and an (approximately) constant duration (Figure 4-9).

We call the $X_i(t)$ "alternating Erlang processes," so named because each consists of draws from two alternating Erlang distributions. As before, let $X_i(t) \in \{0, 1\}$ denote the state of system i at time t; we call state 1 the "stressed" state, and we call state 0 the "normal" state. Each system i stays in the stressed state for a random length of time, and we model this duration using an Erlang random variable of order m_i and mean $a_i = m_i/\mu_i$. Similarly, we model the duration of a normal interval with an Erlang random variable of order n_i and mean $b_i = n_i/\lambda_i$. We assume all normal durations and stressed durations are independent. Each of these three alternating Erlang process is a two-state, semi-Markov process. [Gal96a, p. 174] However, it is more convenient to view them as $(n_i + m_i)$-state Markov processes. Adopting this view for each system i, we decompose the normal state into states 1 to n_i and the

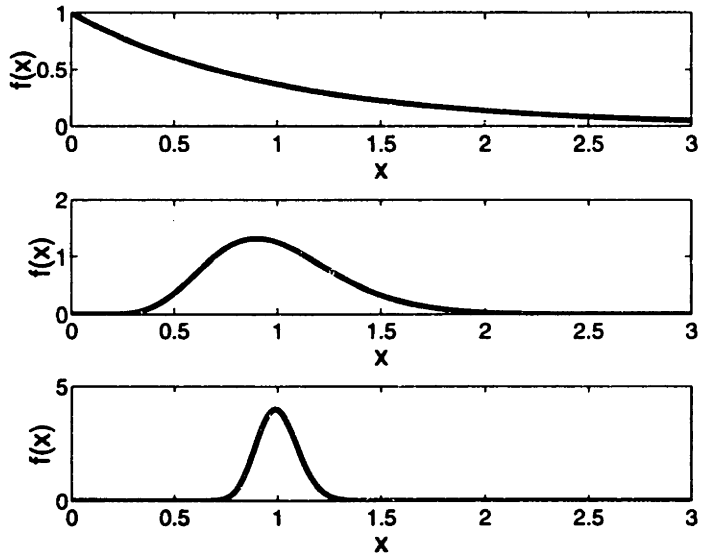

Figure 4-9: The Erlang density function. This plot shows three Erlang probability density functions, all with mean one. The upper plot is a first-order Erlang, which is an exponential distribution. The middle plot is a tenth-order Erlang; due to the Law of Large Numbers, this distribution is approximately normal. The lowest plot is an fiftieth-order Erlang. This distribution is also approximately normal, and has lower variance. By making the Erlang order sufficiently high while keeping the mean fixed, one may make the Erlang variance arbitrarily low, and thus approximate a constant.

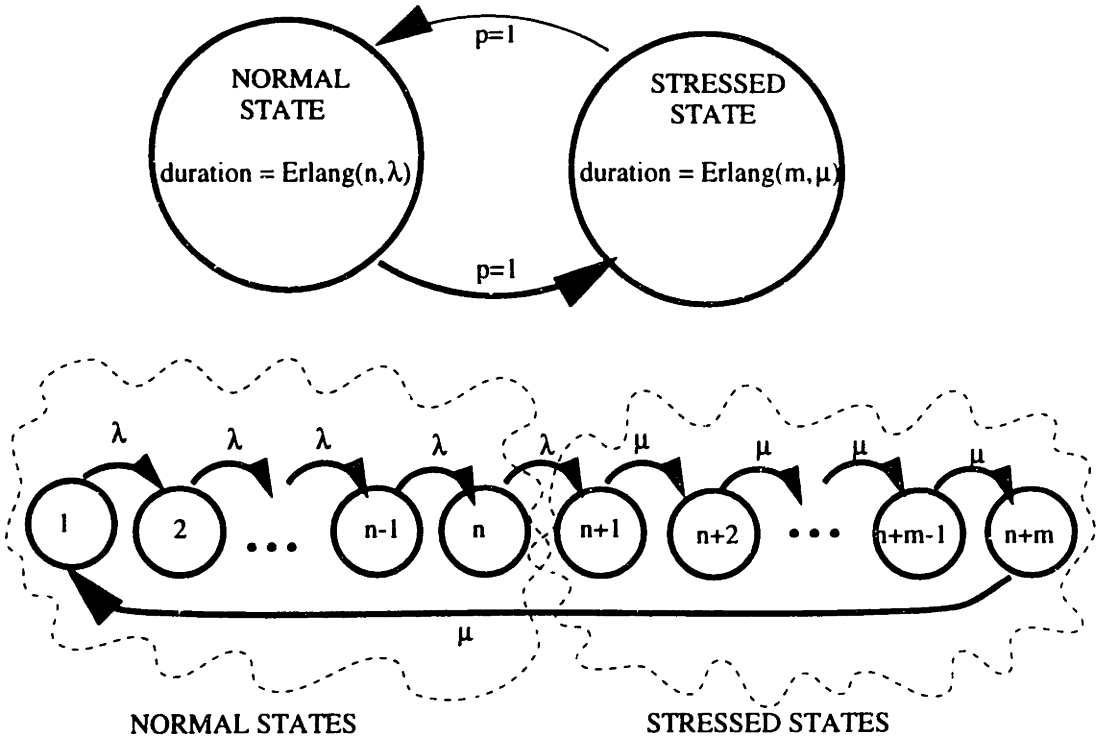

Figure 4-10: Two equivalent views of the alternating Erlang process. In the top diagram, we represent this process as a two-state semi-Markov process. In the lower diagram, we represent this process as a $(n+m)$-state Markov process.

stressed state into states n_i+1 to n_i+m_i (Figure 4-10). Determining the steady-state probabilities for these chains is straight-forward.

$$\begin{aligned} p(k; n_i, \lambda_i, m_i, \mu_i) &= \Pr(\text{system } i \text{ in state } k) \\ &= \begin{cases} \frac{1/\lambda_i}{n_i/\lambda_i + m_i/\mu_i}, & k = 1, \ldots, n_i \\ \frac{1/\mu_i}{n_i/\lambda_i + m_i/\mu_i}, & k = n_i+1, \ldots, n_i+m_i \end{cases} \end{aligned} \tag{4.3}$$

As before, we define the super-process for triple-stress intervals based on the three systems: $\vec{X}(t) = \prod_i X_i(t)$. The super-process is also a semi-Markov process, which we similarly decompose into a $(n_1+m_1)\times(n_2+m_2)\times(n_3+m_3)$-state Markov process. We denote the state of the super-process with a vector $\vec{k} = \{k_1, k_2, k_3\}$; the ith element of this vector represents the state of the ith chain. As the three chains are independent,

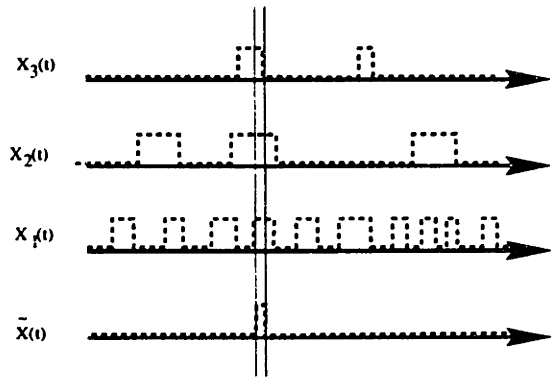

Figure 4-11: Example of the super-process $\tilde{X}(t)$ for triple-stress intervals. This sketch illustrates that $\tilde{X}(t) = 1$ only when $X_1(t) = X_2(t) = X_3(t) = 1$. The horizontal axis corresponds to time, and vertical lines denote the epochs at which the super-process changes state.

the steady-state probabilities of the super-process take product-form.

$$P(\vec{k}; \vec{n}, \vec{\lambda}, \vec{m}, \vec{\mu}) = \prod_{i=1}^{3} p(k_i; n_i, \lambda_i, m_i, \mu_i) \tag{4.4}$$

We wish to determine the mean duration of a triple-stress interval. Consider the super-process just as a triple-stress interval begins. At this instant, one of the three processes just entered the stressed state, and the other two processes were both already in the stressed state. For expository convenience, and without loss of generality, we assume that Process 1 just entered the stressed state, and Processes 2 and 3 were already in the stressed state. This triple-stress interval will continue as long as all three processes remain in the stressed state. It will end as soon as one of the three processes leaves the stressed state.

As process one just entered the stressed state, Process 1 will remain stressed for some duration drawn from an Erlang(m_1, μ_1) distribution. In this example, Processes 2 and 3 have been stressed for some time already. If we decompose their Erlang distributed stressed durations into sums of exponential random variables, and if we know the "stage" of their stress, we can exploit the memorylessness property of the exponential distribution to find the residual stressed durations for Process 2 and 3. For example, let us suppose the stressed duration of Process 2 follows an Erlang(m_2, μ_2)

distribution with $m_2 = 10$. Instead of viewing the event that terminates this stressed interval as the first arrival in a tenth-order Erlang process, we instead view it as the tenth arrival in an ordinary Poisson process of rate μ_2. If, for example, Process 2 was in "stage" six of its ten "stages" at the instant Process 1 became stressed, it must finish this stage and then four more before returning to the normal state. Thus, the residual stressed duration for Process 2 follows a fifth-order Erlang distribution. Likewise, suppose that at the instant Process 1 became stressed, Process 3 was in stage k_3 of its m_3+n_3 stages, where $n_3 \leq k_3 \leq n_3+m_3$. Before Process 3 returns to the non-stressed state, it must finish its current stage (and due to the memoryless property of the exponential distribution, the residual duration has the same distribution as the full duration) and then finish $(m_3 + n_3 - k_3)$ more stages. Accordingly, the residual stressed duration of Process 3 follows an $(m_3+n_3-k_3+1)$th-order Erlang distribution.

As a triple-stress interval of the super-process ends as soon as one of the three processes leaves the stressed state, we see that the duration of a triple-stress interval depends on the minimum of three Erlang random variables.

