Field, E. H., Milner, K. R., Hardebeck, J. L., Page, M. T., van der Elst, N., Jordan, T. H., Michael, A. J., Shaw, B. E., & Werner, M. (2017). A Spatiotemporal Clustering Model for the Third Uniform California Earthquake Rupture Forecast (UCERF3-ETAS): Toward an Operational Earthquake Forecast. *Bulletin of the Seismological Society of America*, 107(3), 1049-1081.
<https://doi.org/10.1785/0120160173>

Peer reviewed version

Link to published version (if available):
[10.1785/0120160173](https://doi.org/10.1785/0120160173)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via SSA at <http://bssa.geoscienceworld.org/content/early/2017/02/24/0120160173>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

A Spatiotemporal Clustering Model for the Third Uniform California Earthquake Rupture Forecast (UCERF3-ETAS): Toward an Operational Earthquake Forecast

by Edward H. Field, Kevin R. Milner, Jeanne L. Hardebeck, Morgan T. Page, Nicholas van der Elst, Thomas H. Jordan, Andrew J. Michael, Bruce E. Shaw, and Maximilian J. Werner

Abstract We, the ongoing Working Group on California Earthquake Probabilities, present a spatiotemporal clustering model for the Third Uniform California Earthquake Rupture Forecast (UCERF3), with the goal being to represent aftershocks, induced seismicity, and otherwise triggered events as a potential basis for operational earthquake forecasting (OEF). Specifically, we add an epidemic-type aftershock sequence (ETAS) component to the previously published time-independent and long-term time-dependent forecasts. This combined model, referred to as UCERF3-ETAS, collectively represents a relaxation of segmentation assumptions, the inclusion of multifault ruptures, an elastic-rebound model for fault-based ruptures, and a state-of-the-art spatiotemporal clustering component. It also represents an attempt to merge fault-based forecasts with statistical seismology models, such that information on fault proximity, activity rate, and time since last event are considered in OEF. We describe several unanticipated challenges that were encountered, including a need for elastic rebound and characteristic magnitude–frequency distributions (MFDs) on faults, both of which are required to get realistic triggering behavior. UCERF3-ETAS produces synthetic catalogs of $M \geq 2.5$ events, conditioned on any prior $M \geq 2.5$ events that are input to the model. We evaluate results with respect to both long-term (1000 year) simulations as well as for 10-year time periods following a variety of hypothetical scenario mainshocks. Although the results are very plausible, they are not always consistent with the simple notion that triggering probabilities should be greater if a mainshock is located near a fault. Important factors include whether the MFD near faults includes a significant characteristic earthquake component, as well as whether large triggered events can nucleate from within the rupture zone of the mainshock. Because UCERF3-ETAS has many sources of uncertainty, as will any subsequent version or competing model, potential usefulness needs to be considered in the context of actual applications.

Electronic Supplement: Figures showing discretization, verification of the *DistanceDecayCubeSampler*, average simulated participation rate, and average cumulative MFD.

Introduction

Long-term earthquake forecasts, which are applicable from decades to centuries, represent our first line of defense with respect to mitigating earthquake risk, especially in terms of informing building codes. We also know, however, that aftershocks and otherwise triggered events can be large and damaging, as demonstrated by several earthquakes including the 2011 M 6.3 Christchurch, New Zealand, event (e.g., [Kaiser et al., 2012](#)). The ability to forecast earthquakes on

shorter time scales, such as days to decades, is known as operational earthquake forecasting (OEF), which also involves the dissemination of authoritative information to inform risk mitigation decisions ([Jordan et al., 2011](#)). The history, motivation, and challenges associated with OEF have been discussed elsewhere ([Jordan and Jones, 2010](#); [Jordan et al., 2011, 2014](#); [Peresan et al., 2012](#); [Wang and Rogers, 2014](#)), and significant progress with OEF has been made in both

Table 1
Some of the Informational Constraints (Observed or Model Inferred) that One Might Use to Forecast the Probability of Triggered Earthquakes

Observation or Model Inference	STEP	UCERF3-ETAS
Global or regional triggering statistics (generic aftershock parameters)	✓	✓
Sequence-specific deviations from generic aftershock parameters	✓	✓
Spatial and temporal variation in aftershock parameters (within a sequence)	✓	✓
Spatial variation in long-term event rates (background rates)	✓	✓
Location of recent events (i.e., areas that are active with microseismicity)	✓	✓
Proximity to active faults, especially with respect to triggering larger magnitude events		✓
The long-term magnitude–frequency distribution inferred for faults		✓
Elastic-rebound implied stress on faults (e.g., time since last event relative to recurrence interval)		✓
Dynamic or static stress changes imposed by any previous events	Implicit?	Implicit?

Also indicated are the constraints used by the short-term earthquake probability (STEP) model (Gerstenberger *et al.*, 2005) and by the Third Uniform California Earthquake Rupture Forecast (UCERF3) epidemic-type aftershock sequence (ETAS). A checkmark means that the constraint is used explicitly with respect to forecasting nonspontaneous (triggered) events, whereas Implicit? means that the influence may be captured by other utilized constraints (e.g., static or dynamic stress changes being manifested in observed seismicity).

Italy following the 2009 L’Aquila earthquake (Marzocchi *et al.*, 2014) and in New Zealand following the 2010 Darfield/Christchurch events (Gerstenberger *et al.*, 2014). Furthermore, a workshop was recently held to discuss the potential usefulness of OEF, at which there was generally broad support for the capability among a wide variety of stakeholders (Field *et al.*, 2016). With the goal of an OEF system for California in mind, we present a spatiotemporal clustering component for the Third Uniform California Earthquake Rupture Forecast (UCERF3) developed by the ongoing Working Group on California Earthquake Probabilities (WGCEP).

Modeling Goals

Because our stated goal is to help communities prepare for potentially destructive earthquakes, viability of a model comes down to its reliability and skill with respect to forecasting larger triggered events.

Table 1 lists a variety of information that one might consider when addressing this question, with the overall challenge being how to integrate such constraints into a system-level model that embodies consilience, not only among these constraints, but also among whatever underlying assumptions are being made throughout the model. Our model also provides forecasts of smaller events, which are also useful for OEF messaging.

An important OEF milestone was the development of the short-term earthquake probability (STEP) model (Gerstenberger *et al.*, 2005), which provides real-time aftershock hazard estimates assuming the following, as combined by Reasenberg and Jones (1989, 1994): a Gutenberg–Richter (GR) magnitude–frequency distribution (MFD; Gutenberg and Richter, 1944), the Utsu relationship for the number of aftershocks as a function of mainshock magnitude (Utsu, 1971), and the modified Omori law for the temporal behavior of aftershocks (Utsu, 1961). The U.S. Geological Survey (USGS) posted STEP forecasts for California starting in

2005, but the system was taken offline in 2010 due to software maintenance issues. Nevertheless, STEP continues to be used to inform decision making in other parts of the world (e.g., Gerstenberger *et al.*, 2014). In lieu of a detailed review of that model here, Table 1 indicates which types of constraints are embodied in STEP, together with those utilized in the UCERF3 epidemic-type aftershock sequence (ETAS) model presented here. Many other candidate OEF models have also been developed and tested since the introduction of STEP (e.g., Werner *et al.*, 2011; Woessner *et al.*, 2011; Nanjo *et al.*, 2012, and references therein; Segou *et al.*, 2013; Zechar *et al.*, 2013, and references therein; Helmstetter and Werner 2014; Marzocchi *et al.*, 2014; Gerstenberger *et al.*, 2014). What makes UCERF3-ETAS unique to all of these is a more explicit and complete incorporation of geologic fault information, as well as the inclusion of elastic rebound (in which rupture probabilities are thought to drop on a fault after experiencing a large event and to grow back with time as tectonic stresses re-accumulate). In other words, our model attempts to merge the point-process-clustering models developed by statistical seismologists (e.g., Ogata, 1988; Reasenberg and Jones, 1989, 1994; Gerstenberger *et al.*, 2005) with the geologically based renewal models typically applied in official long-term time-dependent forecasts (e.g., WGCEP, 1988, 1990, 1995, 2003, 2007; Field *et al.*, 2009, 2015).

All scientifically viable models are ultimately wrong, because in addition to embodying assumptions and approximations, we possess only estimates of the constraints listed in Table 1. What we really hope is to have a model that is nevertheless useful for both risk-mitigation efforts, as already emphasized, and for improving our scientific understanding of earthquakes. To both these ends, an important goal for our model has been the ability to generate synthetic catalogs, also known as stochastic event sets. On the practical side, generating suites of synthetic catalogs enables constructing a complete probability distribution of potential losses, as opposed to being limited to just mean loss estimates when utilizing rate-based forecasts such as STEP. On the scientific side,

synthetic catalogs provide a powerful and perhaps indispensable perspective with respect to testing and evaluating models. For example, simulations have previously revealed that excluding elastic rebound leads to unrealistic aftershock statistics, such as accelerating sequences (e.g., Field, 2012).

Another improvement over STEP is to sample triggered events from the very same population of ruptures defined in the underlying long-term model. For example, we might want the likelihood of sampling an $M \geq 8$ event to depend on the proximity to capable faults (like the San Andreas) rather than assuming the same maximum magnitude and relative likelihood throughout the region. In fact, the California Earthquake Prediction Evaluation Council (CEPEC) is known to convene when a magnitude ~ 5 earthquake occurs near the San Andreas fault (SAF; Jordan and Jones, 2010), but not when such an earthquake occurs well away from known faults. In addition to fault proximity, one might also want to consider the activity rate of a fault (e.g., the mean recurrence interval), as well as whether the fault has had time to reload in an elastic-rebound sense.

Schoenberg and Bolt (2000) suggested one way to combine clustering and elastic-relaxation processes when modeling seismicity but did not consider finite faults or non-GR distributions explicitly. One of the first attempts to consider faults when defining earthquake triggering probabilities was the foreshock model of Agnew and Jones (1991), in which the *a priori* likelihood of having a large event, based on fault information in their examples, is considered in computing the probability of the event being triggered by a smaller prior earthquake. This foreshock model was not a complete OEF model because it did not include aftershocks. Probabilities obtained from the Agnew and Jones (1991) methodology are generally higher than those obtained using Omori–Utsu/GR statistics, which Michael (2012a) demonstrated to be a consequence of assuming a characteristic versus GR MFD when sampling aftershocks. Michael (2012a) also provided a generalized clustering model that relaxes the GR assumption and noted that having a characteristic distribution is the most effective way of getting above-average-triggering probabilities near a fault.

The UCERF3-ETAS model presented here essentially represents an elaborate implementation of the generalized clustering model of Michael (2012a), because the model effectively samples aftershocks according to the nearby MFD implied by the UCERF3 long-term model. Furthermore, and as anticipated by Michael (2012a), we find that the degree of characteristicness (the deviation from GR) is the most important factor when it comes to the likelihood of triggering large events from a given mainshock.

Model Description

To achieve all the goals set out above, the UCERF3 spatiotemporal component utilizes an ETAS model (Ogata, 1988), which is why we refer to the model as UCERF3-ETAS. Elastic-rebound effects are included by building upon

the UCERF3 long-term time-dependent model (UCERF3-TD; Field *et al.*, 2015), which in turn includes multifault ruptures and relaxes fault-segmentation assumptions by virtue of being built upon the time-independent model (UCERF3-TI; Field *et al.*, 2014). Details of these previously published models are reiterated here only to the extent that they are important for understanding UCERF3-ETAS. Those wanting a more complete understanding may wish to consult these previous UCERF3 publications.

Model Overview

The ETAS model is used to produce multiple realizations of how an earthquake sequence may progress over time. Each simulation takes a list of previous $M \geq 2.5$ events as input, which could represent an actual catalog, one or more hypothetical earthquake scenarios, or both. For each input and simulated earthquake, the latter of which includes spontaneous events, we randomly generate some number of $M \geq 2.5$ aftershocks using the ETAS model. Each event is sampled from the long-term UCERF3 model according to the current relative likelihood of each possible rupture, including both elastic-rebound effects on faults and proximity to the mainshock for triggered events. The hope, or assumption, is that ETAS represents an adequate statistical proxy for whatever physics is operating in the system. Furthermore, by incorporating observations of smaller earthquakes to inform the triggering potential of large ones (i.e., by updating the model with actual earthquake data), we hope to capture any static or dynamic triggering effects that may be playing out in the sequence and perhaps the potential implications of any induced seismicity.

Although conceptually simple, the actual UCERF3-ETAS implementation gets complicated with respect to including elastic rebound and dealing with uncertainties associated with modeled faults. Consequently, UCERF3-ETAS has a number of adjustable parameters, or variables, which are listed and defined in Table 2. Note that this is not a complete list of all possible adjustable parameters, which would be much longer, but rather a list of the more important ones explored here. For example, the *TimeSpan* specifies the start time and duration for the desired forecast, and *ProbModel* specifies if, and how, elastic-rebound effects are applied in the model. Three of the adjustable parameters (*ApplyGRcorr*, *ApplyGridSeisCorr*, and *ApplySubSeisForSupraNucl*) control the degree of characteristicness allowed throughout the model, which we already noted as having a first-order influence on triggering probabilities. The last parameter (*TotalRateScaleFactor*) allows one to correct for any overall biases in the total rate of simulated events. All of these variables, which are written in italics throughout this article, are described in more detail below.

The rest of this section describes the various components and algorithms utilized in UCERF3-ETAS, with the goal being one or more synthetic earthquake catalogs for a specified *TimeSpan*. The model has been implemented using

Table 2
The UCERF3-ETAS Model Variables, or Adjustable Parameters, That Are Explored in This Article

Model Variable Name	Description
<i>ApplyGRcorr</i>	This indicates whether a simulation should force all fault sections to be consistent with a Gutenberg–Richter (GR) distribution in terms of total supraseismogenic rates, which is achieved by dividing the nucleation rate of such ruptures by the subsection CharFactor value. Default value is False.
<i>ApplyGridSeisCorr</i>	This tells whether to correct gridded seismicity rates so as to be greater than or equal to the rate of aftershocks implied by the long-term rate of supraseismogenic ruptures. Default value is True.
<i>ApplySubSeisForSupraNucl</i>	This option makes the location where a given supraseismogenic rupture is likely to nucleate from proportional to the long-term rate of smaller events along the fault surface (inside the polygons of associated fault subsections). Default value is True.
<i>ProbModel</i>	This specifies the type of probability model to apply, with the options being: POISSON (time independent); FULL_TD (fully time dependent, including elastic-rebound triggering [ERT]); and NO_ERT (time dependent, but excluding ERT). Default is FULL_TD.
<i>TimeSpan</i>	The start time and duration of the desired forecast.
<i>TotalRateScaleFactor</i>	This defines the amount by which total, long-term rupture rates in the earthquake-rupture forecast (ERF) are scaled to obtain a better match between total simulated and target $M \geq 2.5$ rates. Default value is 1.14 (a 14% increase).

This is not an exhaustive list of all possible model variables (e.g., those in Table 3 could be adjusted as well) but rather only the ones addressed in this study.

OpenSHA, which is an open-source object-oriented platform for conducting seismic-hazard analysis (see [Data and Resources](#)). Care has been taken to implement generic components, which are not specific to California, in order to facilitate applications elsewhere.

The Long-Term Earthquake-Rupture Forecast (ERF)

One of the main model components is a long-term ERF, which by definition specifies every possible earthquake rupture (at some acceptable level of discretization), as well as the associated occurrence probability of each, for a given region and *TimeSpan*. An ERF may also have other adjustable parameters to enable, for example, setting different levels of accuracy or for specifying alternative logic-tree branches. For this study, we utilize one of the UCERF3-TD, and although any one of the 5760 logic-tree-branch options could be chosen, results presented here were obtained using a branch-averaged model that is described in the Figure 1 caption.

Although an ERF can be thought of as a long list of possible earthquake ruptures, the latter are actually bundled inside different earthquake sources. In other words, an earthquake source is a list of earthquake ruptures that are somehow related. Two types of sources are utilized in UCERF3: fault-based sources and gridded seismicity. The first type represents supraseismogenic fault-based sources, each of which involves a unique contiguous set of two or more of the fault subsections shown in Figure 1 (in which supraseismogenic means that the along-strike length of the rupture is greater than or equal to the average down-dip width of the faults involved, and each subsection has a length that is half the down-dip width). Each fault-based source therefore represents a specific rupture area (defined by its unique collection of fault subsections). Rupture magnitudes are determined from the source area, so a fault-based source will have more than one rupture only if multiple magnitudes are assigned for the given area; such aleatory variability is supported as an option in UCERF3 but not utilized in the results shown here, so each fault-based source corresponds

to a single fault-based rupture in this study. As such, there are about 260,000 fault-based sources, including many that represent multifault ruptures.

