The UCERF3 Grand Inversion: Solving for the Long-Term Rate of Ruptures in a Fault System

by Morgan T. Page, Edward H. Field, Kevin R. Milner, and Peter M. Powers

Abstract We present implementation details, testing, and results from a new inversion-based methodology, known colloquially as the “grand inversion,” developed for the Uniform California Earthquake Rupture Forecast (UCERF3). We employ a parallel simulated annealing algorithm to solve for the long-term rate of all ruptures that extend through the seismogenic thickness on major mapped faults in California while simultaneously satisfying available slip-rate, paleoseismic event-rate, and magnitude-distribution constraints. The inversion methodology enables the relaxation of fault segmentation and allows for the incorporation of multifault ruptures, which are needed to remove magnitude-distribution misfits that were present in the previous model, UCERF2. The grand inversion is more objective than past methodologies, as it eliminates the need to prescriptively assign rupture rates. It also provides a means to easily update the model as new data become available. In addition to UCERF3 model results, we present verification of the grand inversion, including sensitivity tests, tuning of equation set weights, convergence metrics, and a synthetic test. These tests demonstrate that while individual rupture rates are poorly resolved by the data, integrated quantities such as magnitude–frequency distributions and, most importantly, hazard metrics, are much more robust.

Introduction

The development of earthquake rupture forecasts in California dates back to the original Working Group on California Earthquake Probabilities (WGCEP, 1988), which considered the probabilities of future earthquakes on different segments of the San Andreas fault (SAF). The most recent California model, Uniform California Earthquake Rupture Forecast (UCERF2) by the 2007 WGCEP (Field *et al.*, 2009), used expert opinion to determine the rates of ruptures on many of the major faults. This expert opinion framework is not compatible with the incorporation of significant numbers of multifault ruptures on a large, complex fault system. Furthermore, the UCERF2 methodology had no way to simultaneously constrain the model to fit both observed paleoseismic event rates and fault slip rates. These limitations were recognized by the leaders of the UCERF2 effort, but there was no methodology to address them at that time.

UCERF3, by the 2014 WGCEP (Field *et al.*, 2013, referred to hereafter as the UCERF3 main report), uses a system-level, algorithmic approach to consider a wider range of possible earthquake ruptures. Some of the primary goals for UCERF3 have been to relax segmentation assumptions, include multifault ruptures, and better match the observed regional magnitude–frequency distribution (MFD). These goals motivated the development of an inversion approach in UCERF3. This inversion, which has become colloquially

known as the “grand inversion,” is described in detail in this paper. In particular, we present details necessary to implement the inversion (e.g., the parallelized simulated annealing methodology and equation set weights), various tests demonstrating reliability, and results that supplement those in the main report.

The purpose of the grand inversion is to solve for the long-term rate of all possible supra-seismogenic ruptures in the fault-system model, in which supra-seismogenic means the rupture length is greater than or equal to the average down-dip width. The grand inversion builds on the methodology first proposed by Andrews and Scherer (2000) and developed further by Field and Page (2011). The rates of earthquakes are constrained by fault slip rates, paleoseismic event rates and average slips, MFDs observed in seismicity, and other *a priori* and smoothing constraints. Because of the size of the numerical problem, we use a simulated annealing algorithm (Kirkpatrick *et al.*, 1983) to invert for rupture rates. The simulated annealing method also has the advantage that it can give multiple solutions that satisfy the data. This allows us to more fully explore epistemic uncertainties. In this paper we describe (1) the constraints used in the inversion, (2) the simulated annealing algorithm that solves the inverse problem, (3) testing of the inversion methodology, and (4) final UCERF3 modeling results.

Setting up the Inversion: Data and Constraints

As described in the main report, UCERF3 uses the logic tree shown in Figure 1 to represent epistemic uncertainties (alternative models). A UCERF3 reference branch is shown in bold, which is not intended to represent the preferred branch, but rather to provide a reference with which to conduct tests and to compare against other models.

The fault models listed in the logic tree define the geometries of fault sections (Dawson, 2013). Fault sections include single faults such as the Cucamonga; larger faults are divided into several sections (e.g., the northern San Andreas is divided into four sections: Offshore, North Coast, Peninsula, and Santa Cruz). These fault sections are further divided into subsections with width approximately 7 km long for vertically dipping faults (see the UCERF3 main report for more details). These subsections are for numerical tractability and do not have geologic meaning.

The inversion methodology solves for rates of ruptures that are consistent with the data and constraints. Some of these constraints differ depending on the inversion model branch; here we describe and present the characteristic branch solution, which is constrained to have fault MFDs that are as close to UCERF2 MFDs as possible. The UCERF2 model assumed a characteristic magnitude distribution on faults (Wesnousky *et al.*, 1983; Schwartz and Coppersmith, 1984). The inversion can also be used to solve for rates consistent with Gutenberg–Richter (GR) MFDs (Gutenberg and Richter, 1944), or the on-fault MFDs can be unconstrained and allowed to be whatever best satisfies the other constraints. Here we describe only the characteristic branch solution and setup, as it more easily satisfies historical seismicity rates than the GR branch (see the UCERF3 main report for more details). The characteristic branch is designed to be as similar as possible to the UCERF2 ruptures rates, while simultaneously satisfying slip rates and paleoseismic data, allowing multifault ruptures, and eliminating the magnitude distribution “bulge” (overprediction relative to seismicity rates) that existed in the UCERF2 model.

Inversion Constraints

Slip Rates. The inversion is constrained to match long-term slip rates for each fault subsection, as given by the deformation model (Dawson and Weldon, 2013; Parsons *et al.*, 2013). This requires computing the average slip on each subsection in each rupture, where the average is over multiple occurrences of the event. The average slip over the entire rupture is computed using a magnitude–length or magnitude–area relationship (Shaw, 2013), and this slip is then partitioned onto individual subsections either uniformly or with a tapered distribution (see the UCERF3 main report). The average slip on a subsection in each rupture, multiplied by the rate of that rupture, must sum over all ruptures to the long-term slip rate for that subsection. The slip rates used in this constraint have been reduced from the slip rates specified in the deformation models to account for subseismogenic-thickness ruptures and

Figure 1. The UCERF3 logic tree and branch weights. The reference branch is shown in bold.

aseismicity (see the UCERF3 main report). This constraint is applied to each fault subsection in both normalized and unnormalized form. For the normalized constraint, each slip-rate constraint is normalized by the target slip rate (so that misfit is proportional to the fractional difference, rather than the absolute difference, between the model slip rates and the target slip rates), with the exception of slip rates below 0.1 mm/yr. This prevents some extremely low slip rates from dominating the misfit. Including both normalized and unnormalized forms of this constraint means we are minimizing both the ratio and the difference between the target and model slip rates; this approach represents a balance between better fitting slip rates on fast faults such as the San Andreas versus smaller slip rates on slower, secondary faults. These constraints can be written as

$$\sum_{r=1}^R D_{sr} f_r = v_s \quad \text{and} \quad \sum_{r=1}^R \frac{D_{sr}}{v_s} f_r = \frac{v_s}{v_s}, \quad (1)$$

in which $v'_s = \max(v_s, 0.1 \text{ mm/yr})$, v_s is the s th subsection slip rate, f_r is the long-term rate of the r th rupture, and D_{sr} is the average slip on the s th subsection in the r th rupture. Depending on the logic-tree branch, D_{sr} is either a uniform distribution along strike or a tapered distribution (Biasi *et al.*, 2013).

Paleoseismic Data. There are two types of paleoseismic data used to constrain the inversion: event rates and mean slips at locations on faults. Both of these data are treated in

Table 1
Inversion Parameter Settings*

Parameter Name	Characteristic Branch Setting	Details
Slip-rate constraint weight	1/100	Weight of slip-rate constraints normalized by slip rate/unnormalized (units: m/yr)
Paleoseismic constraint weight	1.2/0.12	Weight of paleoseismic event-rate constraints/mean slip constraints (which are converted to event rates). Both are normalized by event-rate standard deviation
Regional MFD [†] equality constraint weight	10	Normalized by target rate for each magnitude bin
Regional MFD inequality constraint weight	1000	Normalized by target rate for each magnitude bin
Fault-section smoothness weight	1000	Applied to faults with paleoseismic data. Units of earthquakes/magnitude bin/yr
Parkfield event-rate constraint weight	1000	Units of earthquakes/yr
Fault section MFD constraint weight	0.01	Keeps MFDs on-fault subsections close to UCERF2/proto-UCERF2 MFDs. Normalized by target rate for each magnitude bin for all nonzero bins; rupture rates (1/yr) in zero bins minimized with a weight of 0.1
Minimum rupture rate fraction	0.01	Fraction of on-fault moment that is applied as a minimum rupture rate (water level)
Initial rupture model	Zero-rate starting solution	Starting set of rupture rates for simulated annealing algorithm
Regional MFD form	Trilinear distribution	Target MFD for on-fault ruptures: total region MFD target less off-fault seismicity
Regional MFD transition magnitude	7.85	Magnitude to switch from MFD equality constraint to MFD inequality constraint

*Some of the inversion constraint weights are dependent on the units of the particular constraints to which they apply (and to the extent that units differ, weights are not directly comparable). Units of data vector components are given for unnormalized constraints.