Our goal is to calculate the mean duration of a triple-stress interval. We see that this calculation will require the expected value of the minimum of three independent Erlang random variables. We digress briefly to present our method to calculate such expected values. Let $Y_1(k_1)$ denote an Erlang(k_1, μ_1) random variable, $Y_2(k_2)$ denote an Erlang(k_2, μ_2) random variable, and $Y_3(k_3)$ denote an Erlang(k_3, μ_3) random variable. We define a function $W(k_1, k_2, k_3)$ that provides the expected value of the minimum of our independent three Erlangs as a function of their orders. (Note the three Erlangs may have different rates.)

$$W(k_1, k_2, k_3) \equiv E(\min\{Y_1(k_1), Y_2(k_2), Y_3(k_3)\})$$

We again consider these Erlang random variables in terms of their underlying Poisson processes. We can also view these three independent Poisson processes as one combined process with rate $(\mu_1 + \mu_2 + \mu_3)$. With this view, the probability that the next arrival in the combined process comes from Process 1 is $\mu_1/(\mu_1 + \mu_2 + \mu_3)$, the

probability the next comes from Process 2 is $\mu_2/(\mu_1+\mu_2+\mu_3)$, and the probability the next comes from Process 3 is $\mu_3/(\mu_1+\mu_2+\mu_3)$. With this insight, we can define $W(k_1,k_2,k_3)$ recursively. On average, we must wait $1/(\mu_1+\mu_2+\mu_3)$ for the next arrival. Once it arrives, if it is from Process 1, we then must wait $W(k_1-1,k_2,k_3)$ more. If it is from Process 2, we must wait $W(k_1,k_2-1,k_3)$ more. If it is from Process 3, we must wait $W(k_1,k_2,k_3-1)$ more. These arguments, along with sensible boundary conditions, provide a recursive form for $W(k_1,k_2,k_3)$.

$$\begin{aligned} W(k_1,k_2,k_3;\mu_a,\mu_2,\mu_3) &= \frac{1}{\mu_1+\mu_2+\mu_3} \\ &+ \frac{\mu_1}{\mu_1+\mu_2+\mu_3} W(k_1-1,k_2,k_3;\mu_1,\mu_2,\mu_3) \\ &+ \frac{\mu_2}{\mu_1+\mu_2+\mu_3} W(k_1,k_2-1,k_3;\mu_1,\mu_2,\mu_3) \\ &+ \frac{\mu_3}{\mu_1+\mu_2+\mu_3} W(k_1,k_2,k_3-1;\mu_1,\mu_2,\mu_3) \\ &\quad \text{for } k_1 \geq 1, k_2 \geq 1, k_3 \geq 1, \\ W(k_1,k_2,k_3;\mu_1,\mu_2,\mu_3) &= 0 \text{ otherwise.} \end{aligned} \tag{4.5}$$

We have now modeled the occurrence of stress in the driver component, the schedule component, and signal components using three independent alternating Erlang processes. We have decomposed these two-state semi-Markov processes into (n_i+m_i)-state Markov processes. We have defined a triple-stress super-process on top of these three processes. The super-process is also a Markov process, and we have determined its steady-state distribution. We have shown the duration of a triple-stress interval depends on the minimum of three independent Erlang random variables. and we have presented a method to compute the expected value of such minimums. With these pieces in place, we can finally calculate the expected duration of a triple-stress interval.

Consider a triple-stress interval at the instant it begins. We know one of the processes just entered the stressed state, and the other two processes are already in the stressed state. We now describe the possible states of the three processes at this

instant. We build up a collection of sets to do this. Let the set $\mathcal{S}_i$ denote the stressed states of Process i.

$$\begin{aligned}
\mathcal{S}_1 &= \{k_1 : n_1 + 1 \leq k_1 \leq n_1 + m_1\} \\
\mathcal{S}_2 &= \{k_2 : n_2 + 1 \leq k_2 \leq n_2 + m_2\} \\
\mathcal{S}_3 &= \{k_3 : n_3 + 1 \leq k_3 \leq n_3 + m_3\}
\end{aligned}$$

Let the set $\mathcal{L}_i$ denote the possible states of the three processes, given that a triple-stress interval just began with the entrance of Process i into the stressed state.

$$\begin{aligned}
\mathcal{L}_1 &= \{(k_1, k_2, k_3) : \; k_1 = n_1 + 1, \; k_2 \in \mathcal{S}_2, \quad k_3 \in \mathcal{S}_3 \quad \} \\
\mathcal{L}_2 &= \{(k_1, k_2, k_3) : \; k_1 \in \mathcal{S}_1, \quad k_2 = n_2 + 1, \; k_3 \in \mathcal{S}_3 \quad \} \\
\mathcal{L}_3 &= \{(k_1, k_2, k_3) : \; k_1 \in \mathcal{S}_1, \quad k_2 \in \mathcal{S}_2, \quad k_3 = n_3 + 1 \quad \}
\end{aligned}$$

Let the set $\mathcal{T}$ denote the set of possible states for the three processes, given that a triple-stress interval just began.

$$\mathcal{T} = \mathcal{L}_1 \cup \mathcal{L}_2 \cup \mathcal{L}_3$$

We can use the steady-state probabilities from Equation 4.4 to determine the conditional steady-state probability the system is in state (k_1, k_2, k_3), given that a triple-stress interval just began. Let D denote the event, "a triple-stress interval just started."

$$P(\vec{k}; \vec{n}, \vec{\lambda}, \vec{m}, \vec{\mu} | D) = \begin{cases} \frac{P(\vec{k}; \vec{n}, \vec{\lambda}, \vec{m}, \vec{\mu})}{\sum_{h \in \mathcal{T}} P(\vec{h}; \vec{n}, \vec{\lambda}, \vec{m}, \vec{\mu})}, & \vec{k} \in \mathcal{T} \\ 0, & \text{else.} \end{cases} \tag{4.6}$$

Suppose a triple-stress interval just began, and the system is in state $\vec{k} \in \mathcal{T}$. Let $\vec{e} = [\, 1 \quad 1 \quad 1 \,]$. Given that the system is in state $\vec{k}$, the triple-stress interval will continue for an average duration of $W(\vec{e} + \vec{m} + \vec{n} - \vec{k})$. That is, we calculate the conditional expected duration of the triple-stress interval using the expected value of the minimum of three Erlang random variables, because the triple-stress interval ends as soon as one of the three processes leaves the stressed state. Equation 4.6 provides the

steady-state distribution of $(\vec{k}|D)$, and we may now determine the expected duration of a triple-stress interval.

$$\begin{aligned}\tilde{a} &= \sum_{\vec{k}\in\mathcal{T}} E(\text{duration triple-stress interval}|\vec{k}, D)P(\vec{k};\vec{n},\vec{\lambda},\vec{m},\vec{\mu}|D) \\ &= \sum_{\vec{k}\in\mathcal{T}} W(\vec{e}+\vec{m}+\vec{n}-\vec{k},\vec{\mu})P(\vec{k};\vec{n},\vec{\lambda},\vec{m},\vec{\mu}|D) \qquad (4.7)\end{aligned}$$

Conditional collision probability

We have presented a method to determine $\tilde{a}$, the mean duration of a triple-stress interval. Knowing $\tilde{p}$, the probability the super-process is in the triple-stress state at a random instant, we can calculate $\tilde{b}$, the mean duration of a non-triple-stress interval: $\tilde{b} = \tilde{a}(\frac{1}{\tilde{p}} - 1)$. On average, each signal has $6,500/(\tilde{a}+\tilde{b})$ triple-stress intervals each year. With approximately 14,000 signals in the JR East system, we expect an average of $(14,000)(6,500)/(\tilde{a}+\tilde{b}) \approx 10^8/(\tilde{a}+\tilde{b})$ triple-stress intervals across the system each year.

We wish to estimate the probability that a triple-stress interval progresses to a collision. We shall assume that during triple-stress intervals, the occurrence of collisions follows a temporal Poisson process with rate γ collisions/hour. Accordingly, the probability that one or more collisions occurs during a triple-stress interval of length τ hours is $1 - e^{-\tau\gamma}$. Recall that there have been three collisions to date in the JR East system. In the next section, we shall use the triple-stress model to estimate the annual number of triple-stress intervals across the JR East system, and we shall find that triple-stress intervals occur with much greater frequency than collisions. This suggests that γ is very small. Because γ is small, the probability of two or more collisions in one triple-stress interval is effectively zero.

$$\begin{aligned}\Pr(0 \text{ collisions in a triple-stress interval of length } \tau) &\approx 1-\gamma\tau \\ \Pr(1 \text{ collisions in a triple-stress interval of length } \tau) &\approx \gamma\tau \\ \Pr(2+ \text{ collisions in a triple-stress interval of length } \tau) &\approx 0\end{aligned}$$

For convenience, we approximate this Poisson process with a Bernoulli process, by

implicitly averaging over the variation in the duration of triple-stress intervals. We simply assume each triple-stress interval either has a collision, with probability p_c, or does not, with probability $1 - p_c$.

$$\begin{aligned} \Pr(0 \text{ collisions in a triple-stress interval}) &= 1 - p_c \\ \Pr(1 \text{ collisions in a triple-stress interval}) &= p_c \end{aligned}$$

We would like to estimate p_c, the probability that a collision occurs during a triple-stress interval. The obvious estimator of p_c is the ratio of the total number of collisions to the total number of triple-stress intervals.

Our earlier statistical analysis of signal overrun events suggests that system safety has improved since privatization. To obtain a rough estimate of p_c, however, we shall assume that both $(\tilde{a} + \tilde{b})$ and p_c have held constant since privatization. With these assumptions, and with three collisions in nine years of triple-stress periods, we calculate a rough estimate of p_c.

$$\begin{aligned} p_c &= \Pr(\text{triple-stress interval progresses to a collision}) \\ &\approx \frac{3}{(14{,}000 \text{ signals})(9 \text{ years})(6{,}500/(\tilde{a} + \tilde{b}))} \\ &= (3.6 \times 10^{-9})(\tilde{a} + \tilde{b}) \\ &\approx (3.6 \times 10^{-9})(\tilde{b}) \end{aligned}$$

The reliability of this estimate may not be high. As defined in our model, triple-stress intervals are unobservable, except when a collision occurs. It is thus impossible to collect data on the occurrence of triple-stress intervals. Suppose we err when estimating the annual number of triple-stress intervals, and we obtain a number that is one thousand times too large. We would then estimate p_c using this flawed number, and obtain an estimate of p_c three orders of magnitude too small. With no independent means to check, we would be none the wiser. This raises a serious criticism of the triple-stress model. If triple-stress periods are invisible, how may we verify the model? What experiment or data allows us to falsify the model? We answer that the triple-stress model can be neither verified nor falsified. We suggest that this abstract

model offers useful insight despite these concerns.

Reducing the collision risk

If one accepts the triple-stress model for collision risk, there are two strategies to decrease the expected annual number of collisions:

- improve the driving, schedule, and/or signal systems to reduce the annual number of triple-stress intervals, or
- decrease the probability that a triple-stress interval progresses to a collision.

We focus on the first of these strategies, as it allows us to understand the role of the three systems in collision risk. It also avoids the difficulty of estimating p_c. To reduce the expected number of annual triple-stress periods, the railroad could endeavor to:

- decrease the a_i, reducing the mean duration of a stressed interval, or
- increase the b_i, increasing the mean duration of a normal interval,

for each system in isolation, or for some combination of the three systems.