For the time-dependent model applied here, the probability of each rupture depends on how long it has been since each associated fault subsection experienced a supraseismogenic rupture, as defined by the elastic-rebound implementation described in the UCERF3-TD report (Field *et al.*, 2015). Stated simply, these probabilities are computed using a Brownian passage time (BPT) model after averaging the recurrence interval and time since last event over all the subsections utilized by the given fault-based rupture. The calculation is more complicated for faults where we only know that the last event predated some historical open interval; see the UCERF3-TD documentation (Field *et al.*, 2015) for full details.

Gridded seismicity is the other type of source in UCERF3, in which the region is discretized into about 7700 different 0.1° -by- 0.1° cells (Fig. 1) and a nucleation MFD is assigned to each. The sources represent subseismogenic ruptures on or near the explicitly modeled faults, as well as all ruptures elsewhere to account for unmodeled faults. Each grid-based source has five ruptures for each discrete magnitude in the MFD, one for each alternative focal mechanism represented in the model. The number of ruptures also increases if alternative finite rupture surfaces are included (also an option), although we treat gridded seismicity as point sources for simplicity here. The total number of ruptures from grid-based ruptures is about 1,700,000.

The ETAS Model

Aftershocks are sampled using the ETAS model introduced by Ogata (1988), which represents a generalization of modified Omori aftershock statistics in terms of providing a more detailed accounting of event pedigree. That is, the same statistical seismicity laws are used as in Reasenberg and Jones (1989) and STEP (Gerstenberger *et al.*, 2005), but in ETAS every earthquake can spawn others, so that some

Figure 1. 3D perspective view of the types of earthquake sources utilized in the Third Uniform California Earthquake Rupture Forecast (UCERF3). Each fault-based source/rupture is represented by a collection of two or more contiguous fault subsections, the latter of which are depicted as black-outlined parallelograms. The 0.1° -by- 0.1° grid draped over the entire region defines the gridded seismicity cells described in the text. The colors indicate the average likelihood that each cell or fault subsection will participate in an $M \geq 6.7$ earthquake in a 30-year period following 2014. This study utilizes a branch-averaged UCERF3 long-term time-dependent (UCERF3-TD) model, in which only Fault Model 3.1 is included, only the UCERF3 smoothed seismicity branch is used because it performs better for short-term forecasts, and only the Mid-Aperiodicity branch is used for the elastic-rebound time dependence; all other branches contribute to this average model according to their respective weights (see the UCERF3-TD report, including Fig. 3 therein, for more details on the various branch options). Although the UCERF3 time-independent (UCERF3-TI) and UCERF3-TD models have an option for removing aftershocks according to the [Gardner and Knopoff \(1974\)](#) declustering algorithm, this option has not been applied here, meaning that all earthquakes are included. The minimum magnitude considered by the forecast is also lowered from M 5 to 2.5.

events in an aftershock sequence are not triggered by the mainshock directly, but indirectly through a previously triggered aftershock. Events triggered by the first earthquake are referred to as primary aftershocks, those triggered by one of the latter are secondary aftershocks, and so forth. ETAS makes no distinction between aftershocks and any other type of triggered event.

We use the ETAS formulation introduced by [Ogata \(1998\)](#) and used by [Hardebeck \(2013\)](#), in which the nucleation rate of $M \geq M_{\min}$ earthquakes as a function of time (t) and space (x) is given as

$$\lambda(t, x) = \lambda_0 \mu(x) + \sum_{\{i: t_i < t\}} [k 10^{\alpha(M_i - M_{\min})} (t - t_i + c)^{-p}] \times [c_s (r + d)^{-q}], \quad (1)$$

in which λ_0 is the total rate of spontaneous events (background events) and $\mu(x)$ is the long-term spatial density

of $M \geq M_{\min}$ events. The summation is over all events that have occurred prior to time t , in which the first term in brackets gives the rate evolution of primary aftershocks, with k representing overall productivity, α representing the magnitude dependence of triggering, p representing the temporal decay rate, and c preventing a singularity at $t = t_i$. The second set of brackets in the summation gives the linear density of aftershock triggering, in which r represents distance from the rupture surface, q represents the linear decay rate, d prevents a singularity at $r = 0$, and c_s is a normalization factor that ensures a value of 1.0 when the term is integrated over all space. We adopt the ETAS parameter values also given by [Hardebeck \(2013\)](#), listed in Table 3, which were derived from California seismicity. Note that the α parameter is fixed to the b -value (1.0), which ensures that Båth's law holds independent of mainshock magnitude ([Felzer et al., 2002](#)).

The number of primary aftershocks expected from a parent of magnitude M is

Table 3

The ETAS Parameters Used in This Study, Representing the Preferred Values of [Hardebeck \(2013\)](#) Based on an Analysis of $M \geq 2.5$ California Earthquakes between 1993 and 2011

Parameter	Units	Value	Range
p		1.07	1.0–1.4
c	Years	1.78×10^{-5}	1.00×10^{-6} to 3.16×10^{-4}
q		1.96	> 1.8
d	km	0.79	> 0.63
k	Years $^{p-1}$	2.84×10^{-3}	3.79×10^{-4} to 4.97×10^{-3}
α		1.0	Fixed
$M_{\min}$		2.5	Fixed

The Range gives parameter values with log likelihoods within 5% of the maximum log likelihood, over a reasonable range of values of the other parameters (see [Hardebeck, 2013](#), for details).

$$N(M) = k10^{(M-M_{\min})} \int_{t_1}^{t_2} (c+t)^{-p} dt$$

$$= \frac{k10^{(M-M_{\min})}}{1-p} [(c+t_2)^{1-p} - (c+t_1)^{1-p}], \quad (2)$$

in which t_1 and t_2 represent the start and end times of the *Time-Span* (relative to the origin time of the parent). Table 4 lists the number of primary aftershocks expected in 10 years following various mainshock magnitudes, and assuming parameter values adopted here (Table 3). The number of aftershocks for subsequent generations (secondary and beyond) depends on the MFD from which one samples, and Table 4 includes values assuming the total regional MFD for UCERF3-TI (Fig. 2), which is GR consistent, applies everywhere.

One productivity metric we will use extensively in this article is the likelihood that an earthquake will trigger an event larger than itself, which is also listed in Table 4. The expected number of primary aftershocks that are larger than the mainshock is generally 0.053, which represents a 5.2% likelihood of having one or more such events. The expected number considering all generations is 0.165, which represents a 15% probability of one or more such aftershocks. These likelihoods are reduced at the highest mainshock magnitudes due to the tapering of the regional MFD (Fig. 2). These values are also for a 10-year time period following the mainshock; for comparison, values for 7 days following are 3.3% and 5.6% for primary and all generations, respectively, and values for 1-year following are 4.6% and 11%, respectively. This increase in likelihood with forecast duration serves as a reminder that we should never think “it’s a bit late for this to be an aftershock” ([Michael, 2012b](#), p. 630).

A subtle but potentially important issue is that spontaneous events in the ETAS model represent, in part, a proxy for aftershocks of unknown parents. The extent to which this is true can be quantified by the so-called branching ratio ([Helmstetter and Sornette, 2003](#)), which gives the number of primary aftershocks expected over infinite time for a parent magnitude sampled randomly from the assumed MFD. It turns out that the branching ratio of the parameters for [Har-](#)

Figure 2. The total long-term magnitude–frequency distribution (MFD) for the branch-averaged UCERF3 model used here (black) plus on-fault (red) and off-fault (blue) contributions, which sum to equal the total (black). The on-fault contribution includes subseismogenic ruptures inside each fault-section polygon. The thin dashed lines indicate a $b = 1$ extrapolation, in order to show how the distributions differ from a perfect Gutenberg–Richter (GR). The plot only goes down to M 4.0, even though the model goes to M 2.5.

[debeck \(2013\)](#) is effectively 1.0, meaning each event triggers, on average, one other event, implying that all events must be triggered and that the true rate of spontaneous events is zero ($\lambda_0 = 0$). Of course we do not have a complete history of all past events, so we still need a nonzero value of λ_0 to serve as a proxy for descendants of unknown parents. Furthermore, λ_0 must vary with time over the duration of a forecast. For example, if we have no information on past events, then λ_0 must equal the total rate of events at the beginning of the forecast (all events are spontaneous because there are no known parents), but as time increases λ_0 will go down as the forecast generates events, and λ_0 will eventually become zero at infinite time.

This issue is discussed at length by van der Elst (unpublished manuscript, 2016; see [Data and Resources](#)), and we use the equations therein to compute time-dependent fraction of spontaneous versus triggered events for our forecasts, which explicitly account for magnitude-dependent dates of completeness in an earthquake catalog. We use the UCERF3 California catalog compiled by [Felzer \(2013\)](#), together with magnitude-dependent completeness thresholds defined in her table 9, which thereby provides a list of historical/instrumental events between 1850 and 2012. Using this catalog in a forecast that begins in 2012 and given the total model MFD shown in Figure 2, the initial rate of spontaneous events (λ_0) is 0.30 times the total regional rate; this fractional value evolves to 0.28 at 10 years into the forecast, 0.24 at 100 years, and 0.20 at 1000 years. Again, these values depend on the branching ratio implied by the ETAS parameters, and although van der Elst (unpublished manuscript, 2016; see [Data and Resources](#)) provides additional support for a value near 1.0 for California, the practical implications and main

Figure 3. Subseismogenic (light gray bins) and supraseismogenic (dark gray bins) time-independent incremental nucleation MFDs for a Mojave subsection of the San Andreas fault (SAF). The black line is the combined cumulative MFD; the dashed black line is a perfect GR extension of the subseismogenic incremental MFD with the same maximum magnitude; and the dotted line is the cumulative distribution for the latter. (a) Original (uncorrected) rates for supraseismogenic ruptures, with the CharFactor described in the text being defined as the ratio of solid-black to dotted-black lines at M 6.3 (the minimum magnitude for supraseismogenic ruptures), which gives CharFactor = 3.3. (b) Correct MFD, in which the supraseismogenic rupture rates have been divided by the CharFactor value of 3.3, causing the solid-black and dotted-black lines to agree at M 6.3.

conclusions presented in this study do not change significantly with other published values (e.g., using the ETAS parameters of Hardebeck *et al.*, 2008, which have an implied branching ratio of 0.66 for the regional MFD shown in Fig. 2).

Characteristic Versus Gutenberg–Richter MFDs

For each simulated aftershock, UCERF3-ETAS effectively samples a magnitude from a nucleation MFD, or more technically, from a magnitude density function. One of the main findings of the UCERF-TI effort was that, in order to fit all the data constraints, many faults require a characteristic MFD, which is defined as having elevated rates at high magnitudes relative to a GR extrapolation from smaller magnitudes. Figure 3a shows the average time-independent nucleation MFD for one of the Mojave subsections of the SAF,

for which the supraseismogenic versus subseismogenic contributions are plotted separately. Note that the rates for supraseismogenic events are well above the GR extrapolation from subseismogenic rates, meaning the fault has a characteristic MFD. We quantify the degree of characteristicness by what we call the CharFactor, defined as the cumulative rate at the minimum supraseismogenic magnitude divided by the cumulative rate at that magnitude for a perfect GR distribution; the latter is defined as having the same total $M \geq 2.5$ rate, the same maximum magnitude (highest magnitude with nonzero rate), and a b -value of 1.0. The CharFactor value is 3.3 for the example shown in Figure 3a.

Figure 4a shows a map of CharFactor values for each fault subsection in the UCERF3-TI model, together with the cumulative distribution of values (black curve in Fig. 4d). Values range from a low of 0.0016 (1/625) at one end of the

Table 4

Expected Number of Aftershocks (Exp N) over 10 Years Following Various Mainshock Magnitudes (M_{main})*

Mainshock Magnitude (M_{main})	Exp $N \geq M$ 2.5				Exp $N \geq M_{\text{main}}$		Exp $N \geq M_{\text{main}} - 1.0$	
	Primary	Secondary	Tertiary Plus†	Total	Primary	Total	Primary	Total
2.5	0.0529	0.0362	0.0762	0.1653	0.053	0.165		
3.0	0.1672	0.1144	0.2409	0.5228	0.053	0.165		
3.5	0.5287	0.3616	0.7619	1.6533	0.053	0.165		
4.0	1.6719	1.1436	2.4093	5.2281	0.053	0.165		
4.5	5.2869	3.6164	7.6189	16.533	0.053	0.165	0.529	1.65
5.0	16.719	11.436	24.093	52.281	0.053	0.165	0.529	1.65
5.5	52.869	36.164	76.189	165.33	0.053	0.165	0.529	1.65
6.0	167.19	114.36	240.93	522.81	0.053	0.165	0.529	1.65
6.5	528.69	361.64	761.89	1653.3	0.053	0.165	0.529	1.65
7.0	1671.9	1143.6	2409.3	5228.1	0.050	0.157	0.529	1.65
7.5	5286.9	3616.4	7618.9	16533	0.039	0.121	0.529	1.65
8.0	16719	11436	24093	52281	0.013	0.041	0.529	1.65

*Based on Monte Carlo simulations and assuming the UCERF3-TI total magnitude–frequency distribution shown in Figure 2.

†Tertiary Plus includes the third-through-fifteenth generation of aftershocks.

Figure 4. Subsection CharFactor values, defined as the total rate of supraseismogenic ruptures divided by that implied by a GR extrapolation from small magnitudes (see text for details). (a) Uncorrected values (*ApplyGridSeisCorr* and *ApplySubSeisForSupraNucl*, both set as False). (b) Values for case in which *ApplyGridSeisCorr* = True and *ApplySubSeisForSupraNucl* = False. (c) Values for case when both *ApplyGridSeisCorr* and *ApplySubSeisForSupraNucl* are True. (d) The cumulative distribution of values, in which the black curve corresponds to those shown in (a), blue is for those in (b), and red is for those in (c).

Imperial fault to a high of 161 on the Surprise Valley fault, both of which are labeled in Figure 4a.

The CharFactor values have a direct influence on the likelihood of sampling supraseismogenic ruptures, which are generally the ones we care about from a hazard or risk perspective. For example, the expected number of $M \geq 6.3$ primary aftershocks expected from an $M 5.5$ parent is

~ 0.0084 ($0.053 \times 10^{5.5-6.3}$) if sampling from a GR distribution (dashed line in Fig. 3a), but the expected number becomes ~ 0.028 if sampling from the characteristic MFD in that figure, with the multiplicative increase being exactly equal to the CharFactor value of 3.3 for this fault section.

In other words, CharFactor values represent the factor by which the likelihood of triggering supraseismogenic ruptures

Figure 5. (a) The $M > 2.5$ nucleation rate (per year) in each grid cell for the branch-averaged UCERF3 model used in this study (described in Fig. 1 caption). (b) The long-term rate of aftershocks expected from the long-term rate of supraseismogenic ruptures only. (c) The ratio of (b) to (a), in which values less than 1.0 are plotted as gray, revealing areas where the long-term model (a) must be deficient in $M > 2.5$ earthquakes. (d) The $M > 2.5$ nucleation rate for the corrected model (*ApplyGridSeisCorr* = True), in which rates in (a) have been multiplied by the ratio in (c).

differs from that of a perfect GR distribution, conditioned on the occurrence of a nearby parent event. This implies that if the CharFactor value for the Mojave section in Figure 3 were a factor of 161 (the value for the Surprise Valley fault), then the expected number of $M \geq 6.3$ primary aftershocks triggered from an $M 5.5$ parent would be 1.35. We need to be somewhat careful here because the likelihood of having more than one

event will depend on whether elastic-rebound effects are included; if so, the occurrence of the first supraseismogenic rupture will change/reduce the likelihood that others can follow. Furthermore, secondary, tertiary, and subsequent generations of aftershocks could also trigger an $M \geq 6.3$ event, increasing likelihoods compared with those cited above. Nevertheless, the point stands that a sufficiently high CharFactor value will

Figure 6. The depth distribution of hypocenters assumed for California, derived from the UCERF3 catalog (Felzer, 2013).

imply near certainty of triggering at least one supraseismogenic rupture given a smaller nearby event, which highlights the need to carefully scrutinize the degree of characteristicness allowed throughout the model.