[†]MFD, magnitude–frequency distribution.

a similar fashion; first, the mean slip data are divided by the slip rate at its location to turn it into an effective event rate. The total rate of all ruptures that include a given fault subsection, multiplied by the probability each rupture is paleoseismically visible (Madden *et al.*, 2013; Weldon and Biasi, 2013), must sum to the mean paleoseismic event rate for that subsection. The function that determines the probability an event would be seen in a trench differs for the event-rate data (from sites with timing) and the mean slip data (from sites with offset features). This constraint is applied to a total of 31 event-rate sites and 23 mean slip sites in the paleoseismic database (Weldon *et al.*, 2013). The data are weighted by the errors; however, the mean slip data errors do not contain sampling error, which dominates the total error, so these errors are an underestimate. This is accounted for by down weighting this data (see Table 1). These constraints can be expressed as

$$\sum_{r=1}^R \frac{G_{sr} P_{sr}^{\text{paleo}}}{\sigma_s} f_r = \frac{f_s^{\text{paleo}}}{\sigma_s}, \quad (2)$$

in which $G_{sr} = 1$ if the r th rupture includes the s th subsection and 0 otherwise, P_{sr}^{paleo} gives the probability that the r th rupture will be observed at the s th subsection, f_s^{paleo} is the paleoseismically observed mean event rate for the s th subsection, and σ_s is the standard deviation of the mean observed event rate.

Fault-Section Smoothness Constraint. We constrain the nucleation MFD (which gives the rate at which ruptures of a given magnitude nucleate on a given subsection) along fault

sections that contain paleoseismic data to smoothly vary along the fault. We use a Laplacian smoothing formula that constrains the rate of events nucleating in a given magnitude bin to smoothly vary along strike. This constraint prevents unphysical event-rate spikes or troughs near paleoseismic constraint locations. For each subsection s on a fault with paleoseismic data and its adjacent subsections $s - 1$ and $s + 1$, we apply

$$(R_s^m - R_{s-1}^m) + (R_s^m - R_{s+1}^m) = 0, \quad (3)$$

in which R_s^m is the nucleation rate of events in the m th magnitude bin on the s th subsection. At fault edges (in which a subsection s has one adjacent subsection $s - 1$), this constraint becomes

$$R_s^m - R_{s-1}^m = 0. \quad (4)$$

Parkfield Rupture-Rate Constraint. We constrain the total rate of Parkfield $\mathbf{M} \sim 6$ ruptures, which have been observed to be quasi periodic (Bakun and Lindh, 1985), to match the observed mean recurrence interval of 25 yrs. There are six so-called Parkfield earthquakes included in the constraint, one rupture that includes all eight Parkfield subsections, the two 7-subsection long ruptures, and the three 6-subsection long ruptures in the Parkfield section of the SAF. We choose a range of ruptures, rather than the single-Parkfield-section-long rupture, due to evidence that past $\mathbf{M} \sim 6$ earthquakes in Parkfield have ruptured slightly different areas of the fault (Custódio and Archuleta, 2007). The Parkfield earthquakes are not included in the MFD constraints discussed below, because the target on-fault magnitude distribution for some

branches does not have a high enough rate at **M** 6 to accommodate the Parkfield constraint.

The Parkfield constraint is the only *a priori* event-rate constraint we apply in UCERF3. These types of constraints are simply written as $f_r = f_r^{a \text{ priori}}$, in which $f_r^{a \text{ priori}}$ is the *a priori* rate of the *r*th rupture.

Nonnegativity Constraint and Water Level. Rupture rate, the rate per year that each rupture occurs, cannot be negative. This is a hard constraint that is not included in the system of equations but is strictly enforced in the simulated annealing algorithm, which does not search any solution space containing negative rates. To avoid zero-rate ruptures, we go further and apply a minimum rate (“water level”) for each rupture. The inverse problem with inequality constraints can be transformed into an equivalent problem with a simple nonnegativity constraint. The original inverse problem $\mathbf{Ax} = \mathbf{d}$ with $\mathbf{x} \geq \mathbf{x}_{\min}$, with observation matrix **A**, solution vector **x**, data vector **d**, and minimum rupture rates $\mathbf{x}_{\min}$, is mapped to a new inverse problem with a simple nonnegativity constraint, $\mathbf{Ax}' = \mathbf{d}'$, $\mathbf{x}' \geq 0$. The new data vector is given by $\mathbf{d}' = \mathbf{d} - \mathbf{Ax}_{\min}$ and the rupture-rate mapping is given by $\mathbf{x}' + \mathbf{x}_{\min} = \mathbf{x}$. The nonnegativity constraint for this transformed inversion problem is strictly enforced in the simulated annealing algorithm, which does not search any solution space that contains negative rates.

The minimum rupture rates are magnitude dependent and account for 1% of the on-fault moment in the model. These rates sum to a GR distribution (with a *b*-value of 1.0). Water-level rates within each 0.1 magnitude-unit bin are proportional to the lowest slip rate among the subsections utilized by the rupture.

Fault Section MFD Constraint. This constraint is applied to characteristic branch solutions and constrains the MFDs on fault subsections to be close to the characteristic MFDs used in UCERF2. For the faults that were treated as type A in UCERF2, for which rupture rates were derived from paleoseismic data using expert opinion, we use the final UCERF2 MFDs as the constraint here. For the remaining faults (both type B faults in UCERF2 and newly added faults that were not in UCERF2), we construct an MFD consistent with UCERF2 MFD-methodology, in which 1/3 of the moment is in a GR distribution and 2/3 of the moment is in a characteristic distribution (see the UCERF3 main report for more details). This constraint can be written as

$$\sum_{r=1}^R \frac{M_{sr}^m}{R_s^m} f_r = 1, \quad \text{for all } R_s^m > 0, \quad (5)$$

in which M_{sr}^m is the fraction of the *r*th rupture in the *m*th magnitude bin on the *s*th subsection. Rupture rates for magnitude bins in which $R_s^m = 0$ are also minimized.

Regional MFD Constraint. The on-fault target magnitude distribution for the characteristic branch is a trilinear model as shown in Figure 2. Above the maximum magnitude for

Figure 2. Schematic of magnitude–frequency distributions (MFDs) for the UCERF3 reference branch. The total target MFD (found from seismicity; see Felzer, 2013) is shown in black. It is reduced to account for off-fault earthquakes to give the total on-fault target (orange) and further reduced by removing ruptures with lengths less than the subseismogenic thickness to give the supra-seismogenic on-fault target (blue), which is used to constrain the regional MFD given by the inversion solution. (GR, Gutenberg–Richter values.)

off-fault seismicity, the on-fault MFD equals the total regional MFD (total seismicity rates are extrapolated from historical seismicity, see Felzer, 2013). Below the maximum off-fault magnitude, the on-fault MFD transitions to a lower *b*-value, and then transitions again back to a *b*-value of 1.0 (matching historical seismicity) at about **M** 6.25 (this is the average minimum magnitude for the supra-seismogenic-thickness ruptures, the exact value of which is branch dependent). The rate of **M** ≥ 5 events for the on-fault MFD is set from historical seismicity that is inside our fault polygons (these are zones around each fault and are described in the UCERF3 main report). Thus the entire on-fault MFD is uniquely determined by the maximum magnitude off-fault, the average minimum magnitude for the on-fault (supra-seismogenic) ruptures, historical seismicity rates, and the fraction of seismicity that is considered on-fault (within our fault polygons).

The inversion solves only for rates of earthquakes that have lengths greater than the (local) seismogenic thickness, so the on-fault MFD is reduced at low magnitudes by the rate of subseismogenic ruptures on each fault section. This gives the supra-seismogenic on-fault target, shown in Figure 2, which is the MFD used for the inversion magnitude-distribution constraint.