Relative efficacy of improvements

According to the triple-stress model, decreasing some or all of the a_i and/or increasing some or all of the b_i will reduce the expected number and expected duration of triple-stress intervals each year. Let us first focus on a simpler quantity of interest: $\tilde{p}$, the probability that a signal is in the triple-stress state at a random instant. For convenience, we bundle the six parameters into a vector $\vec{x}$, and express $\tilde{p}$ as a function of this vector. The superscript "t" here denotes the transpose operator.

$$\vec{x} = [\, a_1 \quad a_2 \quad a_3 \quad b_1 \quad b_2 \quad b_3 \,]^t$$

$$\tilde{p} = F(\vec{x}) = \prod_{i=1}^{3} \frac{a_i}{a_i + b_i}$$

To determine how $\tilde{p}$ responds to small changes in the a_i and b_i, we approximate the function F with its first-order Taylor expansion. [Lue89, p. 461]

$$\nabla F(\vec{x}) = \left[\begin{array}{cccccc} \frac{b_1}{a_1(a_1+b_1)} & \frac{b_2}{a_2(a_2+b_2)} & \frac{b_3}{a_3(a_3+b_3)} & \frac{-1}{a_1+b_1} & \frac{-1}{a_2+b_2} & \frac{-1}{a_3+b_3} \end{array} \right]^t \cdot F(\vec{x})$$

$$\begin{aligned} F(\vec{x}+\vec{d}) &\approx F(\vec{x}) + \nabla F(\vec{x})^t \vec{d} \\ &= F(\vec{x}) \left[1 + \frac{\nabla F(\vec{x})^t \vec{d}}{F(\vec{x})} \right] \\ &= F(\vec{x}) \left[1 + \sum_{j=1}^{3} d_j \left(\frac{b_j}{a_j(a_j+b_j)} \right) - \sum_{j=4}^{6} d_j \left(\frac{1}{a_j+b_j} \right) \right] \end{aligned} \tag{4.8}$$

Let $\vec{e_k}$ denote a vector with the kth element equal to one and all other elements zero. We investigate a simple form for the perturbation vector: $\vec{d} = \delta x_k \vec{e_k}$. We perturb each of the six parameters of the triple-stress model in turn with a constant percentage change: that is, we choose $\vec{d}$ so that $(\vec{x}+\vec{d})$ corresponds to a $100\delta\%$ percent change in the kth component of $\vec{x}$. We investigate such perturbations by substituting this $\vec{d}$ back into the Taylor approximation.

$$F(\vec{x}+\vec{d}) \approx \begin{cases} F(\vec{x})[1 + (\delta b_k)/(a_k+b_k)], & k = 1,2,3 \\ F(\vec{x})[1 - (\delta b_{k-3})/(a_{k-3}+b_{k-3})], & k = 4,5,6 \end{cases} \tag{4.9}$$

This implies that for small changes, a one percent increase in

$$\left\{ \begin{array}{c} a_1 \\ a_2 \\ a_3 \\ b_1 \\ b_2 \\ b_3 \end{array} \right\} \text{ leads to a } \left\{ \begin{array}{c} \frac{b_1}{a_1+b_1} \\ \frac{b_2}{a_2+b_2} \\ \frac{b_3}{a_3+b_3} \\ \frac{b_1}{a_1+b_1} \\ \frac{b_2}{a_2+b_2} \\ \frac{b_3}{a_3+b_3} \end{array} \right\} \text{ percent } \left\{ \begin{array}{c} \text{increase} \\ \\ \\ \\ \text{decrease} \\ \\ \end{array} \right\} \text{ in } \tilde{p}.$$

We offer three observations regarding $\tilde{p}$, the probability a signal is in the triple-stress state at a random instant.

- For small percentage changes, a small percentage decrease in a_i produces the same reduction in $\tilde{p}$ as the same percentage increase in b_i.

- When varying each parameter singly, $\tilde{p}$ is most sensitive to parameter changes to the most reliable system. That is, a one percent improvement in the system with the largest $b_i/(a_i + b_i)$ yields the largest percentage improvement in $\tilde{p}$.

- Because $b_i/(a_i + b_i) < 1$, a one percent improvement to any one of the six parameters yields less than a one percent improvement in $\tilde{p}$.

The second of these points merits comment. Intuitively, one might think that the least reliable subsystem is the one that provides the greatest overall benefit if improved. This intuition is true for a system comprised of linked subsystems in series. In such a system, if any one subsystem fails, the entire system fails. Just as a chain is only as strong as its weakest link, such systems are most improved by improving the least reliable subsystem. In contrast, the triple-stress model corresponds to a system with three subsystems in parallel. In such a system, all three subsystems must fail to bring down the entire system. For an analogy, consider supporting a piano with a piece of thread, a piece of twine, and a piece of steel cable. Clearly, one's concern should focus on the strength of the cable, as it bears the brunt of the load. Accordingly, for parallel systems, improving the most reliable subsystem offers the largest benefit.

Cost-effectiveness of improvements

Suppose the railroad is evaluating two proposed safety initiatives which may reduce collision risk. The first initiative promises to improve system i by increasing b_i or decreasing a_i by δ_i percent at a cost of c_i yen. The second initiative promises to improve system j by increasing b_j or decreasing a_j by δ_j percent at a cost of c_j yen. Assuming these estimates of costs and improvements are accurate, and assuming δ_i and δ_j are reasonably small, we may use the previous analysis to determine the more cost-effective of the two initiatives.

- If $\frac{c_i}{c_j} < \frac{\delta_i \frac{b_i}{a_i+b_i}}{\delta_j \frac{b_j}{a_j+b_j}}$, then initiative 1 is more cost-effective.

- If $\frac{c_i}{c_j} > \frac{\delta_i \frac{b_i}{a_i+b_i}}{\delta_j \frac{b_j}{a_j+b_j}}$, then initiative 2 is more cost-effective.

A constant interval inspection model

Before investigating the triple-stress model numerically, we pause to consider the reliability of train brakes in the JR East system. As mentioned earlier, the complete failure of a train's brakes would fit Perrow's description of a "final accident," those single component failures that can cause catastrophe. In the context of a rail system, a train which suffers complete brake failure is likely to collide with something—regardless of the state of the driver system, the schedule system, and the signaling system. We shall argue that the rate of complete brake failure in the JR East system is so small that we may exclude this scenario from our analysis. To this end, we introduce a simple constant interval inspection model. The goal of this model is to relate the probability a component is found failed at inspection, p_I, to the probability a component is failed at a random instant, p_R.

Consider a generic component that operates for an exponentially distributed duration before failing, with a mean time to failure of $1/\alpha$. Let random variable L denote the duration of the component's operating life. We assume this component is inspected or tested every T time units, where T is a fixed constant. We assume that if the component found failed at inspection, it is immediately replaced. Suppose this component is found failed at inspection with some probability p_I. As this component's mean time to failure is exponentially distributed, we know that $p_I = 1 - e^{-\alpha T}$. We use this relationship, the observable p_I, and the known T to estimate the hidden parameter: $\alpha = -\ln(1 - p_I)/T$. For convenience, we take the inspection interval as our unit of time. Thus, $T = 1$ and $\alpha = -\ln(1 - p_I)$, where α is measured in reciprocal inspection intervals.

We next determine the average fraction of an inter-inspection interval the component spends in the operating state. We compute this expected value by straightforward integration, and determine $E(\max\{L, 1\}) = (1 - e^{-\alpha})/\alpha$. This quantity also

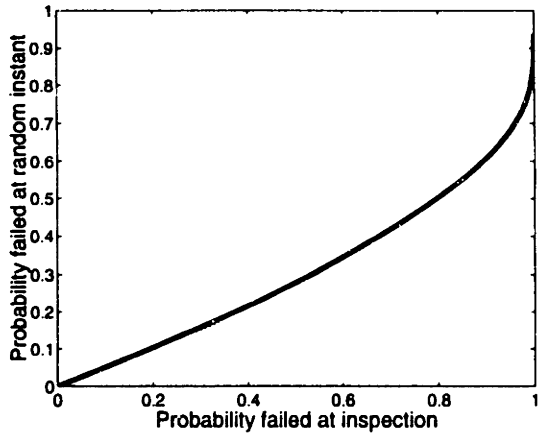

Figure 4-12: The constant interval inspection model. This plot indicates the probability the component is failed at a randomly-chosen instant, as a function of the probability it is found failed at inspection. The dotted line indicates the line $y = 0.5x$.

represents the probability that the component is operating at a randomly-chosen instant. Its complement, $1 - (1 - e^{-\alpha})/\alpha$, provides the probability the component is failed at a randomly-chosen instant. Into this expression we substitute $\alpha = -(\ln(1 - p_I))$ to obtain our fixed-interval inspection model.

$$p_R = \Pr(\text{component failed at random instant}) = \frac{\ln(1 - p_I) + p_I}{\ln(1 - p_I)}$$

This model allows us to relate the probability a component is found failed at inspection to the probability a component is failed at a random instant (Figure 4-12). As p_I becomes small, p_R converges to $p_I/2$, with the limit derived from L'Hôpital's Rule.

$$\lim_{p_I \to 0} p_R = \lim_{p_I \to 0} \frac{\ln(1 - p_I) + p_I}{\ln(1 - p_I)} = \frac{p_I}{2}$$

Total brake failure as a "final accident"

JR East trains have two braking systems: primary brakes, and emergency brakes. The primary brakes are employed first in all braking situations; should the primary brakes fail, the emergency brakes provide automatic back-up. For a train to suffer

a complete loss of brakes, both the primary and the emergency systems would have to fail simultaneously. We shall assume that these two systems fail independently: if the probability the primary brakes have failed at a random instant is p_p and the probability that the emergency brakes have failed at a random instant is p_e, we assume the probability that both have simultaneously failed is $p_p p_e$. One might argue that this is an invalid assumption because these systems are mechanically linked and spatially proximate. For example, a brake pad in the primary braking system might tear free and jam in the emergency braking mechanism, rendering both systems inoperative. However, railroad engineers have designed these systems to be as independent as possible, and it is difficult to assess the degree of dependence between the two systems. We thus shall model the two systems as if they are independent.

In addition to scheduled brake maintenance, JR East tests the emergency brakes on each of its trains every three operating days. As each train receives 120 inspections each year, and as there are roughly 575 operating trains in the JR East system, we estimate that JR East has performed approximately 620,000 such emergency brake checks since privatization. Amazingly, the Safety Research Laboratory reports that no train has ever failed this check. [SRL96b] Let us treat these inspections as independent Bernoulli trials, and let us assume the probability that a train fails inspection is p_f. The probability of no failures in 620,000 inspections is $(1 - p_f)^{620,000}$; to obtain a 99.99% one-sided confidence interval p_f, let us choose p_f to make this probability less than one in ten thousand.

$$(1 - p_f)^{620,000} < 10^{-4} \text{ if } p_f \in [0, 1.5 \times 10^{-5}]$$

Let us adopt this conservative upper bound for the probability a train fails an emergency brake inspection: $p_f = 1.5 \times 10^{-5}$. From the constant interval inspection model, we know that the chance a train has failed emergency brakes at a random instant is one half of the probability that a train has failed emergency brakes at inspection: thus $p_e < 7.5 \times 10^{-6} = p_f/2$.