Several factors influence the CharFactor value inferred for each fault section. The subseismogenic MFD reflects an empirical estimate of $M \geq 2.5$ event rates near the fault (obtained by integrating smoothed seismicity rates over associated fault-section polygons, the latter of which are described more below). The fact that the total rate of $M \geq 5$ events throughout the entire region is uncertain by about 20% (according to UCERF3-TI logic-tree branch options) implies that we should expect considerable observational uncertainty at the fault-section polygon level. Supraseismogenic nucleation rates for each fault section are determined from the UCERF3-TI grand inversion, which incorporates information on nearby fault connectivity, slip-rate estimates, paleoseismic event-rate data where available, and other constraints, such as staying as close as possible to the previous model (UCERF2). The high CharFactor value for the Surprise Valley fault could be biased from either erroneously low $M \geq 2.5$ event-rate estimates in that area, due to a historical lull or network detection issues, or because a lack of known connectivity with neighboring faults produces an erroneously high rate of moderate-sized events when satisfying the slip-rate constraint. The point is that individual CharFactor values are uncertain, and we need to be careful not to propagate any poorly constrained attributes of UCERF3 long-term models into UCERF3-ETAS forecasts, especially if they have a first-order influence on consequent triggering statistics.

One reality check is whether the subseismogenic MFDs near faults are high enough to include the expected long-term rate of aftershocks from supraseismogenic events. Figure 5 reveals that the answer is no near several faults: Figure 5a shows a map of the total long-term event rates implied by UCERF3-TI, which mostly reflects the smoothed seismicity model used for gridded, background seismicity; Figure 5b shows the rate of primary aftershock expected from the

long-term rates of fault-based ruptures, multiplied by a factor of 2 to approximate the contribution of subsequent generations; and Figure 5c shows areas where, and the extent to which, the ratio of the latter to the former exceeds 1.0 (in which UCERF3-TI does not satisfy the expected rate of aftershocks from supraseismogenic events). These areas imply a lack of self-consistency in that supraseismogenic ruptures in the model will trigger more $M \geq 2.5$ events than the long-term model exhibits in the first place, which represents a model inconsistency that results in artificially inflated CharFactor values. For UCERF3-ETAS, we consequently support the correction of gridded seismicity rates so they do not fall below the values shown in Figure 5b, which amounts to multiplying the values in Figure 5a by the ratios in Figure 5c, leading to the values shown in Figure 5d. This correction is made according to whether the *ApplyGridSeisCorr* parameter (Table 2) is set as true or false.

The CharFactor values that result from setting *ApplyGridSeisCorr* = True are shown in Figure 4b, along with the cumulative distribution (blue curve in Fig. 4d), revealing a significant reduction of some of the higher CharFactor values. The maximum value is now 72 on the Robinson Creek fault (labeled on Fig. 4b), and the value for the Surprise Valley fault has gone from 161 to 56.

Up until this point we assumed that supraseismogenic ruptures have an equal probability of nucleating from anywhere on the fault surface. An alternative assumption is that the likelihood of nucleation correlates with the rate of subseismogenic events. In other words, if the rate of little earthquakes is twice as high at one end of a candidate rupture, it might also be twice as likely to nucleate from that end. This assumption serves to further reduce the range of CharFactor values throughout the model, although the effect is relatively mild as shown in Figure 4c (and the red curve in Fig. 4d). We nonetheless support this correction as an option in UCERF3-ETAS, and it is set by the *ApplySubSeisForSupraNucl* parameter value (Table 2). The final value for the Robinson Creek fault is still 72, and the Surprise Valley fault value is now down to 45.

Finally, UCERF3-ETAS includes the option to impose GR throughout the model (by setting *ApplyGRcorr* = True), which effectively divides the nucleation rate of each supraseismogenic rupture by the associated CharFactor value. The result of this correction is shown in Figure 3b for the Mojave S example. Note that although the total rate of supraseismogenic ruptures now matches GR after this correction (at $M \geq 6.3$ in Fig. 3b), rates do not necessarily match at higher magnitudes, given our particular definition of characteristicness.

Overall Calculation Sequence

In addition to an ERF, as described above, the UCERF3-ETAS model takes the following as input: (1) the desired *TimeSpan* and region for the simulation, where the latter is set here as the entire UCERF3 area (Fig. 1); (2) a list of all $M \geq 2.5$ earthquakes that have occurred up until the chosen simulation start time, plus any fake ruptures that might be

Figure 7. This figure poses the question of whether a larger rupture (red) can be triggered from within the area of a smaller supraseismic event (blue area) that has just occurred. The question is relevant because most small aftershocks (stars) will be in the blue area, as illustrated, so the likelihood of triggering the red event will be much higher if the answer is yes.

desired for testing purposes; (3) the average depth distribution of earthquake hypocenters for the region (which assumed here is shown in Fig. 6); and (4) yes or no answers to various calculation options, including whether to include spontaneous events, indirect triggering, and whether to apply the corrections described above (*ApplyGRcorr*, *ApplyGridSeisCorr*, and *ApplySubSeisForSupraNucl*, respectively). Another setting is the type of time dependence to apply in the model, specified by the *ProbModel* parameter (Table 2), the options of which will be described below. Finally, we have a *TotalRateScaleFactor*, which simply scales the total rate of spontaneous events in order to correct any bias between total simulated and target $M \geq 2.5$ rates.

Given these inputs and settings, the algorithm used to produce a simulated catalog of ruptures, the desired outcome, is described in Appendix A. The particularities of the approach described therein stem from the need for a numerically efficient way of sampling ETAS events while honoring an elastic-rebound model on faults. Appendix A assumes the availability of another major model component, the *ETAS_PrimaryAftershockSampler*, which is described in the next section.

Sampling Primary Aftershocks from the ERF

Here, we describe how primary aftershocks are randomly sampled from the ERF, which is achieved via the *ETAS_PrimaryAftershockSampler* component. Again, we want to sample from the population of ERF ruptures according to the current relative probability of each and the distance decay defined by equation (2). To do so, we discretize the entire region into cubes that are about 2 km on a side. Specifically, each 0.1° -by- 0.1° grid cell (Fig. 1) is subdivided into 0.02° increments, and depth is discretized at 2 km increments down to 24 km (see Fig. S1, available in the electronic supplement of this article). This results in 300 cubes per grid cell, and about 2,300,000 cubes for the entire UCERF3 region. When sampling an aftershock for a given parent, first we randomly select a cube from which the triggered event nucleates, and then randomly sample one of the ruptures that can nucleate from within the cube according to the current relative probability of each doing so.

The next step is to define the rate at which each source or rupture nucleates inside each cube according to the ERF. For gridded seismicity, we simply divide each grid-cell rate equally among the 300 associated cubes (the depth dependence of nucleation is handled later). Fault-based ruptures might seem straightforward as well, because one could assume a uniform distribution of nucleation across the surface and then distribute nucleation rates among the cubes crossed by the fault surface according to the rupture area within each. Regrettably, this simple approach does not work for two reasons: elastic-rebound triggering (ERT) and the spatial uncertainty of faults, both of which are discussed next.

Elastic-Rebound Triggering. The ETAS model raises an interesting question with respect to triggering on finite faults, which is illustrated in Figure 7 and can be described as follows: suppose a 50-km-long section of the southern SAF has just ruptured. According to UCERF3-TD, a repeat of this exact rupture has a zero probability of occurrence in the near future because of elastic rebound. However, the probability of having an even longer rupture that overlaps the one that just occurred is nonzero (although reduced according to UCERF3-TD because the average time since last event has gone down). The critical question is this: Can a longer rupture nucleate or be triggered from within the 50 km stretch that just ruptured? The answer to this question has a strong influence on triggering statistics because the vast majority of aftershocks spawned by the first rupture will occur right along the 50 km section, as illustrated in Figure 7, so the probability of triggering the longer rupture will be much higher if the event can nucleate from this previously ruptured stretch. The alternative is that the longer rupture has relatively little likelihood of nucleating along the 50 km section, leaving only aftershocks near the ends of the previous rupture able to do any large event triggering. The latter assumption is included as an option in UCERF3-ETAS depending on the setting of the *ProbModel* parameter; specifically, a setting of *FULL_TD* (full time dependence) means that supraseismic ruptures cannot be triggered from a

zone that just ruptured, and a setting of NO_ERT means it can.

Although it is straightforward to prevent a larger rupture from nucleating within the rupture zone of a very recent large event, the harder question is how this zone transitions back into nucleation viability with time. Regrettably, UCERF3-TD tells us nothing about where fault-based ruptures will nucleate. We therefore make the simple assumption that the relative likelihood that a particular section will nucleate a given rupture is proportional to

$$\frac{A_s \eta_s}{\sum (A_s \eta_s)}, \quad (3)$$

in which η_s is the normalized time-since-last-event of the section (time-since-last-event divided by the average long-term recurrence interval), A_s is the area of the section, and the summation in the denominator is the overall sections utilized by the particular rupture (a normalization that ensures that section-nucleation probabilities sum to 1.0 for the rupture, so that the overall likelihood of the event remains unchanged). That the relative triggering probability increases linearly with time may seem at odds with typically applied renewal models (e.g., BPT), which have a distinctly nonlinear increase. However, we are only talking about where the event might nucleate from, and we are not modifying the overall rupture likelihood defined by the ERF (which does honor the BPT renewal model). That being said, we need to watch for potential biases, particularly when simulating events over time periods that approach or exceed the recurrence intervals of faults. We also acknowledge that a more elegant approach is certainly fathomable, especially with respect to making nucleation likelihood a more integral part of the elastic-rebound formulation. This will be a significant challenge for models that relax segmentation and include multifault ruptures, however, and any progress will likely require at least inferences from physics-based simulators that include Omori-type triggering (e.g., Richards-Dinger and Dieterich, 2012). For UCERF3-ETAS, we apply the triggering probabilities defined by equation (3) when the *ProbModel* parameter is set as FULL_TD; otherwise the time-since-last-event has no bearing on where a rupture nucleates.

There is also observational evidence that smaller subseismogenic ruptures exhibit such ERT effects. Using high-precision double-difference relocations for observed M 4–6.7 earthquakes, van der Elst and Shaw (2015) have shown that aftershocks that are as large or larger than the parent nucleate almost exclusively in the outer regions of the parent aftershock zone, which they interpret as evidence for elastic relaxation. Though applying such a recent finding here may seem premature, it does solve another potential problem with respect to triggering probabilities when transitioning between subseismogenic and suprasedismogenic mainshocks. For example, if an M 6.3 suprasedismogenic mainshock cannot trigger larger events from the subsections of its rupture surface, then the probability of triggering such events could be lower than for a slightly smaller subseismogenic mainshock if the

latter can trigger from anywhere. In other words, we would have a sharp drop in triggering probability as we transition from subseismogenic to suprasedismogenic mainshocks (all other things being equal).

To avoid this problem, we simply apply the rule that non-fault-based ruptures larger than M 4 cannot trigger any suprasedismogenic fault-based ruptures from within a sphere centered on the hypocenter. The sphere radius is set equal to the radius of the source, in which the latter is assumed to be circular with an area (A) given by

$$A = 10^{M-4}, \quad (4)$$

in which M is the magnitude. This source area is consistent with both the Hanks and Bakun (2008) and Shaw (2009, 2013) scaling relationships for $M \lesssim 6.3$ events. That the source radius is applied to a spherical volume around the hypocenter reflects the fact that we will not generally know the actual rupture plane, so we effectively account for any possibility. This rule, which is only applied when *ProbModel* is set as FULL_TD, is admittedly unsophisticated, but nevertheless a seemingly reasonable place to start.

An interesting implication of ERT is a lowering of triggering probabilities relative to GR because larger events can only be triggered from off the ends of the rupture surface (Fig. 7). More specifically, for suprasedismogenic mainshocks, both the expected number of aftershocks and rupture lengths scale similarly with magnitude/length, so the overall density of aftershocks over the rupture surface remains invariant with magnitude. This means that for a long and completely isolated fault, triggering likelihoods will also become invariant with increasing rupture length or magnitude (again, all other things being equal). To quantify this effect, we ran a series of tests for various-sized events on the southern SAF (Mojave section and beyond), where other model variables were set so as to meaningfully isolate the influence of ERT (this included application of CharFactor corrections by setting *ApplyGRcorr* = True and using a time-independent model for overall rupture probabilities). When ruptures were allowed to nucleate from the mainshock rupture zone, the expected number of $M \geq 6.3$ primary aftershocks was exactly equal to that predicted by GR for all mainshock magnitudes. Suppressing such triggering produced the following $M \geq 6.3$ triggering-likelihood reductions relative to GR:

- 0.72 for M 5.0 point-source mainshock on the fault,
- 0.48 for M 5.5 point-source mainshock on the fault,
- 0.24 for M 6.3 point-source mainshock on the fault,
- 0.17 for M 6.3 fault-based mainshock,
- 0.08 for M 7.0 fault-based mainshock,
- 0.06 for M 7.4 fault-based mainshock, and
- 0.09 for M 7.8 fault-based mainshock.

As expected, the reductions relative to GR become greater with increasing mainshock magnitude (because of an increased nucleation exclusion zone), and become approximately invariant above M 7, although not exactly invariant due to nearby branching faults in these examples (e.g., the

Figure 8. Illustration of how supraseismogenic on-fault nucleation rates are distributed over the cubes that are within each fault-section polygon, in which the latter are depicted with black lines on the upper surface, which extend to depth vertically. Images at the top show $M \geq 6.3$ nucleation rates, and those on the bottom depict the CharFactor values described in the main text for cubes near SAF (Mojave S) subsections. (a) and (c) correspond to the case in which supraseismogenic fault-based ruptures only nucleate from within cubes that are intersected by the fault surface (magenta- and red-colored cubes for the San Andreas), leaving other cubes within the subsection polygons to have zero supraseismogenic rates (white). (b) and (d) show the solution adopted here, in which cube nucleation rates decrease linearly out to 10 km and are constant beyond, such that the total nucleation rate in the polygon is preserved, and CharFactor values beyond 10 km are 1.0 (see text for details).

M 7.8 event passes by the Garlock fault, but the M 7.4 does not). These test results are for a particular set of ruptures on one particular fault and may not apply to all other possible mainshocks, especially where there are more or less nearby branching faults. Nevertheless, the tests do give us a general idea of the influence of ERT.

Note that ERT counters the influence of characteristic MFD. In fact, if ERT is real, one will need a CharFactor of about 6.0 just to match GR probabilities for the M 6.3 mainshock case above; this is just to break even with GR, as the CharFactor will have to exceed 6.0 if we think the triggerability should be higher (e.g., because we are on a very active fault).

Thus, there could be long-term characteristic earthquake behavior as required by the grand inversion results from UCERF3 but still could be self-similar foreshock–mainshock–aftershock behavior as observed for southern California, Italy,

New Zealand, and global seismicity (as discussed in [Michael, 2012a](#), p. 2548, and references therein).

Distributing Nucleation Rates Over Fault Section Polygons. The other challenge is that fault surfaces in UCERF3 actually represent a proxy for all supraseismogenic ruptures occurring within a fault-zone polygon, with the latter being predefined in UCERF3-TI for each fault section. For example, the polygon for vertically dipping faults extends out to 12 km on each side of the trace, perpendicular to the strike, as shown for a part of Mojave SAF in Figure 8. The intent was to be more explicit about whether an event like the 1989 Loma Prieta earthquake is a hit or miss with respect to the SAF, or whether the 2010 El Mayor–Cucapah earthquake was a hit or miss with respect to the Laguna Salada fault. A 12 km zone on each side, which is about the same as the seismogenic thickness, also seems reasonable in terms

of the spatial reach of supraseismogenic ruptures with respect to potential elastic relaxation.

To avoid double counting, the maximum magnitude for gridded seismicity sources inside a fault-zone polygon is just less than the minimum magnitude of the associated supraseismogenic ruptures (e.g., the light- vs. dark-gray bins in Fig. 3). If nucleation of a fault-based rupture can only occur from cubes that are intersected by the surface, then all other cubes within the polygon will have a zero nucleation rate above the minimum supraseismogenic magnitude, as illustrated in Figure 8a. Even worse, because the gridded seismicity nucleation rates have been distributed equally among the 300 cubes within each grid cell, the nucleation MFD for cubes intersected by the rupture surface would have a greatly exaggerated rate difference between sub- and supraseismogenic ruptures, effectively increasing the CharFactor values therein by more than an order of magnitude (Fig. 8c).