For each branch, the regional magnitude-distribution constraint is applied as an equality constraint up to and including magnitude 7.85 and is applied as an inequality constraint (with the target on-fault MFD serving as an upper bound) above **M** 7.85. Thus the inversion is not constrained to exactly equal the target MFD at magnitudes for which the true MFD may taper below a strict GR distribution. As it is not known what the MFD looks like at high magnitudes, we allow the inversion to choose whatever roll-off at high magnitudes is consistent with the other constraints.

Equation Set Weights					Data Misfits			
Slip Rate Wt (normalized/unnormalized)	Paleo Wt (event rate / avg. slip)	Regional MFD Wt	Fault Section MFD Wt	Fault Section Smoothness Wt	Regional MFD Misfit	Paleo Event Rate Misfit	Paleo Slip Misfit	Slip Rate Misfit
0.1 / 10	1.2 / 0.12	10	0.01	1000	0.00013	1.64	92.2	0.06946
1 / 100	0.12 / 0.012	10	0.01	1000	0.01645	134.40	204.6	0.00652
1 / 100	1.2 / 0.12	1	0.01	1000	10.50161	8.46	91.6	0.00535
1 / 100	1.2 / 0.12	10	0.001	1000	0.04151	13.57	92.9	0.00839
1 / 100	1.2 / 0.12	10	0.01	100	0.04051	9.53	110.7	0.00845
1 / 100	1.2 / 0.12	10	0.01	1000	0.03963	18.26	117.1	0.00867
1 / 100	1.2 / 0.12	10	0.01	10000	0.05043	35.66	158.8	0.01004
1 / 100	1.2 / 0.12	10	0.1	1000	0.04229	21.88	173.8	0.01580
1 / 100	1.2 / 0.12	100	0.01	1000	0.00004	47.18	183.5	0.02925
1 / 100	12 / 1.2	10	0.01	1000	0.16739	1.11	41.2	0.01869
10 / 1000	1.2 / 0.12	10	0.01	1000	6.97045	87.38	110.9	0.00351

Good
Acceptable
Poor

Figure 3. Data misfits for alternative equation set weights. Misfits are shaded by the quality of the data fit. Thresholds for the misfit shadings were determined by visually examining a different set of runs; these are subjective, because most data do not have formal error bars. The default constraint weights are highlighted with bold in the center row.

For each branch, this constraint is applied for two regions: northern California and southern California. Furthermore, the $M \sim 6$ Parkfield earthquakes are excluded from the constraint, as their rate is quite high and would otherwise be underpredicted by the model. This constraint is implemented in the form

$$\sum_{r=1}^R \frac{M_{gr}^m}{R_g^m} f_r = 1 \quad \text{for } M \leq 7.85, \quad \text{and}$$

$$\sum_{r=1}^R \frac{M_{gr}^m}{R_g^m} f_r \leq 1 \quad \text{for } M > 7.85, \quad (6)$$

in which M_{gr}^m is the fraction of the r th rupture in the m th magnitude bin and the g th region and R_g^m is the rate in the m th magnitude bin and the g th region given by the supra-seismogenic on-fault target MFD.

Tuning of Constraint Weights

Weights for the different constraints used in the inversion are shown in Table 1. Some of the constraint weights have units, so a relatively low weight does not necessarily mean a constraint is not significantly affecting the inversion result. Naturally, there are trade-offs between how well different datasets can be fit and also between data fit and how smooth or close the model is to UCERF2. We select weights for the dataset so that neither slip rates nor paleoseismic data are badly underfit and so that the regional MFD data and the average rate of Parkfield earthquakes are matched nearly exactly. The MFD smoothing constraint and MFD nucleation (UCERF2 MFDs) constraint are set to the highest value possible—increasing the weights beyond these current settings causes a sharp, significant degrade to the other data constraints. Our goal with the constraint weights is to find a model that is as smooth and as close to UCERF2 as possible while fitting all datasets within their uncertainties. Still, small changes in these weights are possible; the question is, do

somewhat arbitrary choices in constraint weighting significantly affect the inversion result?

To test the effect of the constraint weights, we perturb the weights for each of the constraints, individually, by factors of 10 from the default values. The data misfits for these trial runs are shown in Figure 3. The default constraint weights have no poor data fits, and the only other runs in this set that have no poor data fits are either less smooth (smaller along-fault MFD smoothing) or further from UCERF2 (smaller MFD nucleation constraint weight), with the exception of a higher paleoseismic weight, which degrades the regional MFD and makes the model less smooth. The models that are either smoother or further from UCERF2 fit the data somewhat better than the default weights, but not significantly so given the loss of regularization. We can also see that fits to the slip rates and paleoseismic data strongly trade-off against each other; this is in part due to the regional MFD constraint, which limits the extent that the inversion can alter the size distribution along faults to fit both slip-rate and paleoseismic data.

Most importantly, Figure 4 shows that changes in the constraint weights matter little for hazard, even when changes are large enough to cause some data to be poorly fit. The range of hazard implications for the test inversion runs with different constraint weights is far less than the range of hazard spanned by different UCERF3 logic-tree branches. These results imply that it is acceptable to use one set of weights for all inversion runs because changes in these weights (which would have to be smaller than the factors of 10 investigated here in order to not lead to poor data fits or a significantly less-regularized model) are not important for hazard.

The regional MFD constraint has a high weight in the inversion and is fit very well by UCERF3 models; this is to prevent the MFD overprediction (the bulge) that occurred in UCERF2 (Field *et al.*, 2009). This constraint is surprisingly powerful; when removed, the inversion has enormous freedom to fit the slip rates, for example, nearly perfectly. It is informative to relax this constraint to see how it changes the

Figure 4. Hazard implications for the UCERF3 reference branch with alternative equation set weights compared to the range spanned by different branches of the UCERF3 logic tree. AFE, Annual Frequency of Exceedance.

Figure 5. MFDs for a model in which the MFD constraint weight is relaxed, which allows the model to fit other data better at the expense of the fit to the MFD target. This demonstrates that the other constraints could be better fit if more moment were put on faults relative to the amount allowed by the target MFD.

model. MFD fits for the reference branch with the regional MFD constraint weight relaxed by a factor of 10 are shown in Figure 5. We can see that the inversion prefers more moment, interestingly, at all magnitudes. This model, corresponding to the “Relax MFD” branch (which is currently given zero weight in the UCERF3 logic tree), is very well captured by the logic-tree branches that increase the MFD target uniformly.

Defining the Set of Possible Fault-Based Ruptures

A final *a priori* step that must be done before running the inversion is defining the set of possible ruptures. The faults are first discretized into subsections, as described in the

UCERF3 main report. These subsections are specified for numerical tractability and do not represent geologic segments. Subsections are linked together to form ruptures, the rates of which are solved for by the inversion. It is worth noting that here a “rupture” is defined as an ordered list of fault subsections it includes; no hypocenter is specified.

Viable ruptures are generated from the digitized fault subsections via the plausibility filter rules described in detail by Milner *et al.* (2013). The plausibility filter allows ruptures to jump a maximum distance of 5 km between faults, imposes simple rules based on azimuth change along the rupture, and checks for kinematic consistency, based on Coulomb stress modeling, at fault junctions. The plausibility filter defines a total of 253,706 ruptures for fault model 3.1 and 305,709 ruptures for fault model 3.2. By comparison, our UCERF2 mapping has (for a smaller set of faults) 7029 on-fault ruptures.

The plausibility filter either allows ruptures in the rupture set or not—it is a binary filter. One important remaining question is to what extent ruptures allowed by the plausibility rules may require further penalty. Because the inversion fits fault-slip rates and the regional MFD, the frequency of multi-fault ruptures is already constrained. However, one could impose an additional improbability constraint to further penalize any ruptures that are deemed possible but improbable; these constraints are numerically equivalent to an *a priori* rupture-rate constraint with a rupture rate of zero. These constraints could be weighted individually, so some rupture rates could be penalized more harshly (i.e., contribute more to the misfit if they are nonzero) than others. This improbability constraint is not implemented for the UCERF3 rupture set because we do not have a viable model (i.e., a model ready for implementation and agreed to be useful by Working Group members) for such a constraint. Below we present evidence that this constraint is not needed, given empirical data and existing rupture-rate penalties.