The Safety Research Laboratory reports that there have been approximately 10

cases of primary brake failure across the JR East system since privatization, for a rate of approximately one failure per year. [SRL96a] In each of these ten cases, the emergency brakes performed properly and stopped the train. Assuming the independence of the primary and emergency braking systems, the probability a primary brake failure is accompanied by an emergency brake failure is p_e. This implies that the annual rate of total brake failures for the entire system is below 7.5×10^{-6} failures/year. Accordingly, the expected time until a train in the JR East system suffers a complete brake failure is more than 100,000 years. Even if this coarse estimate is off by three orders of magnitude, we remain convinced that it is reasonable to exclude the scenario of total brake failure from our analysis. We next return to the triple-stress model and present some numerical examples.

4.4.4 Numerical example of triple-stress model

The triple-stress model is a conceptual model. It allows one to explore the general relationships between the reliability of the driver system, the schedule system, and the signaling system to collision risk. As we have discussed, we do not trust the model to provide reliable estimates of specific quantities. With this caveat, we demonstrate the triple-stress model. The model has six main parameters: a_1, a_2, a_3, the mean duration of a stressed interval for each the three systems; and b_1, b_2, b_3, the mean duration of a normal interval for each of the three systems. The model has six secondary parameters: n_1, n_2, n_3, m_1, m_2, and m_3. These provide some control over the variance of stressed and non-stressed durations: for system i, the coefficient of variation for the duration of stressed intervals is $1/\sqrt{m_i}$, and the coefficient of variation for the duration of normal intervals is $1/\sqrt{n_i}$.

Of the three systems, we believe the driver system is most often in the stressed state. The driver system is in a stressed state when the driver is unable to respond properly to a dangerous situation. This stress could occur due to limited visibility arising from inclement weather, darkness, or curved areas of track. Less frequently, this stress could arise from driver fatigue, boredom, or sickness. We will assume a driver spends, on average, five minutes of each hour in the stressed state: $a_1 = 5/60$

hour, and $b_1 = 55/60$ hour. We note these figures may be high.

The schedule system is in a stressed state when trains get off schedule. Across the JR East system, this occurs approximately 10,000 times each year, and each disruption lasts an average of forty-five minutes. [SRL96a] We assume that each schedule disruption affects 5% of the signals in the system. Thus, each signal experiences an average of 500 disruptions each year. Recall that the mean number of stressed intervals each year is $6,500/(a_i + b_i)$. This implies that the average duration of a normal period is 12 1/4 hours, as $6,500/(12.25 + 0.75) = 500$. Thus, we estimate $a_2 = 3/4$ hour, and $b_2 = 12$ 1/4 hours.

Of the three systems, the signaling system is most rarely in the stressed state. A signal is in the stressed state when it cannot stop a train from running through a red signal. This fault could arise from problems with the track circuitry, problems with the signal, or problems with the on-board ATS or ATC protection devices. The signaling system is extremely reliable. Signals have failed less than ten times since privatization, and the on-board ATS and ATC controllers share equally low failure rates. [SRL96a] However, the railroad has experienced problems with signals during track repair and maintenance. These operations may involve turning off a signal for some period, or wiring a temporary signal into the system. Such work may leave a signal in a stressed state. With some uncertainty, we will assume the average duration of a normal interval for the signaling system is five years, reflecting a typical inter-repair interval for a signal. Further, we assume the average duration of a stress interval for the signaling system is 3 hours, reflecting the duration of a maintenance project. Thus, we set $a_3 = 3$ hours and $b_3 = (6,500)(5) = 32,500$ hours.

We next need to set the variability of these intervals. Lacking any solid evidence, we will take $n_i = m_i = 1$ for all three systems as our base case. With these values, we model all durations with exponential random variables, and thus all durations have coefficients of variation equal to one. We shall later return to this assumption and investigate the role variance plays in the triple-stress model.

Table 4.5 summarizes our base-case parameters for the triple-stress model. When we evaluate the triple-stress model with these values, we find that:

		Driver system	Schedule system	Signaling system
Normal state	Distribution	Erlang(0.92, 1)	Erlang(12.25, 1)	Erlang(32.500, 1)
	Mean	55 minutes	12.25 hours	5 years
	Coeff. Var.	1	1	1
	Pr(normal)	0.92	0.94	0.999
Stressed state	Distribution	Erlang(0.083, 1)	Erlang(0.75, 1)	Erlang(3, 1)
	Mean	5 minutes	45 minutes	3 hours
	Coeff. Var.	1	1	1
	Pr(stressed)	0.08	0.06	0.001

Table 4.5: Base-case parameters for triple-stress risk model. "Pr(normal)" denotes the probability the system is in the normal state at a randomly-chosen instant, and "Pr(stressed)" denotes the probability the system in the stressed state at a randomly-chosen instant.

- The probability that a signal is in the triple-stress state at a random instant is $\tilde{p} = 4 \times 10^{-7}$.
- The average duration of a triple-stress interval is $\tilde{a} = 4$ minutes.
- On average, a signal experiences a triple-stress interval every $(\tilde{a} + \tilde{b}) = 27$ years.
- With 14,000 signals, the JR East system experiences an average of $14{,}000/27 = 520$ triple-stress intervals each year.
- With three collisions in nine years across the entire system, we estimate that probability a triple-stress period leads to a collision is $p_c = 6 \times 10^{-4}$.

Sensitivity analyses

More interesting than these specific values are the general insights this model provides. We next investigate the sensitivity of $\tilde{p}$, the probability the signal is in the triple-stress state at a random instant, to changes to each of the a_i and b_i, the mean durations of stressed and normal intervals. For each of the three systems, we examine the absolute effect on $\tilde{p}$ of increasing the mean normal duration by 10%, and we examine the effect of decreasing the mean stressed duration by 10% (Figure 4-13). We conduct these sensitivity analyses numerically; we do not use the approximation of Equation 4.8. These analyses suggest the following points:

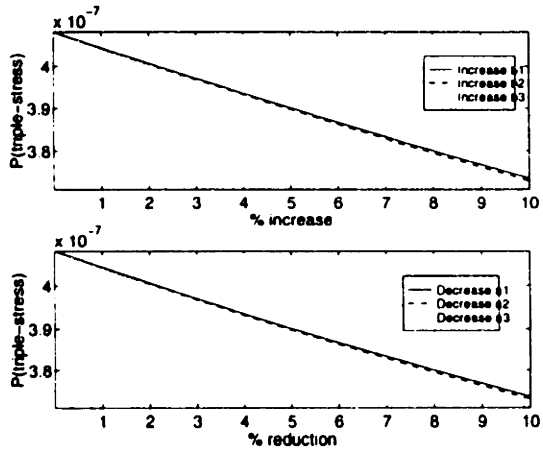

Figure 4-13: 10% improvements to the base-case triple-stress model parameters. These plots present the change in $\tilde{p}$ when one of the six parameters is changed. The three lines in the upper plot present the effect of increasing the mean duration in the normal state for each system. The three lines in the lower plot present the effect of decreasing the mean stress duration for each system.

- Increasing any one of the b_i reduces $\tilde{p}$. For small increases, the magnitude of this reduction is about the same whether it is b_1, b_2, or b_3 that is increased.

- Decreasing any one of the a_i reduces $\tilde{p}$. For small decreases, the magnitude of this reduction is about the same whether it is a_1, a_2, or a_3 that is decreased.

- A small percentage increase to one of the b_i or a small percentage decrease to one of the a_i yield similar reductions to $\tilde{p}$.

- A 10% increase in one of the b_i or a 10% decrease in one of the a_i yield slightly less than a 10% decrease in $\tilde{p}$.

- For larger percentage improvements, the largest reduction in $\tilde{p}$ is gained by applying the improvement to a_3 or b_3, this is, by improving the signaling system.

We compare the percentage reduction in $\tilde{p}$ obtained from 5% and 50% improvements to each of the six main parameters (Figure 4-14 and Figure 4-15). As before, we note that improving any one of the a_i or b_i yields a similar effect on $\tilde{p}$. The largest percentage improvements are to be had by improving a_3 or b_3. We note that a 5%

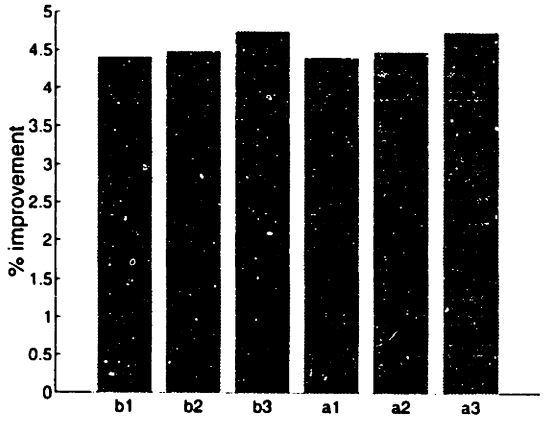

Figure 4-14: 5% improvements to base-case triple-stress model parameters. Each bar presents the reduction in $\tilde{p}$ obtained from a 5% improvement (increases for the b_i, decreases for the a_i) in each of the six parameters.

improvement to one of the parameters yields about a 4.5% improvement in $\tilde{p}$, whereas a 50% improvement to one of the parameters yields only a 32% improvement in $\tilde{p}$.

It is likely that the costs involved in reducing a_i or increasing b_i may vary greatly between systems. We lack data on the relative costs of such improvements. Nonetheless, we may compare the effectiveness of making simultaneous small improvements to each system against the effectiveness of making a single large improvement to one of the systems. If we improve each of the three systems by simultaneously reducing a_1, a_2, and a_3 by 10% each, we observe that $\tilde{p}$ drops by 26%. If we instead improve a_1 alone by 30%, we obtain a 28.3% reduction in $\tilde{p}$. If we improve a_2 alone by 30%, we obtain a 28.7% reduction in $\tilde{p}$. If we improve a_3 alone by 30%, we obtain a 29.9% reduction in $\tilde{p}$ (Figure 4-16). All costs being equal, this suggests that concentrating on improving the reliability of the most reliable system alone (for JR East, the signaling system, as $b_3/(a_3+b+3)$ is larger than $b_1/(a_1+b_1)$ and $b_2/(a_2+b_2)$) is more efficient than dividing an improvement amongst the three systems. We caution that this result depends critically on the assumption of equal costs for percentage improvements across the three systems, and may not hold for more realistic cost structures.