To avoid these problems, we must partition the nucleation rate of each fault-based rupture among all cubes contained within the associated fault-zone polygons. Partitioning nucleation rates uniformly is one option but would imply that the likelihood of triggering a supraseismogenic event is constant over the polygon area, and therefore independent of distance from the actual fault plane. We add some distance dependence by assuming a linear nucleation-rate decay with distance from the fault surface out to 10 km, beyond which the nucleation rates are held constant with an implied CharFactor of 1.0, and such that the total nucleation rate is preserved (i.e., summing rates over the cubes inside the fault-section polygon equal the total nucleation rate for the section). Figure 8b and 8d shows the consequent nucleation rates and implied CharFactor values, respectively, for cubes near SAF (Mojave S) subsections, illustrating that our rate-partitioning algorithm provides both decay with distance from the fault surface and a continuous transition with respect to rates and implied CharFactor values for cubes just outside the polygons (in which CharFactor values are 0.74 because the off-fault MFD [blue curve in Fig. 2] is that much below the GR extrapolation at M 6.3).

The linear partitioning of nucleation rates means the implied CharFactor values for cubes closest to the fault are greater than the section average (e.g., 151 and 103 for Robinson Creek and Surprise Valley faults, respectively), which is consequence of applying monotonic decrease with distance from the fault and a continuous transition at the polygon boundary. This algorithm does not make sense for fault sections that have a CharFactor less than 1.0, because it would imply a monotonic increase with distance from the fault. We therefore distribute nucleation rates uniformly when a fault-section CharFactor is less than 1.0.

Figure 9 shows the implications of this partitioning throughout the region for cubes at 7 km depth. Figure 9a and 9b shows the nucleation rate of $M > 2.5$ and $M \geq 6.3$ events, respectively, Figure 9c shows the implied CharFactor of each cube, and Figure 9d shows the maximum magnitude in each cube.

With the rules for ERT and the distribution of nucleation rates across fault-section polygons hereby established, we are now able to compute the rate at which each ERF rupture nucleates inside each cube (i.e., the nucleation MFD for each cube).

Distance Decay. In sampling ruptures, we also need to account for the distance decay of aftershocks, as given by the second term in brackets in equation (1). This is the linear distance decay, meaning the relative rate or density of triggered events at distance r , rather than the density at any particular point in space. Two complications arise here. First, the zone of potential earthquakes transitions from a sphere, when r is less than the distance to both the free surface and the bottom of the model (24 km here), to something more like the outer surface of a cylinder at large distances. Second, we want the model to honor the depth dependence of seismicity given in Figure 6, meaning that a parent at a depth of 12 km is more likely to trigger an event 4 km above (at 8 km depth) than at 4 km below (16 km depth). Both these complications imply that the triggering density at a point in space depends on the depth of the parent event.

We solve this problem by first assuming that the parent event is located at a cube corner, meaning that it has a depth of 0, 2, 4, ..., or 24 km (Ⓔ Fig. S1). For each of these depths, we create a separate *DistanceDecayCubeSampler*, which contains the relative likelihood that each neighboring cube will host a triggered event, in which the probability density for each cube is calculated in a way that honors both the distance decay (equation 1) and the depth dependence shown in Figure 6. The problem is solved numerically, taking care to properly integrate over the entire volume for the closest cubes (because the nonlinear decay means that the value computed at the center of the cube can be significantly different from the volume-averaged value). The result is illustrated for two different parent depths in Figure 10.

The *DistanceDecayCubeSampler* provides a randomly sampled cube according to the relative probability density of each, and it also gives a random location within the cube, where the latter accounts for any significant rate-density transitions inside the cube (i.e., for those in close proximity to the parent, which most cubes will be). In lieu of further numerical details on how all this is achieved, which can be found in the computer code, Ⓔ Figure S2 shows that random samples are in good agreement with the target distance decay over distances from 50 m to 1000 km.

Sampling Ruptures. Full details on the algorithm used to sample primary aftershocks for a given parent rupture are described in Appendix B. In essence, a cube is sampled according to the distance decay from the parent rupture surface, and a rupture is sampled from within the cube, based on the current relative likelihood that each can nucleate from within. The code does not explicitly sample from a cube MFD but rather follows a more numerically efficient procedure in which it first decides between a gridded seismicity and fault-based source according to the current relative nucleation rate of each. Furthermore, fault-section nucleation rates are calculated from

Figure 9. Implications of how nucleation of supraseismogenic fault-based ruptures are distributed among the cubes that are inside each fault-zone polygon (as described in the text). (a,b) Nucleation rates for cubes at 7 km depth for $M \geq 2.5$ and $M \geq 6.3$ events, respectively. (c) The implied CharFactor value of each cube at 7 km depth. (d) The implied maximum magnitude for cubes at 7 km depth.

and kept in sync with any changes in the elastic-rebound-implied likelihood of each fault-based rupture. Again, further details are given in Appendix B.

Results

UCERF3-ETAS is an admittedly complicated model with a lot of calculation options. Presenting an exhaustive

set of results and sensitivity tests here is therefore not feasible. Instead, we present a limited set of simulations aimed at highlighting both positive features, and more importantly, potential issues and remaining challenges with respect to deploying a potentially useful operational system.

We first present results from long-term (e.g., 1000 year) simulations, which not only provide valuable insights into model behavior but also suggest some corrections that could

Figure 10. Illustration of the *DistanceDecayCubeSampler* described in the text, in which the relative likelihood of sampling an event in each location is shown by the color. (a) and (b) The cases in which the depth of the parent event is 6 and 18 km, respectively.

be made. We then present 10-year probabilities implied by the historical/instrumental catalog. Finally, we use hypothetical earthquake scenarios to quantify the conditional triggering probabilities implied by some events of interest. Although individual simulations can easily be run on a desktop computer, we made extensive use of high-performance computing resources in generating the large ensembles presented below (see Acknowledgments).

1000-Year Simulations

Figures 11–16 show results aggregated from approximately five hundred 1000-year simulations using the default parameter values listed in Table 2, including a *ProbModel* setting of *FULL_TD* (which includes ERT, meaning supraseismogenic ruptures are prevented from nucleating from within recently ruptured zones). The start year for the simulations is 2012, set by when the UCERF3 catalog ends, and the latter was input to the model as potential parents (although this does not really significantly influence 1000-year results). Figure 11 shows the simulated versus target cumulative MFD for the entire region. The match is good below M 5 (within 4%), but there is a 25% underprediction near M 7. Figure 12a shows simulated $M \geq 2.5$ nucleation rates throughout the region at the spatial resolution of the cubes, and Figure 12b shows the ratio of simulated versus target $M \geq 5$ rates in each grid cell, in which the mean ratio is 0.93, the median is 0.79, the minimum is 0.062, and the maximum is 5.4. Areas dominated by gridded seismicity (away from faults) are generally undersimulating the rate of $M \geq 5$ events, and areas near several faults are oversimulating such rates. This is a natural and expected manifestation of the spatial variability of long-term nucleation MFDs throughout the region, where areas exhibiting *CharFactor* values less than 1.0 will undersimulate event rates and vice versa

Figure 11. The average cumulative MFD obtained from five hundred 1000-year UCERF3 epidemic-type aftershock sequence (ETAS) simulations (black line) compared with the target MFD (gray line). Default parameter settings were applied in these simulations, and the target represents the total UCERF3-TI distribution (Fig. 2) modified according to the influence of having *ApplyGridSeisCorr* = True (which increases rates below about M 6.3 by about 4.6%).

for areas with values greater than 1.0. As discussed above, the fractional rate of spontaneous versus triggered events depends on the nucleation MFD, so to match the rates everywhere in Figure 12b, we would really need to vary the fractional rate in space (in addition to time) in order to compensate for the spatial variability of long-term MFDs. This is also why the default value of the *TotalRateScaleFactor* parameter is 1.14; that is, we need to increase the rate of spontaneous events by 14% to better match the total rate of $M \geq 2.5$ events.

Figure 12c shows simulated fault-section participation rates in map view, and Figure 12d shows the ratio of the latter to target values. Figure 13 shows a scatter plot of simulated versus model target participation rates for each subsection, color-coded by the fraction of events on each fault section that were triggered by supraseismogenic fault-based ruptures (as opposed to being spontaneous or triggered by subseismogenic events). On average, 33.5% of fault-based ruptures are spontaneous, 31% are triggered by subseismogenic ruptures (meaning no ancestors were supraseismogenic), and 35.5% of ruptures had at least one supraseismogenic ancestor. The symbol colors in Figure 13 reveal a strong correlation between participation-rate discrepancies and the fraction of events triggered by supraseismogenic events; oversimulated rates generally result from above-average triggering from other fault-based ruptures and vice versa for undersimulated rates.

A number of the outliers are labeled in Figures 12d and 13. The high ones, labeled with orange text, represent faults that are being triggered by the nearby and relatively active SAF. The low outliers, labeled in blue, represent faults where recent earthquakes suppressed occurrences over the 1000-year simulation duration by virtue of elastic rebound (and long recurrence intervals). Also labeled is a point on the SAF (Mojave S), where

Figure 12. Results obtained from 1000-year UCERF3-ETAS simulations based on default parameter settings, in which results represent averages from 500 different runs. (a) $M \geq 2.5$ nucleation rates throughout the region at a resolution of 0.02° in latitude and longitude. (b) Ratio of simulated to model-target nucleation rates for $M \geq 5.0$ events, binned at a 0.1° resolution. (c) The average rate at which each fault section participates in a supraseismogenic rupture. (d) The ratio of (c) to model-target values (the latter from UCERF3-TI).

the simulated rate is 25% above the target, which corresponds to a simulated recurrence interval of 75 years versus the target value of 90 years. This too is a manifestation of triggering by supraseismogenic ruptures. UCERF3-TI had to assign a relatively high rate for M 6.3–6.5 events in this area to satisfy

the relatively short recurrence interval implied by the Wrightwood paleoseismic record. The occurrence of one such event is apparently overtriggering others in the UCERF3-ETAS simulations. Figure 14 shows a histogram of simulated recurrence intervals at this location, where the height of the first bin

Figure 13. The average simulated participation rate of supraseismic events on each subsection, plotted against the model-target value, for the 1000-year UCERF3-ETAS simulations based on default parameter settings. The symbol color indicates the fraction of events that were triggered by supraseismic fault-based ruptures, either directly or indirectly, as opposed to being spontaneous or triggered by subseismic events. The red dashed line represents binned average values, in which the bin width is 0.2 in log10 space, corresponding to 20 bins over the x -axis range.

implies that 9% of the ruptures are occurring within 7 years, and the fact that this bin is the mode of the distribution implies that such a recurrence interval is the most likely one to be observed. Elastic rebound prevents a recurrence of an identical rupture, so the triggered events are different but with spatial overlap. Such recurrence intervals are not likely to be observed paleoseismologically, both because they are short and because M 6.3–6.5 events might go undetected. If we exclude the first bin in computing the mean recurrence interval, the value is 82 years, which is still a bit below the 90-year target.

On average, the match in Figure 13 is relatively good for sections with participation rates greater than 0.004 per year (a recurrence interval of less than 250 years), but below this threshold simulated rates are 25%–33% below targets on average (red-dashed line). One possible explanation is that a 1000-year simulation is not long enough to see the full average-triggering effect on such low-rate faults (i.e., perhaps we need more than about four earthquakes in each simulation). We therefore did some 10,000-year simulations to see if the agreement transition shifts (e.g., from 0.004 per year to a lower value of 0.0004 because the simulation is 10 times longer). The shift was minimal, however, and certainly not enough to explain the full discrepancy. Another benefit of these 10,000-year simulations is that they allowed us to see

Figure 14. Histogram of recurrence intervals for supraseismic events on SAF (Mojave S) subsection 13, obtained from the 1000-year UCERF3-ETAS simulations based on default parameter settings. Bin widths are 7.2 years, and the mean recurrence interval is 75 years.

whether any of the discrepancies discussed above grow with time, and fortunately we did not find any such instability.

So what else might explain the participation-rate agreement transition at 0.004 per year in Figure 13? A log-log plot of participation-rate discrepancy ratios versus the section CharFactor value is shown in Figure 15, together with a linear fit that implies lower rate faults have a lower CharFactor value on average. As explained above, faults with lower CharFactor values will need to have commensurately more spontaneous events in order to match long-term rates, so perhaps this explains the systematic discrepancy at lower rates in Figure 13.

To examine the overall elastic-rebound predictability implied by UCERF3-ETAS, Figure 16 shows a histogram of normalized fault-section recurrence intervals aggregated over all faults and all simulations. The normalization, which amounts to dividing each subsection recurrence interval by the expected mean before adding it to the histogram, is needed to make the comparison meaningful. The normalized recurrence intervals implied by the UCERF3-TD model (without spatiotemporal clustering) and those implied by the RSQSim physics-based simulator (Richards-Dinger and Dieterich, 2012) are also shown, both adopted from Field (2015). As expected, inclusion of spatiotemporal clustering in UCERF3-ETAS increases the aperiodicity (widens the bell curve, or increases the coefficient of variation) and also increases the relative frequency of the shortest recurrence intervals relative to UCERF3-TD (because of spatially overlapping triggered ruptures). The UCERF3-ETAS distribution is qualitatively similar to that of RSQSim, although the latter has a lower coefficient of variation and a smaller fraction in the first bin, implying that RSQSim has less spatial overlap with respect to quickly triggered adjacent ruptures.

We now turn to describing the influence of alternative parameter settings with respect to 1000-year simulations, restricting the discussion to only noteworthy differences.

Figure 15. Supraseismogenic participation-rate discrepancy (simulated over target) versus the CharFactor value for each fault subsection, together with the best-fit line (in log–log space). Simulated values are from the 1000-year UCERF3-ETAS simulations with default parameter settings, and the CharFactor values are those shown in Figure 4c accordingly.

Turning Off Corrections. In accordance with default values, the above simulations used the following setting: *ApplyGridSeisCorr* = True, *ApplySubSeisForSupraNucl* = True, and *TotalRateScaleFactor* = 1.14. Setting the first two to False and the last one to 1.0 (no correction) generally makes the various discrepancies discussed above worse. For example, (E) Figure S3 shows simulated versus target subsection participation rates for this case, in which an increase in the scatter is apparent when compared with the default-parameter case (Fig. 13). In particular, discrepancies are even greater for the more extreme oversimulated cases, with the recurrence interval for the Mojave S subsection discussed above now being down to 45 years (from a value of 75 years for default parameters and compared with the 90-year target). The difference is mostly a manifestation of the *ApplyGridSeisCorr* parameter, in which not applying this correction causes the simulation to produce more earthquakes than exist in the long-term model near several faults, and these surplus events in turn trigger extra supraseismogenic ruptures due to the artificially high CharFactor values. As already mentioned, the default value of the *TotalRateScaleFactor* is 1.14 because a value of 1.0 causes about a 14% undersimulation of total $M \geq 2.5$ event rates (not shown).

ProbModel Parameter. (E) Figure S4 shows the subsection participation scatter diagram for the case in which *ProbModel* is set as NO_ERT, meaning supraseismogenic ruptures are

Figure 16. (a) Normalized subsection recurrence intervals for supraseismogenic events from the 1000-year UCERF3-ETAS simulations, based on default parameter settings, in which the normalization involves dividing each recurrence interval by that expected before adding to the histogram. (b) The normalized recurrence intervals implied by UCERF3-TD (i.e., without spatiotemporal clustering), and (c) the equivalent results obtained from the RSQSim physics-based simulator (Richards-Dinger and Dieterich, 2012). The latter two plots come from Field (2015).

free to trigger from anywhere. All other parameters were set to default values in this case, except *TotalRateScaleFactor* = 1.0 because the total rates are well matched; in fact, the entire MFD is fit well, as shown in (E) Figure S5. In general, simulated rates are shifted upward in (E) Figure S4 compared with the default-parameter case (Fig. 13), and there is less average discrepancy across a wider range of target values (red-dashed

line). However, 49% of the fault-based ruptures are both triggered and have at least one supraseismogenic ancestor (compared with 35.5% above), which manifests as warmer symbol colors in [Figure S4](#) compared with [Figure 13](#). Furthermore, the recurrence interval for the Mojave S subsection is 54 years in this case.