It is not even clear what constitutes a multifault rupture in the case of a complex, possibly fractal, connected fault network. One way to quantify the rate of multifault ruptures is to simply define “multifault” in the context of the names assigned to faults in the UCERF3 database. For this definition of multifault, we count all sections of a fault such as the SAF as a single named fault, even though the individual sections have different names (e.g., Carrizo, Mojave North, etc.). Using the inverted rupture rates for the UCERF3 reference branch model, 40% of $M \geq 7$ ruptures and only 16% of paleoseismically visible ruptures are on multiple faults. We can compare this number to the ruptures in the Wesnousky database of surface ruptures (Wesnousky, 2008); of these, 50%, or 14 out of 28, are on multiple faults. So by this albeit simple metric, the solutions given by the inversion algorithm are not producing more multifault ruptures than are seen in nature.

Another multifault metric from our model that we can compare with empirical data from the Wesnousky database is the rate of ruptures that have no jumps between faults versus 1, 2, or 3 jumps greater than 1 km. This comparison is shown in Figure 6. This figure shows that the inversion constraints greatly reduce the rate of multifault ruptures in the solution relative to their frequency in the rupture set. Also, by this metric, the inversion is, again, underpredicting the rate of multifault ruptures relative to the empirical data. We would even further underpredict the rates of ruptures with jumps if an improbability constraint were added to the inversion.

It is important to note that many of the truly multifault ruptures seen in nature may in fact not be part of our model at all because they could “link up” known, mapped faults with unknown faults (or faults that are known but not sufficiently studied as to be included in our fault model). Our model completely separates on-fault ruptures from background earthquakes; there are no ruptures that are partly on-fault and partly off-fault.

Without the improbability constraint, the inversion is not likely to give a lower rate to a rupture that is only moderately kinematically compatible (but allowed by our filters) relative to a more kinematically favored rupture, unless kinematically less-favored jumps are also disfavored by slip-rate changes. However, the most egregious fault-to-fault jumps are already excluded by the Coulomb criteria described by Milner *et al.* (2013). It appears that the slip-rate constraints and rupture filtering leave only a modest amount of multifault ruptures in the inversion solution; further penalties to multifault ruptures will result in larger deviations from the empirical data.

The Simulated Annealing Algorithm

We use a simulated annealing (SA) algorithm to solve the nonnegative least-squares problem $\mathbf{Ax} = \mathbf{d}$ with the additional constraint $\mathbf{A}_{\text{ineq}}\mathbf{x} \leq \mathbf{d}_{\text{ineq}}$ (this last constraint is due to the regional MFD inequality constraint). The SA algorithm simulates the slow cooling of a physical material to form a crystal. In the same way that annealing a metal reduces defects and allows the material to reach a lower thermody-

Figure 6. Rates of ruptures with different numbers of fault-to-fault jumps for the UCERF3 rupture set (dashed line) and UCERF3 model. Here a “jump” is defined as a pair of adjacent subsections in the rupture that are greater than 1 km apart, in 3D, as defined by our fault model. The empirical data is also shown. By this metric, the inversion model does not have as many multifault ruptures as the empirical data.

namic energy, the simulated annealing algorithm attempts to minimize energy (in simulated annealing parlance this is the summed squared misfit between the data and synthetics) by slowly decreasing the probability of jumps to worse solutions. The algorithm we employ has the following steps:

1. Set $\mathbf{x}$ equal to initial solution $\mathbf{x}_0$. We have tested different initial solutions; the final UCERF3 model uses an initial solution of all zero rupture rates.
2. Lower the parameter T , known as the “temperature,” from 1 to 0 over a specified number of iterations. The temperature is given by the inverse of the iteration number, although different approaches to annealing that specify different cooling functions were also tested. Over each simulated annealing iteration:
 - One element of $\mathbf{x}$ (one rupture rate) is chosen at random. This element is then perturbed randomly. It is here that the nonnegativity constraint is applied, because the perturbation function is a function of the current rupture rate and will not perturb the rate to a negative value. Unlike some simulated annealing algorithms, our algorithm does not use smaller perturbations as the temperature is lowered (this was tested but did not result in faster convergence times).
 - The misfit for the perturbed vector $\mathbf{x}$, $\mathbf{x}_{\text{new}}$, is calculated; from this, the energy of that solution is $E_{\text{new}} = (\mathbf{Ax}_{\text{new}} - \mathbf{d})^2 + E_{\text{ineq}}$, in which E_{ineq} is additional energy from the MFD inequality constraint: $E_{\text{ineq}} = [\min(\mathbf{A}_{\text{ineq}}\mathbf{x}_{\text{new}} - \mathbf{d}_{\text{ineq}}, 0)]^2$. The weight on the inequality constraint is set quite high so that a solution that violates it faces a significant penalty; thus, in practice, this is a strict inequality constraint.

Figure 7. The UCERF3 simulated annealing algorithm. (Visualization courtesy of Tom Jordan, University of Southern California [USC].)

- The transition probability P is calculated based on the change in energy (between the previous state and the perturbed state) and the current temperature T . If the new model is better, $P = 1$. Therefore, a new model is always kept if it is better. If the new model is worse, it is kept with probability P . It is more likely that the solution will be kept early in the annealing process when the temperature is high. If $E < E_{\text{new}}$, then $P = e^{\frac{E - E_{\text{new}}}{T}}$.
3. Once the annealing schedule is completed, the best solution $\mathbf{x}$ found during the search (the solution with the lowest energy) is returned. (This is a common departure from pure simulated annealing, which returns the last state found. In some cases the final state will not be the best solution found, because occasionally solutions are discarded for worse solutions.)

The UCERF3 SA algorithm is shown graphically in the flowchart in Figure 7.

Simulated annealing works similarly to other nonlinear algorithms such as the genetic algorithm (Holland, 1975). One advantage of simulated annealing is that there is a strong theoretical backing: besides the analogy to annealing a physical material, the simulated annealing algorithm will find the global minimum given infinite cooling time, provided the annealing schedule lowers the temperature sufficiently slowly (Granville *et al.*, 1994).

There are several advantages of this algorithm in contrast to other approaches such as the nonnegative least-squares algorithm. First, the simulated annealing algorithm scales well as the problem size increases. (In fact, it would not be computationally feasible for us to use the nonnegative least-squares algorithm to solve a problem as large as the

UCERF3 grand inversion.) Simulated annealing is designed to efficiently search a large parameter space without getting stuck in local minima. Next, quite importantly, for an under-determined problem the simulated annealing algorithm gives multiple solutions (at varying levels of misfit depending on the annealing schedule). Thus both the resolution error (the range of models that satisfy one realization of the data) and the data error (the impact of parameter uncertainty on the model) can be sampled. Finally, simulated annealing allows us to include other nonlinear constraints in the inversion apart from nonnegativity; in our case we incorporate the MFD inequality constraint, which is a nonlinear constraint that cannot be easily incorporated into the perturbation function.

Parallelization of the Simulated Annealing Algorithm

To tackle the computational demands of the UCERF3 inversion for each node of the logic tree, we have implemented a parallel version of the simulated annealing algorithm. This algorithm runs the serial simulated annealing algorithm we have just described over a number of processors for a given number of subiterations or subcompletion time. Then the best solution among these is kept and redistributed over the processors; this process repeats until convergence criteria (a target misfit, a given number of total iterations, or an allotted annealing time) are satisfied. Final production runs for UCERF3 used a convergence criteria of 5 hrs of total (wall-clock) annealing time utilizing five processors on a single node running the parallel SA algorithm with a subcompletion time of 1 s; longer runs for testing purposes are discussed in the next section.

The parallel SA algorithm scales well up to 20–50 processors, but adding processors beyond this does not improve performance. Using the parallelized algorithm on a cluster

Figure 8. Convergence test showing hazard curves from 200 runs of the UCERF3 reference branch. For (a) Los Angeles and (b) San Francisco, the range of hazard curves from multiple runs (shaded area) is so small that it is not visible behind the mean hazard curve over all runs. We also show a hazard curve for (c) Diablo Canyon, because in this case differences between different inversion runs for the reference branch are visible at very low annual frequencies of exceedance; this is the largest variance among the 60 test locations.

results in average speedups of 6–20 relative to the serial algorithm. As the parallelized SA algorithm represents a departure from pure simulated annealing, we tested both the serial and parallelization algorithms to ensure the algorithm difference does not change the solution qualitatively; as discussed in the next section, the only difference we found was convergence speed.

Although a single inversion run can be done on a typical desktop computer, our final UCERF3 results include 10 runs on each branch of a 1440-branch logic tree and therefore require 3000 node-days of computation time. We used the Stampede cluster at the University of Texas and the HPCC cluster at the University of Southern California to run all the inversions required for the UCERF3 branches in under a day.