Our base-case parameters for the triple-stress model assumed $n_i = m_i = 1$, implying that the durations of normal intervals and stressed intervals take an exponential

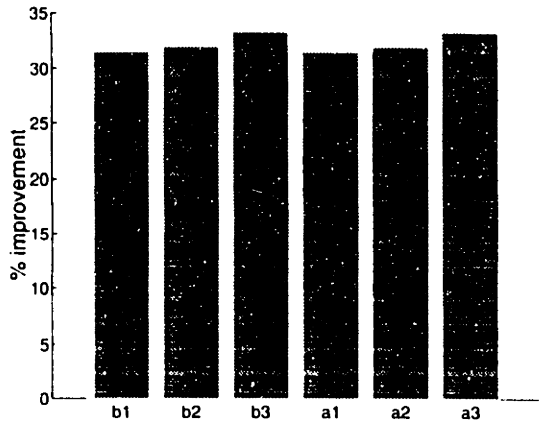

Figure 4-15: 50% improvements to base-case triple-stress model parameters. Each bar presents the reduction in $\tilde{p}$ obtained from a 50% improvement (increases for the b_i, decreases for the a_i) in each of the six parameters.

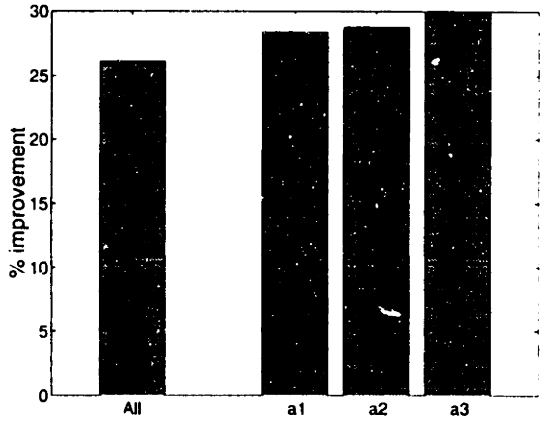

Figure 4-16: Dividing an improvement amongst the three systems. The bar labeled "All" indicates the percentage reduction in $\tilde{p}$ that follows from simultaneously reducing a_1, a_2, and a_3 by 10%. The bars labeled "a_1," "a_2," and "a_3" indicate the percentage reduction following a 30% reduction in only a_1, only a_2, and only a_3, respectively.

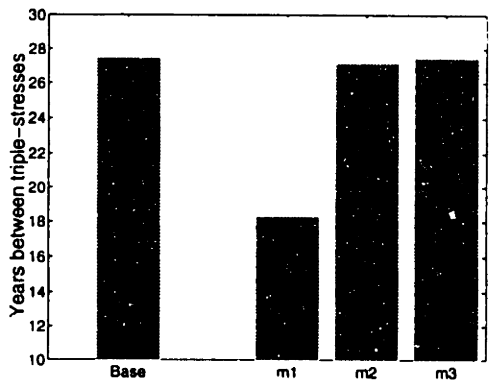

Figure 4-17: The contribution of variance in the triple-stress model. The bar labeled "Base" indicates the value of $6,500/(\check{a}+\tilde{b})$ using the base-case parameters. The bars labeled "m_1," "m_2," and "m_3" indicate the same quantity when m_1, m_2, and m_3 are increased from one to four, respectively.

distribution. Let us reduce the variance of the duration of a stressed interval for each of the three systems in turn. We shall reduce the coefficient of variation of these durations from 1.0 (the exponential distribution) to 0.5 (the fourth-order Erlang distribution) by increasing m_i from one to four. We hold all other parameters constant. From Equation 4.1, we see that changing the variability of the stressed and normal durations does not affect $\tilde{p}$. Such changes do affect $\check{a}$ and $\tilde{b}$, the mean durations of triple-stress and non-triple-stress intervals. We see that reducing the coefficient of variation from 1.0 to 0.5 decreases $6,500/(\check{a}+\tilde{b})$, the mean number of years between the occurrence of triple-stress intervals at a signal (Figure 4-17). For systems 2 and 3, this reduction is negligible: the base value of $6,500/(\check{a}+\tilde{b})$ drops from 27.5 years to 27.2 and 27.4 years, respectively. For system 1, the driver system, halving the coefficient of variation of the duration of a driver stress interval almost doubles the frequency of triple-stress intervals, reducing $6,500/(\check{a}+\tilde{b})$ to 18.3 years.

4.5 Conclusions

This chapter offers two approaches for investigating collision risk in the JR East system. The first approach is a statistical analysis of JR East collision and signal overrun

accidents. The second approach is an abstract model for collision risk, based on the idea that a necessary condition for a collision is the simultaneous occurrence of three types of "stress." We derive the following observations from these two approaches.

- The collision rate across the JR East system is low. JR East enjoys a collision rate at least three times lower than JNR collision rate across the same geographic region.

- The rate of signal overrun accidents appears to have fallen by 10% each year since privatization. Adopting signal overruns as a proxy measure of system safety, we might assume the collision rate similarly declined.

- Signal overruns occur one third as frequently on lines protected by ATS-SN as on lines protected by ATS-S/B. Signal overruns occur one eleventh as frequently on lines protected by ATS-P or ATC as on lines protected by ATS-S/B.

- The ongoing program of upgrades to the signaling systems is sufficient to explain the improvements in the signal overrun rate.

- The three JR East collisions to date do not suggest major safety flaws in the rail system. Based on Perrow's "normal accident" framework, we believe that future collisions will have antecedent causes that appear unique, complex, and accidental.

- To reduce the collision risk, the triple-stress model suggests reducing the mean duration of intervals when drivers are incapable of stopping the train, reducing the mean duration of intervals when the distance between trains is unusually small, and reducing the mean duration of intervals when the signals are not functioning properly. Alternatively, increasing the mean time between such stresses is equally effective in decreasing collision risk.

- The triple-stress model suggests that all of these measures are equally effective in percentage terms for small improvements. For larger improvements, the most benefit is to be gained by improving the most reliable system; for JR East, this

is the signaling system. This result supports the railroad's commitment to its costly but important program of upgrading its signaling systems.

Chapter 5

Conclusions

This thesis investigates three hazards facing JR East: seismic risk, slope failure risk, and collision risk. This concluding chapter addresses in preliminary terms the relative importance of these three hazards. We find the importance of the seismic threat relative to the two other hazards is highly sensitive to one's beliefs regarding the probability of a major Tokyo earthquake. This chapter also summarizes the principal findings of the thesis and offers suggestions for additional research.

5.1 Comparative risk analysis

Like any railroad, JR East faces many potential hazards. Some hazards occur frequently but involve typically minor consequences. Slope failure risk, discussed in Chapter 3, is an example. Other hazards occur less frequently but involve more severe consequences. Collision risk, discussed in Chapter 4, is such a case. Finally, some hazards are sufficiently rare that the railroad has yet to experience them, but could involve extremely severe consequences. Seismic risk from a major earthquakes, as described in Chapter 2, falls into this category. In this section, we discuss how JR East might determine the relative importance of various hazards that have differing consequences and occur on differing time scales. In this chapter, we restrict our attention to the three hazards discussed in this thesis. We propose to measure the negative effect that accidents of different types and sizes have upon the railroad

using a disutility function. Before presenting this approach, let us review the three hazards in qualitative terms.

From Chapter 3, we believe that rainfall-induced slope failures pose minor threat to the railroad. We estimate that the conventional train network experiences approximately 60 slope failures that damage the track each year. The Shinkansen network has experienced no rainfall-induced slope failures to date; due to its design, we believe that the Shinkansen network is essentially immune from such damage. In nine years of operation, the JR East conventional network has suffered only one derailment due to such track damage. No passengers were injured or killed in this derailment. Over the last nine years, the JR East conventional network has experienced a total of 32 operational derailment accidents from all causes; over this period, rainfall-induced slope failure has accounted for $1/32 = 3\%$ of all derailments. [RK95] As the current rainfall alert system slows (or stops) trains during heavy rain, it is likely that trains which would derail at rainfall-induced slope failures would do so while traveling at low speed. We speculate that the rainfall hazard may be a more important contributor to total rail service disruption than to total rail risk.

From Chapter 4, we believe that train collisions pose a larger threat to the railroad than rainfall-induced slope failures. Since 1987, the JR East conventional network has averaged between 10 and 20 signal overrun accidents each year; over this period, the railroad has suffered three collisions, all involving conventional trains. Two of the three collisions were minor, with no passenger injuries or fatalities. One of the three collisions was more serious, involving the death of one passenger and one crew member, and the injury of 116 passengers. We note that all three collisions occurred at reduced speed: in each case, the drivers of the moving trains realized a collision was imminent and were braking to an emergency stop when the collision occurred. The Shinkansen network has experienced no train collisions to date; due to its design, we believe that the Shinkansen network is effectively immune from such accidents.

While JR East has had direct experience with the rainfall hazard and the collision hazard, the railroad has never yet experienced a major earthquake. The 1995 Kobe quake and the 1923 Great Kanto quake did not involve JR East trains, but these

events provide two examples of rail damage in large earthquakes. The Kobe quake struck before the trains began their daily operation, and thus damaged no moving trains. The Kobe quake did topple several parked conventional trains from their tracks, and this quake damaged both the conventional and Shinkansen track in several places. The 1923 Great Kanto quake struck when the trains were operating. This quake derailed or overturned 108 rail cars, killing over 100 passengers. [SRL95h] As discussed in Chapter 2, a large earthquake which struck close to Tokyo could damage many JR East trains and endanger the lives of many JR East passengers.

There is considerable uncertainty regarding the probability of a Tokyo quake over the next year, or even over the next decade. Three tectonic plates collide off the coast of Odawara, about 80 km from Tokyo, and the junction of these plates has generated major earthquakes near Tokyo in 1633, 1703, 1782, 1853, and 1923. The mean interval between these disasters is 72.5 years, with a standard deviation of 4.3 years. As it has been 73 years since the 1923 Great Kanto quake, it could be argued that Tokyo is "due" for a major earthquake sometime in the next decade. [Eco93] Further, the Mihama volcano on the island of Oshima has erupted five to ten years before each of these earthquakes; Mihama erupted most recently in 1986, ten years ago. [Eco91] These observations suggest that the annual probability of a major Tokyo earthquake could be quite high at this time. In 1991, seismologist Tsuneji Rikitake of Nihon University—one of the six "wise men" entrusted with issuing civil defense warnings before a Tokyo earthquake—predicted that Tokyo had a one-in-five chance of being hit within the decade by a catastrophic earthquake of similar magnitude to the Great Kanto quake of 1923. [Eco91] However, earthquake prediction is a highly imprecise science, and such estimates are far from certain. [Bla95]

The challenge in assessing the relative importance of the rainfall hazard, the collision hazard, and the seismic hazard is comparing two relatively well-understood threats with one that is less so. It is clear that the relative importance of the seismic threat will depend on one's beliefs about the annual probability of a major Tokyo earthquake, and one's beliefs about the rail damage in such a disaster. In the absence of solid data regarding this probability and these consequences, a relative risk

assessment necessarily involves some subjective judgment.