Setting *ProbModel* = POISSON (no elastic rebound) produces very unrealistic simulations, and in fact so many aftershocks are triggered that the code never terminates. This is an expected manifestation of *CharFactor* values that exceed 1.0, because a characteristic distribution like that shown in [Figure 3a](#) has a branching ratio well above 1.0, meaning that the number of simulated events will forever increase. Elastic rebound alleviates this problem by dropping the rate, or likelihood, of supraseismogenic ruptures when such an event has already occurred.

One remaining hope for the POISSON option might be to correct all fault-section MFDs to be GR consistent, which we can enforce by setting *ApplyGRcorr* = True. However, [Field \(2012\)](#) has already shown via simpler models that this too produces unrealistic triggering statistics, in that something like 80% of large triggered ruptures simply re-rupture the parent surface, which we do not see in nature. This problem is also exemplified using UCERF3-ETAS in the [Discussion](#) section.

10-Year Probabilities (with No Scenario Mainshock). The 1000-year simulations presented above were intended to provide a better understanding of how the UCERF3-ETAS model works. Though finding perfect agreement between 1000-year simulations and model-target values might be helpful for deeming UCERF3-ETAS potentially useful, such agreement may not be a necessary condition. That is, perhaps it is asking too much to expect the model to behave well over long time periods, especially when we know that the fractional rate of spontaneous events should really vary in both time and space. Furthermore, most practical applications will be interested in what the model says about the next few days or years, and perhaps the model is more reliable over such timeframes.

We therefore examine 10-year probabilities here, rather than long-term rates, based on 10,000 simulations for each case, and again starting at the beginning of 2012 with the UCERF3 catalog provided as input (both of which are now more influential). The main metric of interest here is the probability of one or more supraseismogenic ruptures occurring on each fault section, for which historical elastic-rebound effects will also be more influential. [Figure 17](#) shows a scatter diagram of simulated versus target probabilities for each subsection, in which the simulated values represent the fraction of runs that had one or more occurrence within 10 years and for which the target values are computed from UCERF3-TD; note that the term target here does not imply expected or preferred, but rather the value we get when spatiotemporal clustering is ignored. The most striking observation is that simulated values are, on average, 25% below the target values and 33% below for the higher probability subsections (where target values are greater than 0.03). The fact that probabilities are below target values on most faults is consistent with

the fact that recent history is relatively quiet with respect to large earthquakes, because these low current values will need to average out with elevated probabilities following future large events.

The fraction of supraseismogenic ruptures that are both triggered and have at least one supraseismogenic ancestor is 19%, down from 35.5% in the default 1000-year simulations above. This implies that 10 years is not enough to see full average effects, especially given our recent history. This is borne out by the symbol colors in [Figure 17](#), which indicate the fraction of supraseismogenic ruptures that were triggered by a UCERF3-catalog earthquake, either directly or indirectly. Cases with a high rate of historical event triggering (green symbols, some of which are labeled in [Fig. 17](#)) also tend to have simulated probabilities that are well above the target values.

Some of the high outliers that are not triggered by historical events, such as the King Range and Great Valley 08 cases labeled in [Figure 17](#), represent triggering from a more active nearby fault. The reason for some of the other high outliers is not so obvious, such as for the Goldstone Lake and Mission Ridge cases labeled in [Figure 17](#); these may be related to clusters of high *CharFactor* faults producing bursts of large events, because all such cases have a high fraction of triggering from other supraseismogenic events (not shown). The lowest outliers, subsections of Mohawk Valley, represent a fault that is relatively isolated and therefore cannot be triggered by other fault-based ruptures.

Changing *ProbModel* from FULL_TD to NO_ERT (allowing supraseismogenic triggering from recent rupture zones) produces qualitatively similar results, but for which simulated probabilities are 9% below target values, on average, compared with 25% below for the FULL_TD case above.

Scenario Simulations

An OEF system will presumably be in highest demand following some seismic activity of interest. We therefore turn to implied 10-year earthquake probabilities following a variety of hypothetical scenarios, in which 10,000 simulations have been computed for each case. As mentioned above, the UCERF3 catalog was provided as input, but this does not influence the results shown here. A variety of plots are provided for the first example to give a feel for the simulations. We then restrict attention to the likelihood that each scenario will trigger events that exceed various magnitudes and compare these with probabilities obtained assuming that the total regional GR distribution applies everywhere (as quantified in [Fig. 2](#) and [Table 4](#)). Keep in mind that these scenarios are blind in that they incorporate no actual aftershock data, whereas a real OEF system will have some such data depending on the start time of the forecast.

M 5.0 Scenario on San Andreas (Mojave S). We first present results for an *M* 5.0 mainshock located right on the SAF (Mojave S) section and for default model parameters. The location of this scenario can be seen in [Figure 18](#), which shows the nucleation rate of primary aftershocks averaged

Figure 17. Simulated versus target (UCERF3-TD implied) 10-year probability of one or more suprasedsmogenic events on each subsection, based on 10,000 UCERF3-ETAS simulations using default parameter settings. Symbol colors indicate the fraction of events that were triggered by an event in the historical/instrumental UCERF3 catalog (either directly or indirectly).

over all 10,000 simulations. Figure 19 shows the average nucleation rate for all generations of aftershocks, indicating that some large and distant events have been triggered over the 10-year time period (with the more distant triggers being unique to this set of simulations, as more realizations would be needed to get true averages).

The average rate that each fault subsection participates in a suprasedsmogenic aftershock is shown in Figure 20, and the average rate at which such aftershocks nucleate from each fault section is shown in Figure 21, both of which include all generations of aftershocks. As expected, most suprasedsmogenic aftershocks occur near the M 5 scenario. Triggering does occur at more distant locations, but Figure 21 indicates that 10,000 simulations is not an adequate number to establish average nucleation rates on more distant faults.

Figure 22 shows that the distance decay implied by aftershocks of this scenario are in good agreement with that imposed/assumed, in which the increasing discrepancy beyond about 300 km is due to the fact that we are not simulating events outside the California region (Fig. 1). Likewise, the average temporal decay for primary aftershocks, shown in Figure 23, is in good agreement with that imposed/assumed, and the temporal decay for all aftershocks (also shown) has a lower rate of decay as expected by the temporal delay associated with indirect triggering.

A histogram of the maximum magnitude triggered in each simulation is shown in Figure 24 for both primary and all generations of aftershocks. The modal value is 3.8, consistent with Båth's law that the magnitude of the largest aftershock tends to

Figure 18. Average $M \geq 2.5$ nucleation rate of primary aftershocks for the M 5.0 scenario on the SAF (Mojave S) section (between subsections 10 and 11), obtained from 10,000 10-year simulations using default parameter settings (Table 2). The region was discretized into 0.02° latitude and longitude bins for this plot, and rates should be multiplied by the 10-year duration to get the average number of nucleations in each bin.

be 1.2 magnitude units below that of the mainshock (Richter, 1958). Figure 24 also shows a propensity to trigger larger events, with M 6.3–6.5 aftershocks being particularly likely in this location (for reasons discussed above).

Finally, Figure 25 shows the cumulative MFD of triggered events, or more accurately, the average number of events triggered above each magnitude (not normalized to rates, and note that we nevertheless continue referring to such plots as MFDs to avoid introducing another acronym). Several curves are shown in this figure, including the mean, mode, median, and 2.5th and 97.5th percentiles (the latter implying that 5% of simulations had values above or below these lines). The mean for primary aftershocks only is also shown (green curve), which has values that are about half the total mean at higher magnitudes. The most striking result is that the total mean is well above the median and mode, the latter two of which are similar and imply a maximum magnitude of 3.8–4.0 (again, consistent with Båth's law). The total mean is clearly being pulled high by large event triggering, whereby such events produce additional aftershocks that significantly skew the averages. In fact, the mean curve implies that we should ex-

Figure 19. Same as Figure 20, but including all generations of aftershocks for the M 5.0 SAF (Mojave S) scenario.

pect one $M \geq 5$ aftershock on average, compared with a 15% likelihood of one or more, assuming the total regional GR. The mean likelihood of triggering an $M \geq 6.3$ event is about 10%, and the likelihood of an $M \geq 7.8$ aftershock is about 0.5%. For comparison, the dashed line in Figure 25 shows the total mean number of aftershocks, assuming that the regional GR distribution in Figure 2 applies everywhere.

One way to avoid this mean-value distortion is to, instead, quantify the fraction of simulations that had one or more aftershocks exceeding each magnitude, the result of which is also shown with the red line in Figure 25 (and labeled as “Fract With 1 Or More”). This fraction is 1.0 (100%) at low magnitudes, meaning virtually all simulations included an $M \geq 2.5$ aftershock, with fractions being about 13%, 7%, and 0.5% at $M \geq 5.0$, $M \geq 6.7$, and $M \geq 7.5$, respectively. In terms of an appropriate evaluation metric, it is not yet clear that this fraction is superior to the mean value, so we have quantified both but utilize mean values in the discussions that follow; the difference is relatively small at the more hazardous magnitudes anyway.

For convenience, the mean number of aftershocks expected above various magnitudes, as implied by our simulations, is summarized in Table 5 for all the scenarios examined in this article. The ratio of this value to that expected from a GR distribution, or more specifically, from the assumption

Figure 20. The average rate at which each fault section participates in any aftershock for the M 5.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2). Multiply rates by the 10-year duration to get the average number of participations.

that the total regional GR in Figure 2 applies everywhere is also listed in the table. We refer to the latter as the R_M factor, meaning the ratio of the expected number to the regional GR value at the given magnitude (or the ratio of the solid black line to the dashed black line in Fig. 25). For example, $R_{5.0} = 6.25$ for the scenario discussed above, meaning that we should expect 6.25 times more $M \geq 5.0$ aftershocks than implied by the regional GR average. The R_M factors for the other magnitude thresholds listed are $R_{5.5} = 6.64$, $R_{6.3} = 12.5$, $R_{7.0} = 7.78$, $R_{7.4} = 14.8$, and $R_{7.8} = 43.5$, all of which seem plausible given that the scenario is located right on one of the most active faults in the region.

R_M factors reflect the effective MFD being sampled in the vicinity of the mainshock, which in turn is influenced by (1) the degree of characteristicness in the long-term MFDs; (2) time-dependent probability gains implied by elastic rebound, which are presently about a factor of 1.5 on Mojave S subsections (Field *et al.*, 2015); and (3) the influence of ERT, which influences where supraseismogenic ruptures nucleate from. Tracking the exact influence of each of these is a challenge, given the spatial variability in the long-term MFDs (e.g., as characterized by CharFactor values; Fig. 9c) and the

Figure 21. The average rate at which supraseismogenic aftershocks of any generation nucleate from each fault section for the M 5.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2). Multiply rates by the 10-year duration to get the average number of nucleations.

fact that the occurrence of one event can change the likelihood that others will follow (by virtue of elastic rebound).

Other San Andreas (Mojave S) Scenarios. Simulations were conducted for other scenarios involving SAF (Mojave S) subsections. Specifically, mainshock magnitudes of 5.5, 6.3, 7.0, 7.4, and 7.8 were applied, and both FULL_TD and NO_ERT values for the *ProbModel* parameter were considered. Scenarios with $M > 6.3$ are modeled as supraseismogenic fault-based ruptures, and those with $M < 6.3$ are treated as point sources. One of the M 6.3 scenarios was treated as a supraseismogenic fault-based rupture, and another was treated as a point source for comparison. Plots equivalent to those presented for the M 5 scenario above are also available in [Data and Resources](#) for the other scenarios discussed in this article; in particular, the maps showing the rate of $M \geq 2.5$ primary aftershocks are a good way to see the location/rupture surface of each scenario, and the cumulative MFD plots are a more complete representation of the data summarized in Table 5.

Results for the M 5.5 scenario are qualitatively similar to those for the M 5.0 case, except that R_M factors are reduced

Figure 22. The distance decay for aftershocks of the M 5.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2). The plot includes all generations of aftershocks, but distance is measured relative to the immediate parent of each aftershock. The black and gray solid lines represent the mean and 95% confidence of the mean, respectively, and the dashed gray line represents that imposed/assumed in the model.

by about 35% due to the increased effect of ERT (a larger volume from which supraseismogenic ruptures cannot rupture). For the M 6.3 scenario treated as a point source (non-fault-based rupture), R_M factors range from a value of 1.55 to 8.12, reflecting even further reductions due to ERT. Values for the M 6.3 scenario treated as a fault-based rupture are similar, ranging from 1.47 to 6.26.

Perhaps the most striking result is that R_M factors are below 1.0 for the M 7 scenario, ranging from 0.13 to 0.69, depending on the magnitude; the mean cumulative MFD for this scenario is also shown in Figure 26. The increasing influence of ERT at greater magnitude is influential here, but apparently any remaining elastic-rebound probability gains and characteristicness are not making up for this. R_M factors are even lower for the higher magnitude scenarios, ranging from 0.05 to 0.58 for the M 7.4 mainshock and from 0.01 to 0.44 for the M 7.8 event. As we go to higher magnitudes, elastic-rebound effects are further suppressing the possibility of triggered events.

Table 5 also lists the results for the NO_ERT *ProbModel* option, in which supraseismogenic aftershocks are allowed to be triggered from within the rupture zone of the mainshocks. As expected, R_M values are larger, typically by a factor of 5 compared with the FULL_TD simulations. The smallest change is for the M 7.8 scenario, in which $R_{5.0}$ went from 0.44 to 0.47, and the biggest change is for the M 7.4 scenario, in which $R_{7.8}$ went from 0.05 to 1.72. For the M 7 scenario, expected numbers generally go from being less than GR for the FULL_TD case to being greater than GR for the NO_ERT case (see also Fig. 26). However, numbers generally stay below GR for the M 7.4 and 7.8 scenarios.

San Andreas (Peninsula) Scenarios. Scenarios involving SAF (Peninsula) sections near San Francisco were also

Figure 23. The temporal decay for aftershocks of the M 5.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2). Results are plotted separately for primary aftershocks and all aftershocks (as labeled). The black and gray solid lines represent the mean and 95% confidence of the mean, respectively, and the dashed gray line represents that imposed/assumed for primary aftershocks.

Figure 24. Histogram of the maximum aftershock magnitude for M 5.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2).

investigated, including an M 5.5 point source and M 6.3 and 7.0 fault-based ruptures. In addition to being of practical concern, this fault is also of interest by virtue of having a CharFactor of ~ 0.55 (it is anticharacteristic with respect to $M \geq 6.3$ rates; Fig. 4c) and by having a time-dependent probability gain of ~ 0.64 , meaning it is statistically underdue (see fig. 4b of Field *et al.*, 2015). These facts, coupled with the influence of ERT, predict R_M factors below 1.0. Indeed, for the FULL_TD ProbModel option, the range of values in Table 5 is a low of $R_{6.3} = 0.04$ for the M 6.3 scenario and a high of $R_{5.0} = 0.71$ for the M 5.5 scenario. Switching to NO_ERT increases expected numbers by up to a factor of 6

Figure 25. Cumulative number of aftershocks greater than each magnitude for the M 5.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2). See the legend for a description of the various curves. The Mean, Mode, Median, and fractile curves correspond to all aftershocks (all generations); Primary Mean is the mean for primary aftershocks only; and the Regional GR Mean is the total expected number, assuming the total regional GR distribution in Figure 2 applies everywhere. Fract With 1 or More represents the fraction of simulations that had one or more aftershocks greater than or equal to the given magnitude.

(e.g., the likelihood of the M 6.3 scenario triggering $M \geq 7.8$ events), with a maximum of $R_{7.4} = 0.94$ for the M 5.5 scenario. That is, none of the R_M factors currently exceeds 1.0, not even for the NO_ERT case, meaning that conditional triggering probabilities are currently below regional GR values. In the future, the elastic-rebound probability gain may go from the current value of 0.64 to something like 2.0, but even then only some R_M factors will exceed 1.0 and only up to a value of about 3.0.