Testing of Inversion Methodology

Convergence Properties

The simulated annealing algorithm used in UCERF3 is stochastic; due to random perturbations of the model parameters during the annealing process, the final model will be different each time it is run, even if the starting model and data do not change. This can be an advantage because it allows us to explore the range of models that fit the data on a single branch; however, finding a stable solution can present a challenge. We have extensively tested the convergence properties of the inversion and found that over individual runs, unsurprisingly, individual rupture rates are very poorly constrained and are highly variable. However, averaged parameters such as fault section MFDs are far more robust.

In a single run of the inversion only about 10,000 ruptures (~4%) have rates above the water-level rates. This demonstrates that the data can be fit with a relatively small subset of ruptures. By averaging the rupture rates from multiple runs of the inversion, however, we can obtain many more greater-than-water-level rates because different runs use a different subset of ruptures to match the data. For a single branch, averaging 10 runs increases the number of ruptures above the water level to approximately 38,000 and averaging 200 runs in-

creases this number to approximately 115,000. Different runs of the same branch therefore match the data similarly well by using different sets of ruptures and are sampling the epistemic uncertainty of the problem in this way. Averaging across branches achieves smoother, less compact results; the average of 10 branches for all 720 logic-tree branches in fault model 3.1 produces 93% of ruptures above the water-level rates.

For many faults, stable parameters of the rate a given fault section participates in a rupture of a given magnitude can be determined from averaging only a few inversion runs; the median number of inversion runs needed to resolve the mean rate of $M \geq 6.7$ events on a fault within 10% is 9. A small number of slow-moving faults require on the order of 100 runs to obtain well-resolved magnitude distributions; as discussed in the UCERF3 main report, the worst case is the Richfield fault (a two-subsection-long fault near the Whittier fault), which requires 1294 runs. Because the branch-averaged UCERF3 model, known as Mean UCERF3, averages the results from 14,400 individual SA runs (each logic-tree branch run 10 times), all faults have MFDs that are well resolved.

Hazard results, as shown in Figure 8, are even more stable across inversion runs. Of the 2% in 50 yrs and 1% in 100 yrs annual frequencies of exceedance for peak ground acceleration (PGA) and spectral accelerations at 5 Hz, 1 Hz, and 4 s, the worst uncertainty we found over 10 runs of the UCERF3 reference branch was within 3% of the mean, as described in the UCERF3 main report. Thus individual branches of the logic tree, each of which average over 10 SA runs, are well resolved for all these hazard metrics. Branch-averaged results are extremely well resolved for hazard metrics; in fact inversion nonuniqueness was found to be negligible for the 2% in 50 yrs PGA, 1% in 100 yrs PGA, and 2% in 50 yrs hazard maps, as well as the 1% in 100 yrs 3 s spectral acceleration hazard maps. In fact, a single SA run per branch would be adequate for these hazard metrics.

Synthetic Test

To test the ability of the inversion method to recover a solution, we create a synthetic test based on the inversion

Figure 9. Slip-rate fits from five synthetic test runs (using the zero starting solution). Color shows the solution slip rate divided by the target slip rate, with white being no misfit.

Figure 10. California-wide MFD for five runs of the synthetic test with zero starting solution compared to the input distribution.

synthetics from the UCERF3 reference branch. The data used include the slip rate, paleoseismic, regional magnitude distribution, fault section MFD, and Parkfield event-rate synthetics. The magnitude-distribution smoothing constraint (used on fault sections with paleoseismic data) was applied as in a typical inversion, without using the reference branch synthetics. No errors were added to the input data. Misfits in the synthetic test demonstrate how well the inversion converges and differences between different runs of the synthetic test give a sense of the null space of the problem.

We ran the synthetic test with two different starting models, which are sets of initial rupture rates input to the simulated annealing algorithm. In UCERF3, our starting model is all zeros, that is, each rupture rate starts at a value of 0. In addition, for testing purposes, we also used a starting model of UCERF2 ruptures rates (or more accurately, UCERF2 rates mapped onto the inversion rupture set, we find the closest analogs to UCERF2 ruptures given our discretization of the fault system). The zero starting model allows for a wider range of final models (and thus a more thorough exploration of epistemic uncertainty). By contrast, with the UCERF2 starting model, the final rupture rates show less variation from run to run and are closer to the UCERF2 ruptures rates.

Slip-rate fits for multiple runs of this synthetic test are shown in Figure 9. The total squared misfit of the slip rates is 18 times less than that for the UCERF3 reference branch. The largest systematic discrepancy on the major faults is an overprediction of the slip rate of approximately 10% on the Mojave South section of the southern SAF. In addition, the southern end of the Ozena fault and the Swain Ravine fault have slip rates that are more severely overpredicted and underpredicted, respectively. These systematics are not persistent features of UCERF3 inversion runs.

The synthetic test runs fit the input regional MFD quite well, as shown in Figure 10. The average moment rate of the

Figure 11. San Andreas paleoseismic fits for the synthetic test (zero starting solution), average of five runs. The probability for detecting an event differs for paleoseismic sites that give timing information versus slip information. For this reason, model predictions for event rates at timing sites should be compared with the synthetic input data at timing sites; similarly, model predictions at slip sites should be compared to synthetic input data at slip sites.

synthetic tests is very close to (0.07% lower than) the moment rate of the model used to create the synthetics. UCERF3, on average, underpredicts the moment rate given by the deformation models by 1.3%, and the moment rate on individual branches ranges from 79% to 115% of the target moment. The synthetic tests suggest that moment-rate misfits are due to the data inconsistencies and not the inversion procedure.

Paleoseismic data fits for the synthetic model are shown in Figure 11. The fits are excellent and are not sensitive to the starting model used in the simulated annealing algorithm. Individual fault magnitude distributions, however, do show slight dependence on the starting solution used, as shown in Figure 12. In general, the fault MFDs are more variable for the synthetic test results using the zero starting model. This is not surprising, as the UCERF2 starting model better fits the data than the zero starting model, meaning that more of the initial perturbations will be kept in the latter case because they improve the misfit. If the inversion were to spend more time at high temperature, this dependence on the starting model should weaken; changes in the annealing schedule are discussed further in the [Long Cooling Test](#) section. In Figure 12 it is also apparent that synthetic test MFDs are more similar to the input model for the zero starting solution, which is the starting solution that was used for the synthetic test input model.

Synthetic test results for an isolated fault section, Battle Creek, are shown in Figure 13. In this figure, we can also see that while rates of individual ruptures vary from run to run, averaged quantities such as slip rates and total event rates are much more stable.

Slip-rate misfits have little dependence on the starting model. However, the zero starting solution models give a wider range of solutions, as well as less compact solutions (more ruptures have rates above the water-level rates), so they may be preferable in that they more adequately sample the null space.

Alternative Simulated Annealing Algorithms

We tested a variety of published perturbation functions and cooling functions to ensure that the simulated annealing algorithm we are using is optimized for our problem. The cooling functions, which give the simulated annealing temperature T as a function of iteration number i , include the following:

1. Classical SA ([Geman and Geman, 1984](#)): $T = 1/\log(i + 1)$,
2. Fast SA ([Szu and Hartley, 1987](#)): $T = 1/i$,
3. Very fast SA ([Ingber, 1989](#)): $T = e^{-i+1}$.

The perturbation functions we tested, in units of 1/yr, as a function of a random number $R \rightarrow [0, 1]$, are as follows:

1. Uniform, no temperature dependence: $\Delta \mathbf{x} = 0.001(R - 0.5)$,
2. Gaussian: $\Delta \mathbf{x} = 0.001T^{-1/2}Re^{1/2T}$,
3. Tangent: $\Delta \mathbf{x} = 0.001T \tan(\pi R - \pi/2)$,
4. Power
Law: $\Delta \mathbf{x} = \text{sgn}(R - 0.5)(0.001T)(1 + 1T)^{|2R-1|-1}$,
5. Exponential: $\Delta \mathbf{x} = 0.001T 10^R$.

The normalization factors for each of the above functions have been optimized within a factor of 10 to produce the fastest convergence. We tested each combination of cooling schedules and perturbation functions for both the serial and parallel SA algorithms. For UCERF3 runs we use the fast SA cooling schedule with uniform (no temperature dependence) perturbations. This produces the best final energies (and therefore lowest misfits) for both the parallel and serial simulated annealing algorithms, within the noise (i.e., neglecting the small differences in final energies typically seen over multiple runs of the same inversion).