5.1.1 Relative expected disutility

One approach to assess the relative importance of these three hazards is to use utility theory. [VM44] Such an approach could involve the following steps. Let $i = 1, 2, 3$ denote the rainfall, collision, and seismic hazards, respectively; and let $j = 1, 2$ denote the conventional and the Shinkansen networks, respectively.

1. For each type of hazard i and each network j, determine $E(A_{ij})$, the mean annual number of type i incidents across network j.

2. For each type of hazard i and each network j, determine $E(K_{ij})$, the mean number of derailing passenger trains in a type i incident in network j.

3. For each type of hazard i and each network j, determine the distribution of $\vec{D_{ij}}$, a vector of consequences associated with a derailment in a type i incident in network j. This vector of consequences could include passenger deaths and injuries, damages to rolling stock and infrastructure, and disruption and repair costs.

4. Choose a utility function, $\psi(\vec{d})$, that measures the total (dis)utility of a vector of consequences $\vec{d}$. As accident consequences have negative utility, $\psi(\vec{d}) < 0$.

5. For each hazard i, compute its expected annual disutility, U_i, summing across both conventional and Shinkansen networks.

$$U_i = \sum_{j=1}^{2} E(A_{ij})\ E(K_{ij})\ E(\psi(D_{ij}))$$

6. For each hazard i, compute its contribution to total annual expected disutility,

$$\mathcal{F}_i = U_i / \sum_{s=1}^{3} U_s.$$

We next present a simplified example of such an analysis. Using results from earlier chapters of this thesis, we can estimate the $E(A_{ij})$ and the $E(K_{ij})$. However, this thesis has not addressed the consequences of different types of derailments, such as injury rates, death rates, damage costs, and so on. Evaluating the disutility of these various consequences to JR East is also beyond the scope of this work. Steps 1 and 2 of the previous list are thus easily accomplished, while steps 3 and 4 go beyond the scope of this thesis. (Steps 5 and 6 require only simple calculation.) To demonstrate this relative expected disutility approach, we shall make a strong simplifying assumption for steps 3 and 4. We shall assume that

$$E(\psi(D_{ij})) = 1, \text{ for all } i \text{ and } j.$$

In other words, we shall assume that all derailments of passenger trains cause equal disutility to the railroad, regardless of whether the derailment involves conventional or Shinkansen trains, and regardless of the cause of the derailment. This assumption is clearly unrealistic. For example, we believe a Shinkansen which derails at high speed in an earthquake will have far more severe consequences than a conventional train which derails at low speed at a slope failure. We make this assumption only to permit the demonstration of the expected disutility approach. Because of this questionable assumption, the results of the following calculation should not be taken too literally.

Derailments from rainfall-induced slope failures

Due to its track design, we assume the Shinkansen system is effectively immune from slope failures, and thus $E(A_{12}) = 0$. (Recall $E(A_{ij})$ denotes the annual number of type i incidents in network j, where $i = 1, 2, 3$ denotes rainfall, collisions, and earthquakes, respectively, and $j = 1, 2$ denotes the conventional and the Shinkansen networks, respectively.) Extrapolating from the Suigun line, we estimate that the JR East conventional system experiences an average of $E(A_{11}) = 60$ serious slope failure incidents per year. We estimate the probability that a slope failure leads to a

derailment of a conventional train is 1/500, and we assume the probability of multiple derailments at the same failure is negligible (Section 3.7). Thus, $E(K_{11}) = 1/500$. With these values, the expected annual number of derailments due to slope failure incidents is $U_1 = E(A_{11})E(K_{11}) + E(A_{12})E(K_{12}) = 0.12 + 0 = 0.12$.

Derailments from train collisions

As a first-order approximation, we assume that the Shinkansen system is effectively immune from collisions. due to its advanced signaling system and long headways between trains. Accordingly, we set $E(A_{22}) = 0$. Using the most recent estimate of the signal overrun rate for the conventional system, $\hat{\delta}_{1995} = 7.5 \times 10^{-8}$ overruns per train-km, and the most recent figure for annual system-wide conventional train volume, $v_{1995} = 2.7 \times 10^8$ train-km, we estimate the JR East conventional system currently experiences 20 signal overrun incidents per year. We assume the probability a signal overrun progresses to a collision is 1/65, and thus $E(A_{21}) = 20/65 = 0.31$ expected collisions per year (Section 4.3.2). The 10/19/88 collision involved two freight trains; the 12/5/88 collision involved two passenger trains; and the 6/28/92 collision involved one freight and one passenger train. The average number of derailing passenger trains in a collision to date has been 1.5, and thus we set $E(K_{21}) = 1.5$. With these values, the expected annual number of derailments due to collisions is $U_2 = E(A_{21})E(K_{21}) + E(A_{22})E(K_{22}) = 0.47 + 0 = 0.47$.

Derailments from major earthquakes

JR East trains travel across much of northern Honshu, yet the bulk of the railroad's operations occur close to Tokyo (Section 2.5). For this preliminary case-study, we will only consider major earthquakes which strike close to Tokyo. As previously discussed, there is considerable uncertainty regarding the likelihood of a Tokyo quake. For this example, we shall adopt Rikitake's 1991 estimate of a 20% probability of a major Tokyo earthquake over the next decade at face value (without adjusting for the fact it has been five years since his prediction). Let p denote the annual probability of a major Tokyo earthquake. Assuming each of the ten years constitutes

an independent Bernoulli trial, we translate Rikitake's prediction into an equation, $1-(1-p)^{10}=0.2$, and calculate $p=0.022$. That is, accepting Rikitake's prediction, the expected annual number of major Tokyo earthquakes is $E(A_{31})=E(A_{32})=0.022$. As JR East trains operate 18 hours of each day, the probability is 18/24 that such a quake strikes when the trains are operating. Given such a major Tokyo quake, we assume all trains within 40km of Tokyo face a risk of derailment. From Section 2.4, we assume that if a major Tokyo earthquake strikes when the trains are operating, 230 conventional trains each face of 0.06 probability of derailment, and three Shinkansen trains each face a 0.70 probability of derailment. With these values, the expected number of conventional train derailments in a major Tokyo earthquake is $E(K_{31})=(18/24)(0.06)(230)=10.4$, and the expected number of Shinkansen derailments in a major Tokyo earthquake is $E(K_{32})=(18/24)(0.70)(3)=1.6$. Accordingly, the expected annual number of derailments due to major Tokyo earthquakes is $U_3=E(A_{31})E(K_{31})+E(A_{32})E(K_{32})=(0.022)(10.4+1.6)=0.26$.

Relative expected disutility

We next calculate the relative expected disutility associated with each of the three hazards.

$$\begin{aligned}
&\text{Slope failures:} && \mathcal{F}_1 = \frac{U_1}{U_1+U_2+U_3} = 14\% \\
&\text{Collisions:} && \mathcal{F}_2 = \frac{U_2}{U_1+U_2+U_3} = 55\% \\
&\text{Earthquakes:} && \mathcal{F}_3 = \frac{U_3}{U_1+U_2+U_3} = 31\%
\end{aligned}$$

We find that with this model and these parameter values, collisions comprise the largest fraction of total expected disutility. We make two comments. First, we note that this result is quite sensitive to one's assumption regarding the probability of a major Tokyo quake over the next decade (Figure 5-1). Once this probability exceeds 1/3, the Tokyo earthquake scenario comprises the largest fraction of total expected disutility. Second, and more critically, we note this result is based on an unrealistic disutility function. We speculate that derailments from earthquakes would occur at higher average speeds than derailments from collisions and slope failures, and thus have far more severe consequences than derailments from collisions and slope failures.

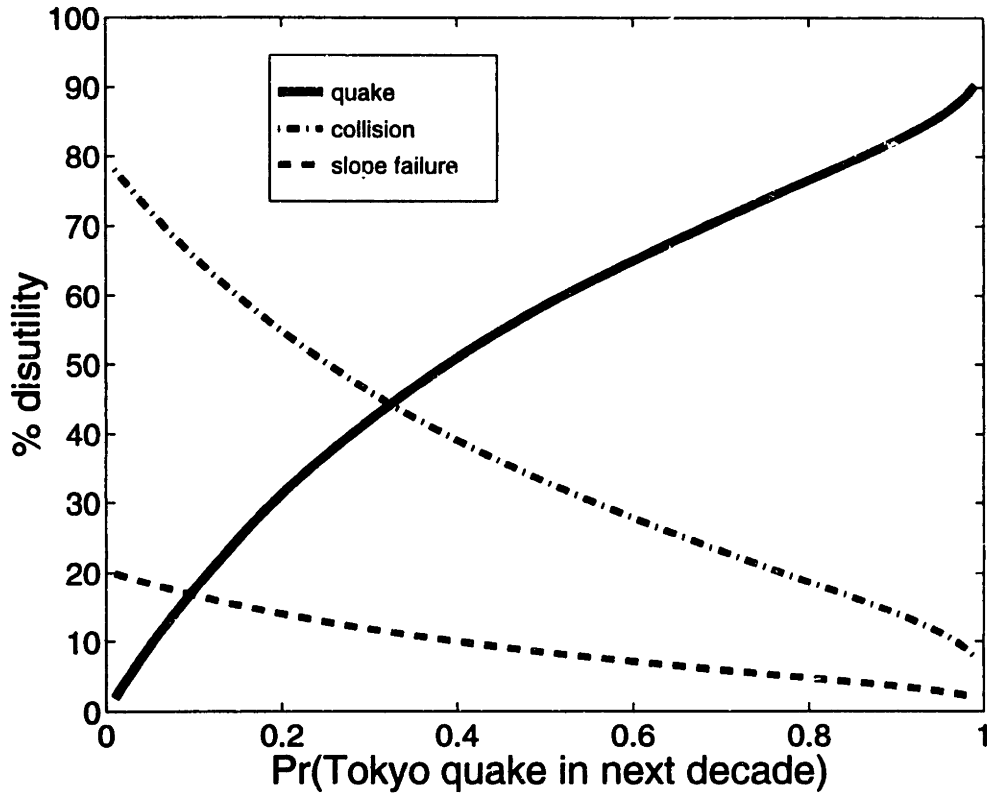

Figure 5-1: Fraction of expected disutility for each of the three hazards, as a function of the probability of a major Tokyo quake in the next decade. This preliminary case-study assumes all types of derailments have equal disutility.

We believe this especially true for Shinkansen derailments. To date, the worst collision accident on the conventional network resulted in one passenger death. To date, the worst rainfall-induced derailment accident on the conventional network resulted in no passenger deaths. A Shinkansen train may carry 1,000 passengers while traveling 240 km/h on an elevated viaduct, 10 meters above the ground. We speculate that should a moving Shinkansen derail in an earthquake, many passengers could face serious injury or death.