Surprise Valley M 5.5 Scenario. As noted above, the Surprise Valley fault has a CharFactor value of 45 (Fig. 4c), which is second only to a subsection of the Robinson Creek fault. Simulations of the latter turned out to be less dramatic than for Surprise Valley due to other influences, so we demonstrate the influence of extreme CharFactor values here using an M 5.5 scenario located right on the Surprise Valley fault. R_M factors range from 2.56 to 4.74 for the FULL_TD case and from 10.2 to 20.5 for the NO_ERT case. Perhaps the most extreme consequence is that the M 5.5 mainshock has a 0.53 probability of triggering an $M \geq 6.3$ event for the NO_ERT case (or a 0.1 chance for the FULL_TD case). This result alleviates any concerns that extreme CharFactor values might lead to unphysical conditional triggering probabilities, because even for this most extreme case we are not saying that an M 6.3 event is certain to happen. This conclusion depends on setting *ApplyGridSeisCorr* = True, because unreasonable conditional probabilities do result without this correction, as discussed above. Note also that the probability

Table 5

Average Expected Number ($N \geq M$) and the Ratio with Respect to Assuming a GR Distribution (R_M) Implied by the Various Mainshock Scenario Simulations

Scenario*	ProbModel	$N \geq 5.0$	$R_{5.0}$	$N \geq 5.5$	$R_{5.5}$	$N \geq 6.3$	$R_{6.3}$	$N \geq 7.0$	$R_{7.0}$	$N \geq 7.4$	$R_{7.4}$	$N \geq 7.8$	$R_{7.8}$
SAF Mojave S M 5.0	FULL_TD	1.0317	6.25	0.3464	6.64	0.1034	12.5	0.0122	7.78	0.0080	14.80	0.0054	43.5
SAF Mojave S M 5.0	NO_ERT	2.8180	17.1	0.9421	18.1	0.2626	31.7	0.0423	27.0	0.0241	44.6	0.0148	119.
SAF Mojave S M 5.5	FULL_TD	2.1791	4.18	0.7261	4.40	0.2066	7.90	0.0248	5.00	0.0151	8.83	0.0110	28.0
SAF Mojave S M 5.5	NO_ERT	8.6379	16.5	2.8433	17.2	0.7338	28.1	0.1210	24.4	0.0760	44.5	0.0502	128.
SAF Mojave S M 6.3 PtSrc	FULL_TD	5.5615	1.69	1.8208	1.75	0.4359	2.64	0.0484	1.55	0.0286	2.65	0.0201	8.12
SAF Mojave S M 6.3 PtSrc	NO_ERT	29.031	8.82	9.5457	9.17	2.4509	14.8	0.3980	12.7	0.2434	22.6	0.1630	65.8
SAF Mojave S M 6.3 FSS	FULL_TD	6.1376	1.86	2.1504	2.07	0.6988	4.24	0.0461	1.47	0.0236	2.19	0.0155	6.26
SAF Mojave S M 6.3 FSS	NO_ERT	26.093	7.93	8.4805	8.15	2.0705	12.5	0.3741	12.0	0.2241	20.8	0.1461	59.0
SAF Mojave S M 7.0	FULL_TD	11.423	0.69	3.1868	0.61	0.1510	0.18	0.0211	0.13	0.0096	0.18	0.0050	0.40
SAF Mojave S M 7.0	NO_ERT	24.218	1.47	6.9864	1.34	0.6886	0.83	0.2082	1.33	0.1539	2.85	0.1096	8.83
SAF Mojave S M 7.4	FULL_TD	23.956	0.58	6.4953	0.50	0.1508	0.07	0.0197	0.05	0.0079	0.06	0.0016	0.05
SAF Mojave S M 7.4	NO_ERT	31.508	0.76	8.7965	0.67	0.4674	0.23	0.1420	0.36	0.0840	0.62	0.0537	1.72
SAF Mojave S M 7.8	FULL_TD	46.032	0.44	12.841	0.39	0.4892	0.09	0.0400	0.04	0.0112	0.03	0.0009	0.01
SAF Mojave S M 7.8	NO_ERT	48.481	0.47	13.708	0.42	0.7399	0.14	0.1377	0.14	0.0410	0.12	0.0060	0.08
SAF Peninsula M 5.5	FULL_TD	0.3693	0.71	0.1008	0.61	0.0050	0.19	0.0016	0.32	0.0009	0.53	0.0001	0.25
SAF Peninsula M 5.5	NO_ERT	0.3865	0.74	0.1082	0.66	0.0079	0.30	0.0026	0.52	0.0016	0.94	0.0000	0.00
SAF Peninsula M 6.3	FULL_TD	2.1673	0.66	0.5889	0.57	0.0339	0.21	0.0123	0.39	0.0072	0.67	0.0001	0.04
SAF Peninsula M 6.3	NO_ERT	2.2387	0.68	0.6015	0.58	0.0391	0.24	0.0139	0.44	0.0079	0.73	0.0006	0.24
SAF Peninsula M 7.0	FULL_TD	10.607	0.64	2.8391	0.54	0.1513	0.18	0.0473	0.30	0.0283	0.52	0.0028	0.23
SAF Peninsula M 7.0	NO_ERT	11.129	0.67	2.9995	0.57	0.1818	0.22	0.0562	0.36	0.0296	0.55	0.0040	0.32
Surprise Valley 5.5	FULL_TD	1.6443	3.15	0.7817	4.74	0.1009	3.86	0.0127	2.56	0.0000	0.00	0.0000	0.00
Surprise Valley 5.5	NO_ERT	5.3264	10.2	2.3494	14.2	0.5334	20.4	0.1018	20.5	0.0002	0.12	0.0000	0.00
San Jacinto Borrego M 5.5	FULL_TD	0.3056	0.59	0.0853	0.52	0.0023	0.09	0.0007	0.14	0.0002	0.12	0.0000	0.00
San Jacinto Borrego M 5.5	NO_ERT	0.3655	0.70	0.1052	0.64	0.0048	0.18	0.0012	0.24	0.0007	0.41	0.0002	0.51
Central Valley M 5.0	FULL_TD	0.1268	0.77	0.0383	0.73	0.0062	0.75	0.0002	0.13	0.0000	0.00	0.0000	0.00
Central Valley M 5.0	NO_ERT	0.1126	0.68	0.0321	0.62	0.0039	0.47	0.0003	0.19	0.0000	0.00	0.0000	0.00
Bombay Beach M 4.8	FULL_TD	0.2091	2.01	0.0581	1.76	0.0085	1.62	0.0034	3.44	0.0020	5.86	0.0007	8.30
Bombay Beach M 4.8	NO_ERT	0.2407	2.31	0.0716	2.17	0.0126	2.41	0.0044	4.45	0.0023	6.74	0.0006	7.03

*PtSrc means treated as point source (gridded seismicity), and FSS means treated as a fault-system-solution rupture (fault based); SAF, San Andreas fault.

Figure 26. Mean cumulative number of aftershocks greater than each magnitude (solid lines) for the M 7.0 SAF (Mojave S) scenario, obtained from 10,000 10-year simulations using default parameter settings (Table 2). Results are shown separately for *ProbModel* set as full-time dependence (FULL_TD) versus NO_ERT (in which ERT means elastic-rebound triggering; as labeled). The dashed line represents the result assuming the total regional GR MFD in Figure 2 applies everywhere, and the dotted lines represent modal values from the simulations.

of having an M 5.5 mainshock in this area is very low in the first place, which is why it has a high CharFactor value.

San Jacinto (Borrego) M 5.5 Scenario. This case exemplifies the consequence of having a particularly low CharFactor, because the value is 0.12 at the hypocenter of this scenario (and the time-dependent probability gain is near 1.0). The $R_{6.3}$ value is 0.09 and 0.18 for the FULL_TD and NO_ERT cases, respectively. In other words, the conditional probability of triggering a supraseismogenic rupture, given something smaller nearby, is well below the GR average. This result reflects a relatively high rate of microseismicity near this fault, with the low CharFactor value effectively preventing these numerous events from over-triggering supraseismogenic ones (relative to the rate of the latter implied by the long-term model).

Central Valley M 5.0 Scenario. This case exemplifies a mainshock that is far away from any modeled faults. Here, there should be minimal difference between the FULL_TD and NO_ERT cases because these parameters only influence supraseismogenic on-fault triggering. Simulations imply that $R_{5.0} = 0.72$, with values decreasing at high magnitudes, which is consistent with the fact that the off-fault MFD (blue curve in Fig. 2) tapers more quickly than the total MFD (black curve) at higher magnitudes. The number of $M \geq 5$ primary aftershocks is consistent with that in GR, but the number of $M \geq 5$ events triggered in subsequent generations is deficient, due to fewer larger events being triggered by virtue of the greater tapering. This is precisely why the long-term simulations underpredict the target rate of events in off-fault areas (Fig. 12b), because the fractional rate of

spontaneous versus triggered events should really be adjusted according to the off-fault MFD.

Bombay Beach M 4.8 Scenario. This scenario corresponds to one studied extensively by Michael (2012a) in terms of its likelihood of triggering an $M \geq 7$ earthquake on the Coachella section of the SAF. The M 4.8 event, which actually occurred on 24 March 2009, is located about 4 km from the fault-surface representation in UCERF3. The probability gain in this area is 1.9 according to UCERF3-TD, and the CharFactor value is 0.65 for the closest subsection (and note that the CharFactor went from above 1.0 to below it when we set *ApplyGridSeis-Corr* = True; Fig. 4a vs. Fig. 5b,c). Here, the metric of interest is the likelihood of triggering an $M \geq 7$ event within 3 days, for which Michael (2012a) found a 0.00035–0.013 range, depending on whether a GR or characteristic MFD is assumed. The value we obtained is 0.0012, meaning that this fraction of the simulations exhibited one or more triggered $M \geq 7$ events within the first 3 days. Both the FULL_TD and NO_ERT *ProbModel* options produced similar values for this relatively small mainshock. The R_M factors based on 10-year forecasts range from a low of $R_{7.0} = 1.62$ to a high of $R_{7.8} = 8.3$ (Table 5).

Discussion

Main Challenges

The goal of this study has been to bring fault-based information into OEF, which has traditionally been based solely on GR and Omori–Utsu statistics (Jordan *et al.*, 2011). As stated in the Introduction, UCERF3-ETAS essentially represents an elaborate implementation of the Michael (2012a) generalized clustering model in that it effectively samples earthquakes according to the spatial distribution of MFDs implied by UCERF3-TD. In so doing, we explicitly include information on fault proximity, activity rate, and susceptibility from an elastic-rebound perspective.

A number of significant challenges have arisen in constructing the model, including the following:

1. How fault-based nucleation rates transition into the surrounding region, especially given uncertainties in the fault surfaces themselves,
2. How and if elastic rebound suppresses the likelihood of other ruptures nucleating from within the rupture zone of previous events,
3. That the degree of characteristicness implied on faults varies widely, and that much of this may represent uncertainty rather than reality,
4. That the previously inferred long-term rate of smaller events near some faults is not consistent with the number of aftershocks expected from larger ruptures on that fault, implying that CharFactors must be biased high,
5. That if one wants to match the long-term model, then the fractional rate of spontaneous versus triggered events in the ETAS model must not only vary in time but also in space because the nucleation MFDs vary spatially.

6. That one must consider the depth distribution of seismicity, as well as the finite thickness of the seismogenic layer, when matching the distance decay of aftershocks.

Although most of us thought that UCERF3-ETAS would be relatively straightforward to construct, in retrospect it is surprising that the model works as well as it does, given all the challenges encountered and the approximate approaches that were used to solve them. The question is, however, whether the model is good enough to be useful, or useful enough to be worth operationalizing.

Potential Model Usefulness

UCERF3-ETAS simulations look realistic in terms of the overall productivity and spatiotemporal decay of aftershocks, and the results maintain a plausible degree of elastic-rebound predictability (Fig. 16). However, the 1000-year simulations do exhibit differences from long-term model rates, some of which can be attributed to a need for a spatially dependent fractional rate of spontaneous earthquakes, which is not presently supported in the model.

Supraseismogenic ruptures can only occur in one of three ways: (1) spontaneously, (2) triggered by a smaller nearby event that is not itself a descendant of any supraseismogenic rupture, or (3) triggered by a supraseismogenic rupture, either directly or indirectly. Discrepancies between simulated and model-expected fault-section participation rates correlate most strongly with the degree of triggering from other supraseismogenic fault-based ruptures, with the highest outliers representing lesser faults adjacent to, and consequently being triggered by, the much more active SAF (e.g., Figs. 12 and 13). Likewise, a well-isolated fault will exhibit undersimulated rates by virtue of only rarely being triggered by a neighboring fault. Although these more extreme outliers are easily explained, we cannot yet claim to understand every discrepancy that the model produces.

The triggering of supraseismogenic ruptures is influenced by the interaction of several factors, including the population of nearby faults, the degree of characteristicness on those faults, the elastic-rebound-implied probability of each rupture, and whether elastic rebound further restricts in which triggered ruptures can nucleate. At this point, it is difficult to say whether long-term simulation discrepancies are a concern with respect to potential usefulness. Either way, 10-year probabilities may be more reliable, because the latter seem to be reasonable when based on the historical/instrumental catalog as input (above average triggering probabilities on faults that are near recent large events and below average elsewhere). Applying a spatially variable fractional rate of spontaneous versus triggered events would certainly help to reduce the remaining discrepancies; how much so remains to be seen, and it begs the question of how close is close enough, which should probably be answered before any such complexity is added to the model, especially when the revised model will still be an approximation.

Perhaps the more troubling aspect of UCERF3-ETAS is the conditional triggering probabilities implied by scenarios. At first blush, one might expect the number of expected aftershocks to always be greater than regional GR-implied values when the mainshock is near a more active fault and below average when not near any known fault (e.g., reflecting the fact that CEPEC convenes when a magnitude ~ 5 earthquake occurs near the SAF, as noted in the [Introduction](#)). This is not borne out by UCERF3-ETAS, as many of the fault-based scenarios imply probabilities below the GR value, including an $M 7$ event on the Mojave S section of the SAF (a fault considered to be locked and loaded).

As noted by [Michael \(2012a\)](#), the degree of characteristicness has a first-order influence on conditional triggering probabilities near a fault. For example, the San Jacinto fault has CharFactor values well below 1.0, thereby implying conditional probabilities that are lower than GR. CharFactor values vary widely throughout the region, even along a given fault (e.g., Fig. 9c), which means that we will have commensurate variability in the conditional triggering probabilities implied by an event.

If all the CharFactor values are real (not due to uncertainties), then we need to retain them in order for the model to correctly reproduce both long-term rates and short-term probabilities. For example, a high CharFactor is needed precisely because small events are rare in that area, and the high value compensates accordingly to produce the correct rate of large events (or the correct portion triggered by smaller events). The same goes for low CharFactor faults. For example, if we ignored the low value on the San Jacinto fault and instead set $\text{CharFactor} \geq 1.0$, then the conditional probability for a large event being triggered from the next small earthquake might not look unreasonable. However, the aggregate conditional probability implied by several such events, which would surely occur given the high rate of microseismicity, might approach an unreasonably high value near 1.0, which could be a basis for rejecting the model.

The greater concern is that CharFactor values are highly uncertain and that we are consequently propagating noise. If we assume that supraseismogenic rates are correct but that subseismogenic rates represent noise, then high CharFactor values will overstate conditional probabilities (because there are really more small events than we thought) and vice versa for low CharFactor values. The problem with this interpretation is that our small event rates are based on a model that is one of the most skillful with respect to forecasting the location of future small events ([Helmstetter et al., 2007](#); [Zechar et al., 2013](#)), so we know the rates cannot represent complete noise.

The more likely situation is that we have some uncertainty in the nucleation rate of both sub- and supraseismogenic ruptures and that these are getting propagated into our conditional probabilities. However, resolving these uncertainties is not likely to convert the conditional triggering probabilities near all faults to a value greater than the GR equivalent, especially if ERT is a reality. The latter remains

an open question, because we do not yet have a basis for rejecting either the FULL_TD or NO_ERT *ProbModel* options. Furthermore, even if ruptures cannot nucleate from the zone of a recent rupture, allowing them to do so may still provide a better statistical proxy overall. We also do not really know what large event-triggering statistics really are given the paucity of observations, so expecting faults to have greater conditional probabilities than GR may be naïve.