Figure 12. San Andreas Mojave South MFDs for the synthetic test input and five models with (a) the starting model of zero rates and (b) UCERF2 starting model. Note that the peak at M 6.25 is due to aliasing (the magnitude difference between the two- and three-subsection-long ruptures results in one skipped 0.1-unit magnitude bin). (c) Slip-rate misfits for five runs for each starting model are shown.

Long Cooling Test

The simulated annealing algorithm more adequately searches the solution space with slower cooling schedules, although this also requires more computational time. Early in the simulated annealing, when temperature is high, jumps to worse solutions are more likely to be taken. This is necessary in order to avoid getting stuck in local minima. Our default simulated annealing parameters have been chosen to achieve the lowest misfit for simulated annealing runs that take 4–8 hrs; given the computational capacity we have at our disposal this allows us to compute the entire UCERF3 model, with 10 runs of each branch, in under a day using supercomputers. To ensure this annealing schedule is adequate, we ran a series of long cooling tests with the reference branch. Each was run for 40 hrs in total, using both the serial and parallel algorithms, for the standard cooling schedule, and for schedules in which the temperature is lowered 2, 5, and 10 times slower.

Figure 14 shows the simulated annealing energy (proportional to the squared misfit) versus time for the slow

cooling tests. The parallel algorithm outperforms the serial algorithm, which is not surprising because it uses multiple processors and is searching more of the space (and represents more total computational time summed over all processors). However, the difference between the parallel and serial algorithm is less pronounced at longer times, which suggests that while the parallel algorithm converges faster, given enough time, the serial algorithm can catch up.

The speed of the cooling, at least within the range we tested, has negligible effect on the final energies at very long times. The differences at 40 hrs between the test runs are within the variability seen for multiple runs with the same parameters. The same is true at 8 hrs, with the exception of the very slowest cooling schedule for the serial simulated annealing algorithm. However, at even shorter times the slower cooling functions perform less well. Thus we are not finding lower minima with the slower cooling, but we are converging to similar minima at a slower pace. This test can give us confidence that we are not quenching the system too fast. Although longer annealing times could improve our

Figure 13. Synthetic test results for the Battle Creek fault. (a) Participation MFDs (these show rates of events rupturing the segment, even if they nucleated elsewhere) for five runs with the zero-rate starting model compared to the model used to create the synthetics. (The spikiness in the MFDs is aliasing due to the magnitude differences between the 2-, 3- and 4-subsection-long ruptures.) (b) Yearly participation rates on the fault along with individual ruptures and rates (traces of ruptures are sorted and colored by rate and plotted above the fault) for the five runs; lower insets show slip-rate misfits.

misfits somewhat, each additional amount of computational time comes at less and less benefit to the misfit.

We visually inspected event-rate and slip-rate synthetics for the different long cooling tests and found no differences in the final models apart from typical run-to-run variability. In particular, there were no systematic differences in the types of solutions found with the slower cooling or with the serial versus parallel algorithm. Most importantly, it does not appear that we are neglecting a better set of minima by using our current annealing schedule.

Modeling Results

We now present results for the Mean UCERF3 model, which is the branch-weighted average of 10 runs for each of the 1440 branches of the UCERF3 logic tree shown in Figure 1. For simplicity, in map-based plots below, we neglect the fault model 3.2 branches.

In many of the figures below we compare the UCERF3 model to UCERF2, or rather a mapping of UCERF2 into our rupture set. For this mapped UCERF2 model, we set the magnitudes of ruptures to their mean UCERF2 values and

average (with 50% weight each) a tapered along-strike slip distribution and a uniform slip distribution for each rupture (these are the same weights used for the along-strike slip distributions in UCERF3).

UCERF3 Fits to Data

UCERF3 is constrained to match slip-rate targets, as defined by the deformation models, for each fault subsection. Figure 15 shows fits to the slip rates for UCERF2, Mean UCERF3, and branch-averaged solutions for each deformation model. The UCERF2 model tends to overpredict slip rates in the centers of fault sections and underpredict slip rates on the edges of fault sections. This is due to two effects: the 50% weight on the tapered along-strike slip distribution, and the floating ruptures used in UCERF2, the set of ruptures for each magnitude bin allowed at different positions along a given fault section, which overlap more in the center of fault sections. Slip-rate fits for the UCERF3 models are on the whole good, although individual branches have total moment rates ranging from 79% to 115% of the total on-fault target moment (this is the moment given by the deformation

Figure 14. Simulated annealing energy (summed squared residuals) versus time for (a, b) standard and (c, d) a cooling schedule 10 times slower, using the serial simulated annealing (a, c) and parallel simulated annealing (b, d) algorithms.

model, less creep and subseismogenic-rupture moment). Mean UCERF3 underpredicts the target moment rate (as determined by the deformation model slip rates) by 1.3%.

The two biggest slip-rate overpredictions are on the northern end of the Elsinore fault and in the center of the Creeping section of the San Andreas. The model overpredicts the target slip rate on the Elsinore fault due to a high paleoseismic event rate on that fault. The outlier in the Creeping section occurs where the mean slip rate is reduced by 80% to 5 mm/yr in order to account for creep. The inversion is unable to match this rapid along-strike slip-rate change; the solution only reduces the slip rate to approximately 10 mm/yr. The slip-rate reduction in the center of the Creeping section limits the rate that ruptures can propagate through this fault section and results in average repeat times for wall-to-wall San Andreas ruptures that extend from the (center of) SAF-Offshore to the SAF-Coachella section of 150,000 yrs. Ruptures that extend from the SAF-North Coast to SAF-Mojave South section

occur approximately every 2500 yrs in UCERF3, and ruptures that extend (just barely) through the Creeping section, from SAF-Parkfield to SAF-Santa Cruz, occur every 900 yrs.

All other UCERF3 slip-rate overpredictions are within 20% of the targets, and virtually all those greater than 10% can be explained by paleoseismic constraint inconsistencies. Furthermore, mean slip rates on all faults are within the error bounds given by the geologic deformation model, even though these error bounds were not used in the inversion.

Slip-rate and paleoseismic data fits for the SAF are shown in Figure 16. The target slip rate is underpredicted on the Cholame section of the San Andreas, and this is systematic across different UCERF3 branches. If the regional MFD constraint is removed (thus allowing the inversion as many ruptures in a given magnitude range as needed to minimize misfits to other data), the slip rate on the Cholame section can be fit nearly perfectly. However, given the limited budget for moderate-size earthquakes throughout the region, this slip

Figure 15. Slip-rate misfits for UCERF3, the UCERF2 mapped solution, and UCERF3 branch averages for each deformation model (see Parsons *et al.*, 2013 for deformation model details). Note that these plots show ratios of the model slip rates to the target, thus deformation models with very small target slip rates appear to have large misfits. This happens in particular with the Average Block Model (ABM), which has slip rates below 10^{-17} mm/yr on some faults. Inset: Histograms of normalized slip-rate misfits show that UCERF3 has smaller slip-rate residuals than UCERF2 on average.

rate cannot be fit without increasing misfits to other constraints. Contributing to the problem is the tendency for many of the branches to overpredict the slip rate on the Parkfield section (and thus preventing more Cholame ruptures from continuing to the north). This is because much of the slip rate for the Parkfield section is taken up in Parkfield $M \sim 6$ ruptures, which are constrained to have a mean recurrence interval of 25 yrs. The mean slip in the Parkfield ruptures ranges from 0.27 to 0.35 m for the Shaw (2009) scaling relation and the Hanks and Bakun (2008) scaling relation to 0.57–0.73 m for the Ellsworth-B (WGCEP, 2003) scaling relation. Thus 10–29 mm/yr is taken up on the Parkfield section just in Parkfield $M \sim 6$ ruptures, which leaves very little, if any, remaining moment for other ruptures on this fault section. The slips given by the scaling relations for the Parkfield ruptures are quite high, and perhaps too high, because the Parkfield ruptures tend to have larger areas and smaller slips than $M 6$ ruptures elsewhere (Arrowsmith *et al.*, 1997). The UCERF3 scaling relations are not necessarily designed to work properly on highly creeping faults (to achieve magnitudes of approximately 6 for the Parkfield earthquakes,

in UCERF2 and UCERF3 average aseismicity was set to 80% and 70%, respectively, on the Parkfield section). In UCERF3 the Parkfield constraint is weighted quite highly relative to the slip-rate constraint; thus the model fits the historical rate of Parkfield earthquakes very precisely at the expense of slip-rate fits; this is preferable because hazard is more sensitive to event rates than slip rates. In future models, adjustments to the scaling relations for highly creeping faults, and possibly relaxation of the Parkfield event-rate constraint, could mitigate this problem.