To go beyond this preliminary case-study, a more realistic characterization of derailment consequences by type and network (the $\vec{D_{ij}}$) and their associated disutilities (the function $\psi(\vec{d})$) is required. Reasonable damage metrics could include the number of passenger injuries, the number of passenger fatalities, the cost of destroyed rolling stock, the cost of destroyed infrastructure, and the cost of service disruption. JR East

could construct a disutility function for these consequences that expresses the railroad's values and priorities. Several different disutility functions could be explored, allowing the railroad to assess the sensitivity of these calculations to this factor. As in our case-study, the railroad could use these $\vec{D}_{ij}$ estimates and this $\psi(\vec{d})$ function to calculate the relative contribution of each category of hazard to total expected disutility. We speculate that such a calculation may reveal that the Tokyo earthquake scenario poses the greatest threat to JR East.

5.2 Contributions of thesis

This thesis has addressed three safety hazards faced by JR East. We have developed a collection of stochastic models to assess the risk of these hazards. Our modeling style is driven by data availability. In situations with ample historical data, such as slope failures and signal overruns, we opt to draw inferences from the data themselves as much as possible. In situations where historical data are lacking, such as catastrophic earthquakes, we construct models from "first principles," utilizing physics, geometry, and probability theory. We use sensitivity analysis to determine critical parameters; and we suggest how our models might inform policy decisions.

In Chapter 2, we consider the number and speed of conventional and high-speed JR East trains that might crash at track failures in the aftermath of a catastrophic earthquake. Even though our models omit shaking risk, we still find that trains caught within the damage region of a large quake face a significant probability of crashing at a quake-induced track failure. For conventional trains inside the region of quake damage, we estimate this probability is 0.06 per train per earthquake; for Shinkansen trains within the region of quake damage, we estimate this probability is 0.75 per train per earthquake. We find that most of the various refinements to our crash risk models—tunnels, station stops, acceleration and deceleration profiles, and so forth—produce only small changes in these values. We suggest the Sanyo Shinkansen was "doubly fortunate" in the Kobe quake: not only did the Kobe tunnel shield 40 km of track from damage, but the quake struck before the trains began

service for the day. Even though Shinkansen trains face a much larger per-train risk of crashing than conventional trains if caught in the damage region of an earthquake, many more conventional trains face this risk. Due to this difference in train density, our analysis of the JR East conventional and high-speed networks suggests that the conventional system would suffer more crashing trains than the Shinkansen system. Our crash speed model suggests that trains which crash while braking tend to crash at high impact speeds, averaging more than 66% of full cruising speed.

In Chapter 3, we analyze the current and alternative slope failure alert rules. Our analysis suggests that the current rainfall alert system performed quite well along the Suigun line between 1991 to 1994, at least in comparison to the suite of alternative rules we investigated. We believe that overlapping gauge control regions may be an important factor in this performance. Within the class of single exponential rules we considered, we find that rules using short, one- or two-hour half-lives offer superior performance to those rules using longer half-lives. We advocate the use of "complementary cumulative plots" as a comprehensive presentation of the full safety vs. delay spectrum of any rainfall index. We note that implicit in a choice of a rainfall rule is a trade-off between safety and delay; by modeling the rate of slope failures as a function of a rainfall index, we may make this tradeoff explicit. Analysis of historical derailments indicates that even the most serious derailments to date have had typically small consequences; extrapolating from the Suigun line to the entire network, we estimate that probability a slope failure causes a derailment is less than one in five hundred.

In Chapter 4, we investigate the risk of train collisions using two methods. First, we analyze the signal overrun and collision data directly. We find the rate of signal overruns has fallen 10% annually since privatization. Assuming a proportionality between overruns and collisions, we conclude the collision rate also declined similarly. We suggest that the ongoing program of upgrades to the signaling systems may explain this improvement. Second, we introduce a conceptual "triple-stress" model to study collision risk. Drawing upon Charles Perrow's normal accident framework, we believe that future collisions will have antecedent causes that appear unique, complex,

and accidental. This triple-stress model suggests that reducing the mean duration of intervals when drivers are incapable of stopping the train, reducing the mean duration of intervals when the distance between trains is unusually small, and reducing the mean duration of intervals when the signals are not functioning properly are various approaches to reduce the likelihood of a collision. The triple-stress model also indicates that large improvements are most beneficial when applied to the most reliable system; this conclusion supports JR East's commitment to upgrading its signaling systems.

5.3 Directions for additional research

Regarding the seismic hazard and the collision hazard, we believe this thesis takes these problems about as far as one may comfortably go within the operations research paradigm. We believe JR East would benefit from additional detailed seismology and railroad engineering studies in these areas. In the area of earthquake risk, we suggest that the railroad may benefit from additional research on advanced braking systems. improved early warning systems, seismic vulnerabilities of the infrastructure across the network, infrastructure-strengthening programs, and improved crash survivability of rolling stock. In the area of collision risk, we suggest the railroad may benefit from additional research on signaling system upgrades, with specific attention to quantifying the costs vs. the benefits of such improvements.

Regarding the slope failure hazard, a number of interesting statistical problems remain for study. Though not demonstrated in this thesis, we speculate that a better rainfall risk model could offer the railroad superior safety and delay performance using the hardware now in place. Such improvements would only require modifying software in the train control centers in Tokyo; beyond the costs of the research itself and such programming changes, these improvements would entail little expense. An improved rainfall risk model might consider local variations in topology, geology, and vegetation, as well as the long-term weathering processes involved in slope stability.

Another intellectually interesting problem concerns system recovery from acci-

dents and delays. Currently, when the JR East rail network gets off schedule, teams of skilled managers in the train control centers make real-time tactical decisions to return operations to normal as quickly as possible. To do this, the managers order the dispatchers to cancel certain trains, hold others, run certain local trains as express trains, and so forth. We envision that a real-time decision-support system could aid these managers. Such a decision-support system might be based on an optimization model; such a system might assist managers quickly evaluate proposed recovery plans before ordering the dispatchers to implement them.

Finally, we believe the railroad may benefit from a unified risk management framework. Such a framework would provide a consistent and rational way to assess the costs and benefits of various risk mitigation options. As a corporation which spends large sums of money each year on a wide range of risk-reduction projects, JR East has an effective but unconscious version of such a framework. We suggest that the very process of formalizing a unified framework for safety management would provide the railroad with valuable insights into the corporation's own values and beliefs about risk. Such a risk management framework could help the railroad make cost-effective decisions about safety investments in a manner most consistent with corporate values.

Bibliography

[Arn90] Steven Arnold. *Mathematical Statistics.* Prentice-Hall, Englewood Cliffs, New Jersey, 1990.

[B+93] Kathleen Bradley et al. *National Transportation Statistics Annual Review: Historical compendium, 1960–1992.* U.S. DOT Research and Special Programs Administration, Washington, DC, September 1993.

[BAC90] A. M. Birk, R. J. Anderson, and A. J. Coppens. A computer simulation of a derailment accident. *Journal of Hazardous Materials*, 25:121–165, 1990.

[BC88] Rajan Batta and Samuel Chiu. Optimal obnoxious paths on a network: transportation of hazardous materials. *Operations Research*, 36(1):84–92, 1988.

[Beg95] Sharon Begley. Lessons of Kobe. *Newsweek*, pages 20–30, January 30 1995.

[Ben95] Fouad Bendimerad. The Great Hanshin Earthquake. Technical report, Risk Management Solutions and Failure Analysis Associates, Menlo Park, California, January 1995.

[Bie93] Vicki Bier. Statistical methods for the use of accident precursor data in estimating the frequency of rare events. *Reliability Engineering and System Safety*, 41:267–280, 1993.

[Bin93] Alan J. Bing. Collision avoidance and accident survivability: Collision threat. Final report DOT/FRA/ORD-93/02.I, Arthur D. Little, March 1993.

[BKJ92] Dennis Bley, Stanley Kaplan, and David Johnson. The strengths and limitations of PSA: where we stand. *Reliability Engineering and System Safety*, 38:3-26, 1992.

[Bla95] Sandra Blakeslee. Hopes for predicting earthquakes, once so bright, are growing dim. The New York Times, August 8 1995. p. C2.

[BM87] Nelson Blachman and Robert Machol. Confidence intervals based on one or more observations. *IEEE Transactions on Information Theory*, 33(3):373-382, May 1987.

[Box92] George Box. *Bayesian inference in statistical analysis.* Wiley, 1992.

[BSO95] Hans Bohnenblust, Joseph Sussman, and Amedeo Odoni. Perspectives of a safety management plan. Technical report, Massachusetts Institute of Technology, Cambridge, Massachusetts, June 1995.

[Cai80] Nel Caine. The rainfall intensity-duration control of shallow landslides and debris flows. *Geografiska Annaler*, 62A:23-27, 1980.

[Car89] A. Carrara. Landslide hazard mapping by statistical methods: a "black box" model approach. Technical report, CNR-CIOC Faculty of Engineering, Bologna, Italy, 1989.

[CE95] Susan Cannon and Stephen Ellen. Rainfall conditions for abundant debris avalanches in the San Francisco Bay region, California. *California Geology*, pages 267-272, December 1995.

[CO87] Tat Chi Chow and Robert Oliver. How many reactor accidents will there be? *Nature*, 327:20-21, May 7 1987.

[COV90] Tat Chi Chow, Robert Oliver, and G. Anthony Vignaux. A Bayesian escalation model to predict nuclear accidents and risk. *Operations Research*, 38(2):265-277, March-April 1990.

[Dra88] Alvin Drake. *Fundamentals of Applied Probability Theory.* McGraw Hill, New York, 1988.

[Eas93] JR East. Annual report, 1993.

[Eas95] JR East. Shinkansen rail system timetable, 1995.

[Eco91] The Economist. Waiting for the big one, December 7 1991. pages 101–103.

[Eco93] The Economist. A fishy tail, September 4 1993. page 37.

[Eco95] The Economist. Fear of trembling: a survey of earthquake engineering, April 22 1995. pages 3–12.

[Edw86] A. Edwards. How many reactor accidents? *Nature*, 324:417–418, December 4 1986.

[Fai76] William Fairley. Evaluting the "small" probability of a catastrophic accident from the marine transportation of liquified natural gas. In William Fairley and Frederick Mosteller, editors, *Statistics and Public Policy*, pages 331–353. Addison-Wesley, Reading, Massachusetts, 1976.

[Gal93] Robert Galganski. Collision avoidance and accident survivability: Accident survivability. Final report DOT/FRA/ORD-93/02.III, Arthur D. Little, March 1993.

[Gal96a] Robert Gallager. *Discrete Stochastic Processes.* Kluwer Academic Publishers, Boston, 1996.

[Gal96b] Kevin Galvin. Twelve die as trains collide in Maryland. *The Boston Globe*, page 1, February 17 1996.

[GHM94] Martha Grabowski, John Harrald, and Thomas Mazzuchi. Risk modeling for rare events: the development and use of a vessel oil spill model for the Puget Sound and the Straits of Juan de Fuca. Technical report, The George Washington University, November 1994.