This brings us to what may be the biggest conceptual problem with UCERF3-ETAS, which applies to the Michael (2012a) generalized clustering model as well. That is, we are using a long-term MFD to define conditional triggering probabilities, when in reality that MFD is extremely time dependent. For example, a fault will exhibit its highest rate of seismicity immediately following a large event, at which time the rate or likelihood of another large event on the fault is lowest due to elastic rebound. As time goes on, the rate or probability of another big one steadily increases, while the rate of smaller events decays back down indefinitely, perhaps to some background rate. Of course, the occurrence of nearby events will modulate these rates somewhat, but the overall MFD time evolution should remain on average. If this behavior represents reality, which UCERF3-ETAS assumes it does, then why would conditional probabilities depend on some hypothetical long-term MFD that may never apply at any point in time? Trying to use the fully time-dependent MFD in defining conditional triggering probabilities would not help, because it would significantly lower the likelihood of one fault-based rupture triggering another one farther down the fault (because aftershocks from the first spilling onto the surface of the second would temporarily lower the CharFactor in that area, making the triggering less likely).

Perhaps UCERF3-ETAS can nevertheless be a useful proxy, especially if properly tuned in terms of, for example, changing CharFactor values to produce acceptable conditional probabilities. However, we do not presently know what values to tune the model to, not even the acceptable range, which is again due to the paucity of large event-triggering data. This is a general problem that will plague the evaluation of any candidate OEF model. We also want to avoid having to customize UCERF3-ETAS for every earthquake that occurs, especially if doing so requires expert judgment.

Another question is whether UCERF3-ETAS might be useful with respect to OEF for induced seismicity. The answer would seem to be yes, in that the model makes no distinction between natural and induced earthquakes, hence if the rates of the latter are changed, then so too will consequent triggering probabilities. The question is whether UCERF3-ETAS is the best approach in terms of an OEF for induced events. For example, Llenos and Michael (2013) found that having a temporally varying rate of spontaneous events provides a clear improvement when using ETAS to model induced events, especially when swarm-type behavior is present; UCERF3-ETAS does not presently support such rate variation. However, if one wants to consider proximity to

faults, then something like UCERF3-ETAS will be needed to model induced seismicity.

Finally, potential usefulness will need to be ascertained in the context of specific applications, such as public preparedness, emergency response, building inspection and tagging, building-code adjustment, setting earthquake-insurance rates, and pricing insurance-linked securities such as catastrophic bonds and reinsurance. UCERF3-ETAS may be an adequate approximation for some of these uses but perhaps not for others. An interesting question is what tests will be necessary and sufficient in terms of deeming usefulness. Running the model for all significant historical California ruptures is an obvious task that should be pursued, but it is not yet clear exactly how we would deem success or failure. Furthermore, we have so far struggled to produce a single viable UCERF3-ETAS logic-tree branch, whereas multiple sources of uncertainty will need to be explored and quantified. Operationalization will also present unique challenges, perhaps including a need for real-time access to high-performance computing. In the interest of brevity we do not discuss these issues further here, but suffice it to say, considerable work remains with respect to ascertaining usefulness.

Scientific Implications

One primary contribution of this study is in providing further evidence that one cannot combine statistical triggering models, such as ETAS, with finite-fault models without also including elastic rebound. This assertion, which was first made by Field (2012) using relatively simple models, is supported here with an example from UCERF3-ETAS. We already explained how characteristic MFDs combined with the POISSON *ProbModel* option, imply a branching ratio much greater than 1.0 and therefore producing runaway sequences. What remains to be exemplified here is that even GR-consistent faults also suffer problems. Consider the *M* 7 Mojave S scenario discussed above but in which *ProbModel* = POISSON and *ApplyGRcorr* = True (making all faults GR-consistent, as exemplified in Fig. 3b). Figure 27 shows that, according to this model, if a supraseismogenic rupture is triggered as a primary aftershock, it has a 77% chance of having more than half of its rupture surface overlap with that of the parent event. We just do not see this kind of re-rupture occurring in nature, at least not anywhere near 77% of the time. Furthermore, and as discussed above, van der Elst and Shaw (2015) have presented observational evidence that aftershocks as large or larger than the parent nucleate almost exclusively in the outer regions of the parent aftershock zone, giving further evidence of some elastic-relaxation process.

Though we are very confident that this need for elastic rebound will stand the test of time, there are legitimate questions about whether UCERF3-ETAS has gotten it right. The recurrence interval distributions shown in Figures 14 and 16 exhibit a plausible degree of elastic-rebound predictability. One potentially important metric is the relative height of

Figure 27. For the SAF Mojave S M 7 scenario (thick black line), this figure shows all possible suprasedimental aftershocks that have at least half of their rupture surface overlapping that of the parent (offset to the northeast for visibility). For *ProbModel* = POISSON, the color indicates the conditional probability of each being sampled as a primary aftershock, given some suprasedimental rupture has been triggered. The total conditional probability that one of these ruptures will be sampled as a primary aftershock is 77%.

the first bin (the fraction of very short recurrence intervals), which is to some degree a measure of the spatial overlap of ruptures triggered on adjacent sections of a fault. Is this the correct amount of potential overlap? Should it be more or less? We do not really know, and it seems that our best hope for addressing such questions is to try to collect relevant observational data and to study the behavior implied by physics-based simulators.

Another scientific conclusion is that some degree of characteristicness is required on at least some faults. That is, if you believe that the conditional probability of triggering a large event should be significantly higher if you are near a seemingly capable fault, then you are also saying that the fault has a relatively characteristic MFD, at least in the context of the generalized clustering model (Michael, 2012a). What this study adds is that you will need even more characteristicness if ERT is a reality, because the suppression of triggering from the rupture zone of the parent significantly lowers the consequent large event-triggering probabilities (illustrated in Fig. 7).

Although these scientific inferences seem robust, we also acknowledge that the approach we have taken here may be fundamentally flawed. Perhaps once you exceed the scale of the seismogenic layer, and thereby enter the range of magnitudes that cause the most damage, maybe ETAS statistics are no longer reliable or perhaps maybe not at the level of specificity represented in UCERF3-ETAS. Is our model thereby misleading? One key to addressing this question will be a comparison with more physics-based approaches. For example, one improvement might be to use Coulomb stress-change calculations (e.g., King *et al.*, 1994), or some other reasoning, to condition which aftershocks can nucleate, as opposed to the isotropic distribution assumed here. We emphasize that the viability of any such enhancement should be tested by evaluating implied synthetic earthquake catalogs, because we hope to have demonstrated that such analyses are a powerful, if not indispensable, way to identify model inconsistencies. If we want both more physics and synthetic catalogs, then physics-based simulators that include both elastic rebound and Omori–Utsu-type triggering seem the logical way to go (e.g., RSQSim; Richards-Dinger and Dietrich, 2012). The challenge will be including off-fault seismicity in such models, as well as inferring the conditional triggering probabilities implied by smaller events (e.g., for the *M* 4.8 Bombay Beach scenario).

This study also raises some questions with respect to model testing. As mentioned, the model that specifies the long-term rate of smaller earthquakes in UCERF3-ETAS has demonstrated skill in the formal prospective tests conducted by the Collaboratory for the Study of Earthquake Predictability (CSEP; Helmstetter *et al.*, 2007; Zechar *et al.*, 2013; Helmstetter and Werner, 2014; Rhoades *et al.*, 2014). However, we had to change these rates near several faults (via the *ApplyGridSeis-Corr* parameter; Fig. 5) in order to prevent artificially high CharFactor values from giving unreasonable conditional triggering probabilities in those areas. We might also be tempted to adjust the CharFactor values elsewhere to produce more reasonable results, which would likely involve changing the long-term rate of smaller events in those areas too. The cost of these model adjustments, aimed at improving performance at the larger more hazardous magnitudes, might very well result in relatively poor performance in CSEP tests because the latter are dominated by smaller earthquakes. Clearly, one key to improved testing and/or tuning of any OEF model will be a better empirical quantification of large event-triggering statistics, especially on a global scale.

Conclusions

UCERF3-ETAS can be viewed as an attempt to either add spatiotemporal clustering to a modern fault-based model (UCERF3-TD) or to fold fault-based information into a state-of-the-art statistical seismology model (ETAS). Perhaps our most significant finding is that both elastic-rebound effects and characteristic MFDs near faults are apparently required to produce realistic behavior, indicating a potential resolution

with respect to debates on whether either one is actually needed (e.g., Kagan *et al.*, 2012; Tormann *et al.*, 2015, 2016; Bürgmann *et al.*, 2016).

The UCERF3-ETAS simulations evaluated thus far produce very plausible results, or at least we cannot yet reject those based on preferred parameter settings. Some of the results are somewhat surprising, however, at least in comparison with the naïve assumption that conditional triggering probabilities should always be greater than regional averages when a mainshock is near an active fault. One first-order influence on conditional triggering probabilities is the degree of characteristicness in MFDs near faults, which varies widely throughout California (at least according to UCERF3). Another important question is whether larger aftershocks can be triggered from within the rupture zone of the parent event, because currently viable answers to this question can change triggering likelihoods by an order of magnitude (e.g., Fig. 26). A third challenge is rectifying large uncertainties in the physical manifestation of faults with the fact that most triggering occurs within a few kilometers of the mainshock.

The bottom line is that there are many uncertainties associated with UCERF3-ETAS, and there are many aspects of the model that we could likely improve upon. However, this will always be the case with any such forecast, hence the relevant question is whether the model is good enough to be potentially useful or useful enough to be worth operationalizing. This question can only be answered in the context of actual uses, because it may be good enough for one application but not another. Any further development of UCERF3-ETAS should, therefore, at least be informed by anticipated uses.

Going forward, one key to improving OEF will be the collection of more global large event-triggering data, both in terms of tuning and testing models. Another will be the use of physics-based simulators to help, for example, narrow the viable range of elastic-rebound behavior. Finally, our results do not imply that simply ignoring faults in OEF will produce a more scientifically plausible or useful model.

Data and Resources

Plots for all the scenario simulations presented in this article are available at <http://www.WGCEP.org/UCERF3-ETAS> (last accessed November 2016). All calculations were made using OpenSHA (<http://www.OpenSHA.org>; last accessed April 2016; Field *et al.*, 2003), which in turn utilizes Generic Mapping Tools (<http://gmt.soest.hawaii.edu>; last accessed January 2012) and JFree-Chart (<http://www.jfree.org/jfreechart/>; last accessed March 2012) for making plots. The unpublished manuscript by N. J. van der Elst (2016). “Including orphaned aftershocks of long-past mainshocks in the ETAS model”.

Acknowledgments

The Southern California Earthquake Center (SCEC) Community Modeling Environment, University of Southern California Center for High-Performance Computing and Communications, and the Texas Advanced

Computing Center provided access to the high-performance computing. This study was sponsored by the California Earthquake Authority, the U.S. Geological Survey (USGS), the USGS John Wesley Powell Center for Analysis and Synthesis, and the SCEC. SCEC is supported in part by the National Science Foundation under Cooperative Agreement EAR-1033462 and by the USGS under Cooperative Agreement G12AC20038. We are also grateful to Andrea Llenos and two anonymous *Bulletin of the Seismological Society of America* reviewers for their very thoughtful and helpful comments. The SCEC Contribution Number for this article is 7163. Any use of trade, product, or firm names is for descriptive purposes only and does not imply endorsement by the U.S. Government.

References

- Agnew, D. C., and L. M. Jones (1991). Prediction probabilities from foreshocks, *J. Geophys. Res.* **96**, 11,959–11,971.
- Bürgmann, R., N. Uchide, Y. Hu, and T. Matsuzawa (2016). Tohoku rupture reloaded? *Nat. Geosci.* **9**, 183–184.
- Felzer, K. R. (2013). Appendix K: The UCERF3 earthquake catalog, *U.S. Geol. Surv. Open-File Rept. 2013-1165-K*, and *California Geol. Surv. Special Rept. 228-K*.
- Felzer, K. T., T. W. Becker, R. E. Abercrombie, G. Ekström, and J. R. Rice (2002). Triggering of the 1999 M_w 7.1 Hector Mine earthquake by aftershocks of the 1992 M_w 7.3 Landers earthquake, *J. Geophys. Res.* **107**, 2190, doi: [10.1029/2001JB000911](https://doi.org/10.1029/2001JB000911).
- Field, E. H. (2012). Aftershock statistics constitute the strongest evidence for elastic relaxation in large earthquakes—Take 2, presented at the 2012 Meeting of the SSA, San Diego, California, 17–19 April.
- Field, E. H. (2015). Computing elastic-rebound-motivated earthquake probabilities in unsegmented fault models: A new methodology supported by physics-based simulators, *Bull. Seismol. Soc. Am.* **105**, 544–559, doi: [10.1785/0120140094](https://doi.org/10.1785/0120140094).
- Field, E. H., R. J. Arrowsmith, G. P. Biasi, P. Bird, T. E. Dawson, K. R. Felzer, D. D. Jackson, K. M. Johnson, T. H. Jordan, C. Madden, *et al.* (2014). Uniform California Earthquake Rupture Forecast, version 3 (UCERF3)—The time-independent model, *Bull. Seismol. Soc. Am.* **104**, 1122–1180, doi: [10.1785/0120130164](https://doi.org/10.1785/0120130164).
- Field, E. H., R. J. Arrowsmith, G. P. Biasi, P. Bird, T. E. Dawson, K. R. Felzer, D. D. Jackson, K. M. Johnson, T. H. Jordan, C. Madden, *et al.* (2015). Long-term, time-dependent probabilities for the Third Uniform California Earthquake Rupture Forecast (UCERF3), *Bull. Seismol. Soc. Am.* **105**, 511–543, doi: [10.1785/0120140093](https://doi.org/10.1785/0120140093).
- Field, E. H., T. E. Dawson, K. R. Felzer, A. D. Frankel, V. Gupta, T. H. Jordan, T. Parsons, M. D. Petersen, R. S. Stein, R. J. Weldon II, and C. J. Wills (2009). Uniform California Earthquake Rupture Forecast, Version 2 (UCERF 2), *Bull. Seismol. Soc. Am.* **99**, 2053–2107, doi: [10.1785/0120080049](https://doi.org/10.1785/0120080049).
- Field, E. H., T. H. Jordan, and C. A. Cornell (2003). OpenSHA: A developing community-modeling environment for seismic hazard analysis, *Seismol. Res. Lett.* **74**, 406–419.
- Field, E. H., T. H. Jordan, L. M. Jones, A. J. Michael, M. L. Blanpied, and Other Workshop Participants, (2016). The potential uses of operational earthquake forecasting, *Seismol. Res. Lett.* **80**, 1–10, doi: [10.1785/0220150174](https://doi.org/10.1785/0220150174).
- Gardner, J. K., and L. Knopoff (1974). Is the sequence of earthquakes in southern California with aftershocks removed Poissonian? *Bull. Seismol. Soc. Am.* **64**, 1363–1367.
- Gerstenberger, M. C., G. H. McVerry, D. A. Rhoades, and M. W. Stirling (2014). Seismic hazard modeling for the recovery of Christchurch, New Zealand, *Earthq. Spectra* **30**, 17–29.
- Gerstenberger, M. C., S. Wiemer, L. M. Jones, and P. A. Reasenber (2005). Real-time forecasts of tomorrow’s earthquakes in California, *Nature* **435**, 328–331.
- Gutenberg, B., and C. F. Richter (1944). Frequency of earthquakes in California, *Bull. Seismol. Soc. Am.* **34**, 185–188.
- Hanks, T. C., and W. H. Bakun (2008). M-log A observations of recent large earthquakes, *Bull. Seismol. Soc. Am.* **98**, 490.