Paleoseismic mean recurrence intervals (MRIs) for all UCERF3 timing sites are shown in Table 2. Five of the UCERF2 MRIs are outside of the (UCERF3) 95% confidence bounds on the UCERF3 observed rate; all of the UCERF3 MRIs are within the bounds. All but five of the paleoseismic sites are closer to the mean observed MRIs in UCERF3 versus UCERF2.

One readily apparent systematic in the paleoseismic misfits is that paleoseismic synthetics on the southern SAF are lower than the mean observed rates (recurrence intervals in the model are, on average, longer than observed). This is also

Figure 16. San Andreas slip-rate and paleoseismic data fits for (a,b) UCERF2, and (c,d) UCERF3. Paleoseismic mean slip data have been converted to proxy event-rate data. The paleoseismic data in both subplots shown is from the UCERF3 model; UCERF2 mean paleoseismic event-rate data and error bounds were quite different in some locations. UCERF2 event rates plotted with UCERF2 data are shown in the UCERF3 main report.

Table 2
Paleoseismic Mean Recurrence Intervals (MRIs) for UCERF2 and UCERF3 Compared to UCERF3
Paleoseismic Data and Error Bounds

Paleo Site Location (Fault Model 3.1)	Observed MRI (yrs)	Upper 95% Confidence (yrs)	Lower 95% Confidence (yrs)	UCERF2 MRI (yrs)	UCERF3 MRI (yrs)
Calaveras (North), subsection 6	618	1189	321	365	611
Compton, subsection 2	2658	6072	1164	N/A	2766
Elsinore (Glen Ivy), subsection 1	179	262	122	441	244
Elsinore (Julian), subsection 4	3251	128205	881	1278	2526
Elsinore (Temecula), subsection 4	1019	94340	11	708	1077
Whittier, subsection 4	3197	126582	867	1206	2510
SAF* (Big Bend), subsection 6	149	269	82	199	245
Garlock (Central), subsection 6	1435	3293	625	858	1229
Garlock (West), subsection 6	1230	2891	524	882	1264
Green Valley, subsection 3	293	542	159	N/A	334
Hayward (North), subsection 3	318	492	206	196	326
Hayward (South), subsection 3	168	217	129	211	178
SAF (North Coast), subsection 28	870	34364	236	250	401
SAF (Santa Cruz Mts), subsection 5	110	182	66	242	164
SAF (North Coast), subsection 16	306	431	218	247	288
SAF (North Coast), subsection 3	264	408	170	250	254
SAF (Offshore), subsection 8	188	248	142	225	195
Puente Hills, subsection 2	3506	5238	2346	2516	3490
San Gregorio (North), subsection 8	1019	40323	276	622	955
Rodgers Creek, subsection 2	325	785	135	N/A	351
San Jacinto (Anza), subsection 3	312	549	177	256	363
San Jacinto (Superstition Mtn), sub. 0	508	1687	153	401	509
SAF (Carrizo), subsection 1	115	205	64	187	193
SAF (San Bernardino South), subsection 6	205	354	119	312	289
SAF (Coachella), subsection 1	178	321	99	162	259
SAF (Coachella), subsection 1	277	449	171	162	259
SAF (Mojave South), subsection 9	149	225	99	158	180
SAF (San Bernardino North), subsection 2	173	284	106	190	249
SAF (San Bernardino South), subsection 1	205	345	122	244	264
SAF (San Geronio Pass-Garnet Hill), sub. 0	261	409	167	394	306
SAF (Mojave South), subsection 13	106	148	76	174	138

*San Andreas fault.

the case with the UCERF2 model, although the UCERF2 paleoseismic data (see fig. 20 in the UCERF3 main report) had higher means and larger error bars on this fault. The southern SAF paleoseismic data could be better fit by increasing the weight on the paleoseismic data, but this would degrade the slip-rate fit and lead to (further) overfitting of the paleoseismic data in other areas. In fact, the reduced chi-square value for the UCERF3 paleoseismic event-rate fits is 0.72. Because this is less than 1, this means that given the uncertainties, on the whole UCERF3 overfits the paleoseismic data. Although event rates on the southern SAF are systematically lower than the observed rates, these sites do not represent independent data. These sites are all seeing overlapping time periods on the San Andreas. Therefore, if activity is above or below the long-term average on the southern San Andreas, it could be above or below average at all sites. In addition, many events are seen in multiple, neighboring trenches, so the event rates measured at nearby different paleoseismic sites are not independent. Given the correlated data, the spatial systematics in the paleoseismic misfits are expected.

On-fault, off-fault, and total magnitude distributions for UCERF2 and UCERF3 are shown in Figure 17. All branches

Figure 17. California-wide MFDs for the UCERF2 and mean UCERF3 models. The UCERF3 model does not have an overprediction around M 6.5–7 (the bulge) that was present in UCERF2. Note that above M 7.8 the inversion target MFD only specifies an upper bound.

Figure 18. Segmentation on the San Andreas fault (SAF) system for (a, b) UCERF2 and (c, d) UCERF3, for (a, c) ruptures with $M \geq 7$ and (b, d) all supra-seismogenic-thickness ruptures. Note that panels (c) and (d) include the Brawley and Imperial faults on the southern end; the Brawley seismic zone was not a fault-based source in UCERF2. Points show the rate at which neighboring subsections do not rupture together; this rate is normalized by the total rate of ruptures involving those two subsections. Thus, when the line reaches one, there is strict segmentation and no ruptures break through that location.

of UCERF3 fit the target MFD quite well. At high magnitudes the regional MFD tapers away from the $b = 1$ extrapolation (which does not contribute to the regional MFD misfit because the MFD constraint becomes an inequality constraint above $M 7.85$). The background MFD is similar, but smoother, than the background MFD in UCERF2. Most importantly, the inversion methodology eliminates the bulge problem in UCERF2—the overprediction of earthquakes around $M 6.5$ – 7 . Through testing changes to the rupture set, we determined that fitting the MFD constraint and thus eliminating this bulge required multifault ruptures not included in UCERF2. By allowing faults to “link up,” the UCERF2 $M 6.5$ – 7 overprediction has been removed through the accommodation of seismic moment in larger ruptures. It is also worth noting that were UCERF2 methodology used with UCERF3 ingredients, the bulge problem would have gotten

even worse due to additional faults and on-fault moment in UCERF3.

Model Segmentation

We compare segmentation on the SAF between UCERF2 and UCERF3 in Figure 18. In general, UCERF3 allows for many more fault connections than UCERF2; however, due to slip-rate incompatibilities many ruptures that include these connections have low rates relative to ruptures on single, contiguous faults. Thus in UCERF3, some segmentation is a result of the data rather than being strictly imposed through model parameterization. On the SAF, UCERF2 allowed for ruptures between different fault sections (e.g., from SAF Cholame to SAF Carrizo), except through the Creeping section. UCERF3 preserves some of the segmentation of UCERF2,

Figure 19. Fault participation rates for different magnitude ranges, for the mapped UCERF2 model (top panels) and mean UCERF3 (bottom).

but it is not as strict at many of the fault junctions. The segmentation present in UCERF3 is also magnitude dependent.

Model Participation Rates

The rate at which each fault participates in ruptures in a given magnitude range, for UCERF2 and UCERF3, is shown in Figure 19. Because of the added faults and increased fault system connectivity in UCERF3, many more faults participate in larger earthquakes than in UCERF2. For example, in the mapped UCERF2 model only the southern and northern SAF (separately) rupture in $M \geq 8$ earthquakes; the UCERF3 model includes $M \geq 8$ ruptures all along the San Andreas, including (rarely) through the Creeping section, and also on the Garlock, San Jacinto, and Hosgri faults.

Of particular importance for seismic hazard is the distribution of event sizes at a point. Magnitude–frequency participation distributions for selected faults of interest, the South section of the Hayward fault and Mojave South section of the SAF, are shown in Figure 20. The increased connectivity of the fault system allowed in UCERF3 is readily apparent in these examples; multifault ruptures allow for larger magnitudes on many major faults. Further visualizations of ruptures including the Hayward North and SAF Mojave South fault sections are shown in Figures 21 and 22.