[Gla96] Malcolm Gladwell. Blowup. *The New Yorker*, pages 32–36, January 22 1996.

[GR84] Theodore Glickman and Donald Rosenfield. Risks of catastrophic derailment involving the release of hazardous materials. *Management Science*, 30(4):503–511, 1984.

[H+93] John Harrison et al. Collision avoidance and accident survivability: Collision avoidance and safety systems. Final report DOT/FRA/ORD-93/02.II, Arthur D. Little, March 1993.

[Ham93] Lawrence Hamilton. *Statistics with STATA 3.* Duxbury Press, Belmont, California, 1993.

[Han84] A. Hansen. Landslide hazard analysis. In Denys Brunsden and David Prior, editors, *Slope Instability*, Landscape Systems: a series in geomorphology, chapter 13, pages 523–603. John Wiley and Sons, New York, 1984.

[Heb96] H. J. Hebert. Flames killed many victims of train crash. *The Boston Globe*, page 3, February 20 1996.

[HR74] David Halliday and Robert Resnick. *Fundamentals of Physics.* Wiley and sons, New York, second edition, 1974.

[ID87] R. M. Iverson and R. P. Denlinger. The phsyics of debris flows: a conceptual assessment. In *Erosion and Sedimentation in the Pacific Rim*, IAHS publication 165, pages 155–165, Corvallis Symposium, August 1987.

[IL86] S. Islam and K. Lindgren. How many reactor accidents will there be? *Nature*, 322:691–692, August 21 1986.

[Ing94a] Ármann Ingólfsson. *Earthquake forecasts: the life-saving potential of last-minute warnings.* PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, September 1994. Doctorate in Operations Research.

[Ing94b] Ármann Ingólfsson. Estimating the probability of rare events. Technical report, Massachusetts Institute of Technology, Cambridge, Massachusetts, September 1994.

[Ing94c] Ármann Ingólfsson. Exploratory analysis of level crossing accident patterns. Technical report, Massachusetts Institute of Technology, Cambridge, Massachusetts, September 1994.

[IRJ95] International Railway Journal. Engineers work around the clock to restore earthquake-damaged railways, March 1995.

[JEN95] Japan Economic Newswire. Buckled shinkansen pillars contained pieces of wood, January 24 1995.

[JICA93] Japan International Cooperation Agency. Textbook for the group training course in railway signal, telecommunications, and information system engineering, January 1993. TITC-JR-93-35.

[Joh72] Chalmers Johnson. *Conspiracy at Matsukawa.* University of California Press, Berkeley, 1972.

[K+87] David Keefer et al. Real-time landslide warning during heavy rainfall. *Science*, 23:921–925, 13 November 1987.

[KAT84] Kazuhiko Kawashima, Koh Aizawa, and Kuzuyuki Takahasi. Attentuation of peak ground motion and absolute acceleration response spectra. In *Eighth World Conference on earthquake engineering*, volume 2, pages 257–263, San Francisco, California, 1984.

[KMBS82] Sungwoo Kim, Ronald Mauri, Joesph Blair, and David Skinner. A regression analysis on frequencies of rail freight accidents in the United States. *The Logistics and Transportation Review*, 18(2):99–116, 1982.

[Knu94] Mike Knutton. In a derailment, the TGV triumphs. *Railway Age*, page 26, February 1994.

[Kri95] Nicholas Kristof. Death toll in Japanese quake nears 2,000: Hundreds trapped in ruins as fires rage. *The New York Times*, page 1, January 18 1995.

[KS87] Sumiji Kobasi and Masakazu Suzuki. The critical rainfall (danger index) for disasters caused by debris flows and slope failures. In *Erosion and Sedimentation in the Pacific Rim*, IAHS publication 165, pages 201–211, Corvallis Symposium, August 1987

[Kun72a] Masaharu Kunieda. Several problems about rolling stock which can run at curves at high speed. *Reports of the Railway Technical Research Institute, Japanese National Railways*, 13(639), 1972.

[Kun72b] Masaharu Kunieda. Theoretical study on the mechanics of overturn of railway rolling stock. *Reports of the Railway Technical Research Institute, Japanese National Railways*, 13(651), 1972.

[Lue89] David Luenberger. *Linear and nonlinear programming.* Addison-Welsey, Reading, Massachusetts, 1989.

[MA95] Carl Martland and Sudhir Anandarao. Risk assessment of level crossings in the JR East system. Technical report, Massachusetts Institute of Technology, Cambridge, Massachusetts, June 1995.

[MR66] Robert Machol and Jonathon Rosenblatt. Confidence interval based on a single observation. *Proceeings of the IEEE*, 54(8):1087–1088, 1966.

[Nas95] Thomas Olivier Nasser. A framework to monitor the safety performance of transportation systems. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, February 1995. Masters of Science in Transportation.

[OS94] Katsuya Okada and Tomoyasu Sugiyama. A risk estimation method of railway embankment collapse due to heavy rainfall. *Structural Safety*, 14:131–150, 1994.

[Pap95] Achilleas Papadimitriou. Optimization of the seismic early warning system for the Tohoku Shinkansen. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, July 1995. Masters of Science in Civil and Environmental Engineering.

[Per84] Charles Perrow. *Normal accidents: living with high-risk technologies.* Harper-Collins, 1984.

[PV95] Achilleas Papadimitriou and Daniele Veneziano. Optimization of the seismic early warning system for the Tohoku Shinkansen. Research report, Massachusetts Institute of Technology, Cambridge, Massachusetts, August 1995.

[R+75] N. C. Rasmussen et al. Reactor safety study. Technical Report Wash-1400 (NUREG 75/014), U.S. Nuclear Regulatory Commission, Washington, DC, 1975.

[RK95] Alan Rimm-Kaufman. A statistical analysis of JRE non-collision derailments. Technical report, Massachusetts Institute of Technology, Cambridge, Massachusetts, July 1995.

[Ros89] Sheldon Ross. *Introduction to Probability Models.* Academic Press, San Diego, California. fourth edition, 1989.

[SE92] Rom Singh and John Elkins. Derailment preventation. *Railway Age*, pages 53–76, April 1992.

[Shi93] Heki Shibata. On fundamental approach to probabilistic risk assessment of railway passengers in accident. In *Proceedings of the International Conference on Speedup Technology for Railway and Mag-Lev Vehicles*, pages 562–568, Yokoyama, Japan, November 1993.

[Shi94] Heki Shibata. The conceptual flow of PSA on railway train crashing. Technical report, Yokohama National University, Yokohama, Japan, 1994.

[SK95] Jonathan Sapsford and Masayoshi Kanabayashi. Leader of cult is hunted in Tokyo subway poisonings. *Wall Street Journal*, page A8, March 23 1995.

[Soc74] Society for Industrial and Applied Mathematics. *Conference on Reliability and Fault Tree Analysis: theoretical and applied aspects of system reliability and safety assessment*, Berkeley, California, 1974. University of California Press.

[SRL95a] JR East Safety Research Laboratory. 1976–1987 JNR derailment accident database, 1995. Unpublished.

[SRL95b] JR East Safety Research Laboratory. 1987–1995 JRE derailment accident database. 1995. Unpublished.

[SRL95c] JR East Safety Research Laboratory. 1987–1995 JRE signal overrun database, 1995. Unpublished.

[SRL95d] JR East Safety Research Laboratory. 1991–1994 JRE five-minute precipitation database for the Suigun line, 1995. Unpublished.

[SRL95e] JR East Safety Research Laboratory. Derailment accident reports, 1995. Unpublished, in Japanese. Translated by Nobuyuki Yokoyama.

[SRL95f] JR East Safety Research Laboratory. JRE rainfall alert database for the Suigun line, 1995. Unpublished.

[SRL95g] JR East Safety Research Laboratory. JRE slope failure database for the Suigun line, 1995. Unpublished.

[SRL95h] JR East Safety Research Laboratory. List of railway earthquake disasters, 1995. from "Railway and Earthquakes" and "Railway Civil Engineering," distributed by the Railway Facilities Association.

[SRL95i] JR East Safety Research Laboratory. Notes on rainfall data analysis. Unpublished, 1995.

[SRL95j] JR East Safety Research Laboratory. Safety investment from 1989–1993. Unpublished spreadsheet, 1995.

[SRL95k] JR East Safety Research Laboratory. System profile database, 1995. Unpublished.

[SRL95l] JR East Safety Research Laboratory. Videotape of rail damage in the Kobe earthquake, 1995. Unpublished.

[SRL96a] JR East Safety Research Laboratory. Personal communication, Makoto Shimamura, Chief Researcher, 1996.

[SRL96b] JR East Safety Research Laboratory. Personal communication, Masahiko Horiuchi, Chief Researcher, 1996.

[SS82] Allan Sampson and Robert Smith. Assessing risks through the determination of rare event probabilities. *Operations Research*, 30(5):839–866, September–October 1982.

[SSB90] Anthony Starfield, Karl Smith, and Andrew Bleloch. *How to model it: Problem solving for the computer age.* McGraw-Hill Publishing Company, 1990.

[STA92] STATA Corp., Santa Monica, California. *STATA Reference Manual: Release 3*, fifth edition, March 1992.

[SV94] B. Sirangelo and P. Versace. A real time forecasting model for landslides triggered by rainfall. In *Workshop on Climate Change and Hydrological hazards in the Mediterranean area*, pages 1–32, Guigno, Italy, 1994.

[Swa90] Morris Swartz. An improved method for setting upper limits with small numbers of events. *Nuclear instruments and methods in physics research.* A294:278–281, 1990.

[Swo87] Cary Swoveland. Risk analysis of regulatory options for the transport of dangerous commodities by rail. *Interfaces*, 17(4):90–107, 1987.

[TAP96] The Associated Press. Maryland crash prompts rules for better passenger safety. *The Boston Globe*, page 6, February 21 1996.

[USD93] USDOT Bureau of Transportation Statistics, Washington, DC. *National Transportation Statistics Annual Report*, September 1993.

[USGS95] United States Geological Survey. Reducing landslide hazards in the United States. USGS Fact Sheet 112-95, 1995.

[Uts74] Tokuji Utsu. Space-time pattern of large earthquakes occurring off the pacific coast of the Japanese islands. *J. Phys. Earth*, 22:325–342, 1974.

[Van81] Mark Vangel. Poisson regression models for estimating the effect of tidal stress on earthquake occurrences. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, August 1981.

[VM44] J. Von Nuemann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, Princeton, New Jersey, 1944.

[Wal96] Matthew Wald. Investigators citing design flaw and engineer error in accident. *The New York Times*, page 1, February 19 1996.

[WW95] Raymond Wilson and Gerald Wieczorek. Rainfall thresholds for the initiation of debris flows at La Honda, California. *Environmental and Engineering Geoscience*, 1(1):11–27, Spring 1995.