- Hardebeck, J. L. (2013). Appendix S: Constraining epidemic type aftershock sequence (ETAS) parameters from the Uniform California Earthquake Rupture Forecast, Version 3 catalog and validating the ETAS model for magnitude 6.5 or greater earthquakes, *U.S. Geol. Surv. Open-File Rept. 2013-1165-S*, and *California Geol. Surv. Special Rept. 228-S*.
- Hardebeck, J. L., K. R. Felzer, and A. J. Michael (2008). Improved tests reveal that the accelerating moment release hypothesis is statistically insignificant, *J. Geophys. Res.* **113**, no. B08310, doi: [10.1029/2007JB005410](https://doi.org/10.1029/2007JB005410).
- Helmstetter, A., and D. Sornette (2003). Importance of direct and indirect triggered seismicity in the ETAS model of seismicity, *Geophys. Res. Lett.* **30**, no. 11, 1576, doi: [10.1029/2003gl017670](https://doi.org/10.1029/2003gl017670).
- Helmstetter, A., and M. J. Werner (2014). Adaptive smoothing of seismicity in time, space, and magnitude for time-dependent earthquake forecasts for California, *Bull. Seismol. Soc. Am.* **104**, 809–822.
- Helmstetter, A., Y. Y. Kagan, and D. D. Jackson (2007). High-resolution time-independent grid-based forecast for $M \geq 5$ earthquakes in California, *Seismol. Res. Lett.* **78**, 78–86.
- Jordan, T. H., and L. M. Jones (2010). Operational earthquake forecasting: some thoughts on why and how, *Seismol. Res. Lett.* **81**, no. 4, 571–574.
- Jordan, T. H., Y.-T. Chen, P. Gasparini, R. Madariaga, I. Main, W. Marzocchi, G. Papadopoulos, G. Sobolev, K. Yamaoka, and J. Zschau (2011). Operational earthquake forecasting: state of knowledge and guidelines for implementation, final report of the International Commission on Earthquake Forecasting for Civil Protection, *Ann. Geophys.* **54**, no. 4, 315–391, doi: [10.4401/ag-5350](https://doi.org/10.4401/ag-5350).
- Jordan, T. H., W. Marzocchi, A. J. Michael, and M. C. Gerstenberger (2014). Operational earthquake forecasting can enhance earthquake preparedness, *Seismol. Res. Lett.* **85**, 955–959.
- Kagan, Y. Y., D. D. Jackson, and R. J. Geller (2012). Characteristic earthquake model, 1884–2011, RIP, *Seismol. Res. Lett.* **83**, 951–953, doi: [10.1785/0220120107](https://doi.org/10.1785/0220120107).
- Kaiser, A. E., C. Holden, R. J. Beavan, R. D. Beetham, R. A. Benites, A. Celentano, D. Collet, W. J. Cousins, M. Cubrinovski, G. D. Dellow, et al. (2012). The M_w 6.2 Christchurch earthquake of February 2011: Preliminary report, *New Zeal. J. Geol. Geophys.* **55**, no. 1, 67–90, doi: [10.1080/00288306.2011.641182](https://doi.org/10.1080/00288306.2011.641182).
- King, G. C. P., R. S. Stein, and J. Lin (1994). Static stress changes and the triggering of earthquakes, *Bull. Seismol. Soc. Am.* **84**, 935–953.
- Llenos, A. L., and A. J. Michael (2013). Modeling earthquake rate changes in Oklahoma and Arkansas: Possible signatures of induced seismicity, *Bull. Seismol. Soc. Am.* **103**, 2850–2861, doi: [10.1785/0120130017](https://doi.org/10.1785/0120130017).
- Michael, A. J. (2012a). Fundamental questions of earthquake statistics and estimation of earthquake probabilities from possible foreshocks, *Bull. Seismol. Soc. Am.* **102**, 2547–2562, doi: [10.1785/0120090184](https://doi.org/10.1785/0120090184).
- Michael, A. J. (2012b). Do aftershock probabilities decay with time? *Seismol. Res. Lett.* **83**, 630–632, doi: [10.1785/0220120061](https://doi.org/10.1785/0220120061).
- Marzocchi, W., A. M. Lombardi, and E. Casarotti (2014). The establishment of an operational earthquake forecasting system in Italy, *Seismol. Res. Lett.* **85**, 961–969, doi: [10.1785/0220130219](https://doi.org/10.1785/0220130219).
- Nanjo, K. Z., H. Tsuruoka, S. Yokoi, Y. Ogata, G. Falcone, N. Hirata, Y. Ishigaki, T. H. Jordan, K. Kasahara, K. Obara, et al. (2012). Predictability study on the aftershock sequence following the 2011 Tohoku-Oki, Japan, earthquake: First results, *Geophys. J. Int.* **191**, 653–658.
- Ogata, Y. (1988). Statistical models of point occurrences and residual analysis for point processes, *J. Am. Stat. Assoc.* **83**, 9–27.
- Ogata, Y. (1998). Space-time point-process models for earthquake occurrences, *Ann. Inst. Stat. Math.* **50**, no. 2, 379–402.
- Peresan, A., V. G. Kossobokov, and G. F. Panza (2012). Operational earthquake forecast/prediction, *Rend. Fis. Acc. Lincei* **23**, 131–138.
- Reasenber, P. A., and L. M. Jones (1989). Earthquake hazard after a mainshock in California, *Science* **243**, 1173–1176.
- Reasenber, P. A., and L. M. Jones (1994). Earthquake aftershocks: Update, *Science* **265**, 1251–1252.
- Rhoades, D. A., M. C. Gerstenberger, A. Christophersen, J. D. Zechar, D. Schorlemmer, M. J. Werner, and T. H. Jordan (2014). Regional earthquake likelihood models II: Information gains of multiplicative hybrids, *Bull. Seismol. Soc. Am.* **104**, 3072–3083.
- Richards-Dinger, K., and J. H. Dieterich (2012). RSQSim earthquake simulator, *Seismol. Res. Lett.* **83**, 983–990.
- Richter, C. F. (1958). *Elementary Seismology*, W. F. Freeman, San Francisco, California.
- Schoenberg, F., and B. Bolt (2000). Short-term exciting, long-term correcting models for earthquake catalogs, *Bull. Seismol. Soc. Am.* **90**, 849–858.
- Segou, M., T. Parsons, and W. Ellsworth (2013). Comparative evaluation of physics-based and statistical forecasts in northern California, *J. Geophys. Res.* **118**, 6219–6240, doi: [10.1002/2013JB010313](https://doi.org/10.1002/2013JB010313).
- Shaw, B. E. (2009). Constant stress drop from small to great earthquakes in magnitude-area scaling, *Bull. Seismol. Soc. Am.* **99**, 871, doi: [10.1785/0120080006](https://doi.org/10.1785/0120080006).
- Shaw, B. E. (2013). Appendix E: Evaluation of magnitude-scaling relationships and depth of rupture: Recommendation for UCERF3, *U.S. Geol. Surv. Open-File Rept. 2013-1165-E*, and *California Geol. Surv. Special Rept. 228-E*.
- Tormann, T., B. Enescu, J. Woessner, and S. Wiemer (2015). Randomness of megathrust earthquakes implied by rapid stress recovery after the Japan earthquake, *Nat. Geosci.* **8**, 152–158.
- Tormann, T., B. Enescu, J. Woessner, and S. Wiemer (2016). Reply to “Tohoku rupture reloaded?” *Nat. Geosci.* **9**, 183–185.
- Utsu, T. (1961). A statistical study on the occurrence of aftershocks, *Geophys. Mag.* **30**, 521–605.
- Utsu, T. (1971). Aftershocks and earthquake statistics (III)—Analyses of the distribution of earthquakes in magnitude, time, and space with special consideration to clustering characteristics of earthquake occurrence (1), *J. Facul. Sci.* **3**, 379–441.
- van der Elst, N. J., and B. E. Shaw (2015). Larger aftershocks happen farther away: Nonseparability of magnitude and spatial distributions of aftershocks, *Geophys. Res. Lett.* **42**, 5771–5778, doi: [10.1002/2015GL064734](https://doi.org/10.1002/2015GL064734).
- Wang, K., and G. C. Rogers (2014). Earthquake preparedness should not fluctuate on a daily or weekly basis, *Seismol. Res. Lett.* **85**, 569–571.
- Werner, M. J., A. Helmstetter, D. D. Jackson, and Y. Y. Kagan (2011). High-resolution long-term and short-term earthquake forecasts for California, *Bull. Seismol. Soc. Am.* **101**, 1630–1648.
- Woessner, J., S. Hainzl, W. Marzocchi, M. J. Werner, A. M. Lombardi, F. Catalli, B. Enescu, M. Cocco, M. C. Gerstenberger, and S. Wiemer (2011). A retrospective comparative forecast test on the 1992 Landers sequence, *J. Geophys. Res.* **116**, no. B05305, doi: [10.1029/2010JB007846](https://doi.org/10.1029/2010JB007846).
- Working Group on California Earthquake Probabilities (WGCEP) (1988). Probabilities of large earthquakes occurring in California on the San Andreas fault, *U.S. Geol. Surv. Open-File Rept.*, p. 62.
- WGCEP (1990). Probabilities of large earthquakes in the San Francisco Bay Region, California, *U.S. Geol. Surv. Circular*, p. 51.
- WGCEP (1995). Seismic hazards in southern California: probable earthquakes, 1994–2024, *Bull. Seismol. Soc. Am.* **85**, 379–439.
- WGCEP (2003). Earthquake probabilities in the San Francisco Bay region: 2002–2031, *U.S. Geol. Surv. Open-File Rept. 2003-214*.
- WGCEP (2007). The Uniform California Earthquake Rupture Forecast, Version 2 (UCERF 2), *U.S. Geol. Surv. Open-File Rept. 2007-1437*.
- Zechar, J. D., D. Schorlemmer, M. J. Werner, M. C. Gerstenberger, D. A. Rhoades, and T. H. Jordan (2013). Regional earthquake likelihood models I: First-order results, *Bull. Seismol. Soc. Am.* **103**, 787–798.

Appendix A

Overall Calculation Sequence

The UCERF3-ETAS calculation sequence is described here, with the goal of producing a synthetic catalog of events given the following model inputs: an earthquake-rupture forecast (ERF); the desired simulation *TimeSpan*; the region of interest (see Fig. 1 for the results presented here); a list of all $M \geq 2.5$ earthquakes that have occurred up until the

chosen simulation start time (*ObsEqkList*), plus any fake ruptures that are desired for testing purposes; the average depth distribution of earthquake hypocenters for the region (*SeisDepthDistribution*, for which that assumed here is shown in Fig. 6); and values for the other model parameters (*ApplyGRcorr*, *ApplyGridSeisCorr*, *ApplySubSeisForSupraNucl*, *TotalRateScaleFactor*, and *ProbModel*; see Table 2).

Given these inputs, we take the following steps to produce the desired list of simulated ruptures (*SimulatedRupsList*):

1. We compute the total expected regional rate of $M \geq 2.5$ events for the specified *TimeSpan* by summing the equivalent Poisson rate of each rupture in the ERF (the latter provides rupture probabilities P_r that can be converted to rate as $-\ln(1 - P_r)/\Delta T$, in which ΔT is the forecast duration). This total rate is dominated by gridded seismicity because the minimum magnitude of fault-based ruptures is generally around M 6.3, almost four magnitude units above the minimum magnitude of gridded seismicity (M 2.5 here).
2. Then, we reset the forecast duration in the ERF to 1 or 2 years (regardless of the specified simulation duration) and compute the total equivalent Poisson rate for each gridded seismicity and fault-based source. The results are stored in an array (*SourceRatesArray*), and the values for fault-based sources will get updated during the simulation to account for any elastic-rebound probability changes. In what follows, we assume that the relative probability of fault-based ruptures remains constant (unchanged by elastic rebound) between the occurrence times of such events, which is a safe assumption, given the average interval between UCERF3 fault-based ruptures is about 2 years. The results are therefore not sensitive to the exact duration specified here, as long as it is less than a few years.
3. Next, we generate a *SpontaneousEventSampler*, which holds the current yearly rate of each rupture and provides random sampling among them, based on their relative rates.
4. For each event in the *ObsEqkList*, we then compute the expected number of primary aftershocks N from equation (2) and randomly sample an actual number assuming a Poisson process. For each primary aftershock, we then randomly sample an event time according to the temporal decay represented in equation (1). Then, we create an *EqkRupture* for each event, set the origin time and parent information accordingly, and put it into a chronologically ordered list for further processing (*EventsToProcessList*) because other attributes, such as magnitude and location, will be determined later. Again, this process is repeated for every earthquake in the *ObsEqkList*.

Then, we randomly sample the origin time of spontaneous events, considering the time dependence in λ_0 over the *TimeSpan*, as described in the main text. For each, we create an associated *EqkRupture* and put it into the *EventsToProcessList* (again, other attributes will be filled in later).

5. Next, we create an *ETAS_PrimaryAftershockSampler*, which randomly selects a primary rupture from the ERF

for a given parent event, taking into account both the relative rates of events throughout the model and the distance decay of aftershocks. This component, which is described extensively in the main text, has access to the *SourceRatesArray*, defined here in order to stay informed about any elastic-rebound probability changes that occur during the simulation.

6. Now, we loop over the *EventsToProcessList*, which is in chronological order, to fill in the other information about each event (because they currently contain only origin time and parent information). For each event we do the following:
 - If the event is spontaneous (no parent):
Sample an ERF rupture from *SpontaneousEventSampler* and fill in event attributes for the triggered earthquake accordingly (e.g., magnitude, rake, and rupture surface). If it is a gridded seismicity rupture, randomly choose an epicenter assuming a uniform distribution within the grid cell, and then randomly sample a hypocentral depth from the *SeisDepthDistribution* (Fig. 6).
 - Otherwise it is a triggered event, so do the following:
Sample an ERF rupture from the *ETAS_PrimaryAftershockSampler* and fill in the event attributes accordingly (e.g., magnitude, rake, and rupture surface). Again, details on how this is done are given in the main text.
 - Move the event from *EventsToProcessList* to *SimulatedRupsList*.
 - Sample a number of primary aftershocks and their origin times for the event, as outlined in step (4) above (but only over what remains of the simulation time span) and add each to the *EventsToProcessList* (which again, maintains chronological order within the list).
 - If the sampled event is a fault-system rupture (supraseis-mogenic) and the *ProbModel* parameter is not set as POISSON, apply elastic-rebound probability changes in the ERF by setting the date of last event on all subsections of the rupture to the origin time and change the start time of the forecast to the origin time of the sampled event as well. Then, update the *SourceRatesArray* (the pointer held by *ETAS_PrimaryAftershockSampler*) and the *SpontaneousEventSampler* accordingly, which will reduce probabilities for events that utilize the sections just ruptured, and will update the probabilities of other events due to time having marched on.
 - Continue looping through *EventsToProcessList* until it is empty.

Appendix B

Algorithm for Sampling Primary Aftershocks from the ERF

This section provides further details on how the *ETAS_PrimaryAftershockSampler* component samples aftershocks from the ERF. For a given parent event, the algorithm proceeds as follows:

1. Define the point on the parent rupture surface from which the triggering occurs. If the mainshock is a fault-based rupture, randomly sample a point from that surface (depth dependence of seismicity is not applied in this sampling, but it could be). For non-fault-based ruptures, use the hypocenter if $M \leq 4$, or randomly sample a point from within a sphere centered on the hypocenter, for which the radius of the sphere is given by equation (4) and the sampling distribution is uniform within. This is the *ParentLocation*, or again, the point from which the triggering occurs.
2. Translate this location to the nearest cube corner (essentially moving the *ParentLocation* by no more than ~ 1.7 km, which is about half the distance between opposite cube corners). This is referred to as the *TranslatedParentLocation*.
3. According to the depth of the *TranslatedParentLocation*, use the appropriate *DistanceDecayCubeSampler* to sample a neighboring cube from which the triggered event nucleates (see main text for details).
4. From the chosen cube, randomly sample whether the rupture is from a gridded seismicity source or from one of any fault subsections therein (according to current relative nucleation rates for each and accounting for any elastic-rebound effects on the faults).
5. If the sampled event is a gridded seismicity rupture, use the *DistanceDecayCubeSampler* to randomly sample a hypocentral location inside the cube, and translate this location by the exact opposite amount applied in step (2). For now, gridded seismicity events are treated as point sources, but finite surfaces could be assigned if desired. Also sample a magnitude and focal mechanism according to their relative likelihoods.
Otherwise, the sampled event is from a fault subsection, so randomly sample one of the ruptures that can nucleate from the subsection according to the current relative likelihood of each doing so.
6. Return information about the sampled event.

U.S. Geological Survey
Denver Federal Center
PO Box 25046, MS 966
Denver, Colorado 80225-0046
field@usgs.gov
(E.H.F.)

Southern California Earthquake Center
University of Southern California
3651 Trousdale Parkway #169
Los Angeles, California 90089-0742
(K.R.M., T.H.J.)

U.S. Geological Survey
345 Middlefield Road, MS 977
Menlo Park, California 94025
(J.L.H., A.J.M.)

U.S. Geological Survey
525 S. Wilson Avenue
Pasadena, California 91106
(M.T.P., N.v.)

Lamont Doherty Earth Observatory
Columbia University
Palisades, New York 10964
(B.E.S.)

School of Earth Sciences and Cabot Institute
University of Bristol
Wills Memorial Building, Queen's Road
Bristol BS8 1RJ, UK
(M.J.W.)

Manuscript received 2 June 2016