Timing Correlations between Neighboring Paleoseismic Sites

The UCERF3 models are directly constrained to fit paleoseismic event rate and mean slip data at points along

faults. One set of data that is not directly included, however, is the correlation of event dates between adjacent paleoseismic sites. This is a nonlinear constraint that is difficult to include directly in the inversion; however, we can check to see how well the inversion matches this independent data.

We compare the fraction of events that are correlated between paleoseismic sites (the total number of events for which the age probability distributions are consistent between sites divided by the number of events) with the fraction of the time those fault subsections rupture together (in paleoseismically visible ruptures) in the inversion. This comparison for the southern SAF is shown in Figure 23. The 95% probability bounds on the data show sampling error only, given the number of events observed; they do not take into account, for example, the probability that closely spaced events in time could have consistent ages. For this fault, UCERF2 and UCERF3 each fall outside of the 95% confidence bounds of the data for one site pair. However, they are each below the bounds, which may be acceptable because not all timing correlations in the paleoseismic data necessarily represent the same event. Some correlated events could be separate events beyond the time resolution of the paleoseismic data. Being above the 95% confidence bounds is more problematic—UCERF2 falls above the bounds at one site on the northern SAF and both UCERF2 and UCERF3 are above the bounds at one site on the San Jacinto fault, although UCERF2 misses the data by more at this site.

In conclusion, Mean UCERF3 fits the paleoseismic timing correlations slightly better than UCERF2, although this may not be true for every individual branch of UCERF3. The

Figure 20. Magnitude-participation distributions for selected faults for UCERF2 and mean UCERF3. These distributions only include ruptures of length greater than the subseismogenic thickness. Note that in general, multifault ruptures in UCERF3 allow for higher magnitude ruptures on many faults. These plots include the aleatory variability in magnitude present in UCERF2; without this magnitude smoothing, UCERF2 results would be less smooth than UCERF3.

Figure 21. Visualization of ruptures involving the Hayward South fault section. (a) Faults that rupture with the Hayward South are colored by the total rate those ruptures occur. (b) A visualization of those ruptures; the trace of ruptures is plotted floating above the fault trace, ordered by rate. Segmentation can be seen as common stopping points for ruptures; this segmentation is not prescribed explicitly by the model, but rather is a consequence of the data and constraints.

Figure 22. Visualization of ruptures involving the Mojave South fault section of the SAF. (a) Faults that rupture with any of the subsections of Mojave South are colored by the total rate those ruptures occur. (b) A visualization of those ruptures; the trace of ruptures is plotted floating above the fault trace, ordered by rate. Common stopping points for these ruptures include the SAF Creeping section and the San Jacinto/SAF boundary.

Figure 23. Comparison of timing correlations between paleoseismic sites and the rate these sites rupture together in UCERF2 and UCERF3 for the southern SAF. As discussed in the text, UCERF3 matches the correlation data slightly better than UCERF2, on average, even though these data are not directly constraining the inversion.

performance of UCERF3 is surprisingly good, given that it outperforms UCERF2 without including this data directly as a constraint.

Misfits for Individual Logic-Tree Branches

We can compute misfits for each constraint of each logic-tree branch (as well as a total misfit that is a weighted average of these); examination of these misfits allows us to determine which branch choices are preferred by the inversion in the sense that the misfits are smaller. In general, branches that have lower misfit are branches in which the total moment balancing is less constrained. For example, better misfits are typically obtained with the UCERF3 spatial seismicity distribution, which leaves more available moment on the faults; similarly, a higher total seismicity rate for the region (which, again, leaves more available moment to distribute on the faults) gives lower misfits. The Average Block Model (ABM; [Parsons et al., 2013](#)) typically gives the worst misfits among the deformation models; it has the largest target moment on the faults (and this makes moment balancing more constrained because UCERF3 models underpredict the total on-fault moment).

These misfits also presumably include information about correlations between logic-tree branches. For example, one scaling relation (that, say, tends to give high slips) may work poorly with a deformation model that has high slip rates but well with a deformation model that gives lower slip rates. Logic trees employed in probabilistic seismic-hazard analysis (PSHA) typically do not have a mechanism to include correlations between branches; that is, the choice on one level of the logic tree is independent of the choice at another level. One could devise, however, a weighting scheme that uses the total misfit for each branch of the logic tree as a likelihood function to update the prior weights determined from expert opinion using Bayes' rule. In this way, single paths through the logic tree that do not fit the data as well could be down weighted.

We are not currently employing such a Bayesian scheme in UCERF3; we are using only *a priori* weights determined by expert opinion. Branch misfits do not necessarily represent model likelihood; for example, the branches using uniform slip along a rupture outperform those branches that use a tapered slip distribution. This is despite the fact that the tapered distribution far better reflects the average slip distribution of ruptures seen in the data. The uniform distribution of slip has lower misfits because it makes it easier to fit the step-like changes in slip rates along faults. These slip-rate changes are probably not physical either; they are simply a modeling simplification. However, in this case, the inversion misfits are not telling us that the uniform slip distribution is correct, just that it is more consistent with other data simplifications in the model.

Additional Results

Many more plots than are feasible to present here have been generated and are available in the link given in [Data and](#)

[Resources](#). Plots available include magnitude distributions on every fault in the model, paleoseismic fits for all faults with paleoseismic data, and hazard results using a range of metrics.

Concluding Remarks

The inversion approach developed for the UCERF3 project relies on data inputs and models from many other project tasks to determine the rates of on-fault ruptures. Many of these inputs will be uncertain and subject to debate. However, it is important to note that all these uncertainties existed in the previous methodology for determining rupture rates as well. The inversion methodology described here eliminates the need for prescriptive assignment of rupture rates. Importantly, we have shown that given a more realistic set of possible ruptures, the data are not sufficient to uniquely constrain individual rupture rates. However, all the hazard metrics we tested were quite robust to particulars of our inversion algorithm, whether it be the equation set weights chosen, cooling schedule, or starting model employed. The stochasticity of the simulated annealing method allows multiple models to be explored, but these models do not differ significantly in terms of hazard implications, especially when compared to the influence of logic-tree epistemic uncertainties.

In California, we have particularly rich geologic, geodetic, paleoseismic, and seismic datasets. The inversion methodology developed for UCERF3 could be applied to other regions as well, provided there were sufficient data to constrain the solution. At minimum, slip rates on faults and a regional MFD are needed; however, it may be most useful where paleoseismic data are available. The inversion has a particular advantage over previous PSHA methodology in that it can simultaneously fit slip-rate and event-rate (for example, paleoseismic rate) datasets.

The grand inversion provides a means to easily update the model as new data become available. In addition, the inversion methodology provides a mechanism to constrain the model by multiple datasets concurrently. This was lacking in UCERF2—expert opinion did not simultaneously satisfy slip rates and event rates, and the magnitude distribution of the final model was inconsistent with the observed distribution and the well-supported assumption that the regional MFD will be characterized by a GR distribution ([Felzer, 2013](#)). The inversion allows all these constraints to be satisfied to the extent that they are compatible.

The inversion can also be used as a tool to determine when a set of constraints is not compatible. Earlier inversion models, for example, showed that some deformation models (since updated) were not compatible with historical seismicity rates and the assumption of full coupling on the faults. This led to significant revisions of the deformation models, which brought the inversion ingredients into better alignment. The GR branch of the model also demonstrates incompatibilities with model ingredients: given the current connectivity of the model, on-fault deformation model

moments, if fully seismic, are not compatible with low historical seismicity rates if we assume an on-fault b -value of 1.0 and assume 58% of seismicity is on the faults (see the UCERF3 main report for more details). These branches have been given zero weight in UCERF3 due to the low on-fault coupling coefficients that they imply. One way these incompatibilities might be mitigated is with increased connectivity, which would allow more on-fault moment in the model.

The degree to which modeled fault connectivity affects the final model results is quite significant. There are several aspects of the model that suggest the connectivity assumed in UCERF3, while greater than connectivity assumed in past models, is an underestimate of the true connectivity in nature. By several metrics, the rates of multifault ruptures in the model underpredict their observed rates. Furthermore, many of the misfits in the model might be improved if maximum magnitudes, particularly on secondary faults, were higher. For example, more moderate-size earthquakes on the southern San Andreas would bring the synthetic paleoseismic rates closer to the observed rates. However, the regional magnitude distribution constraint only allows a finite budget of moderate earthquakes, and many are instead used to match slip rates on faults that cannot rupture in large, $M \sim 8$ earthquakes due to the connectivity assumptions.

Because of the global nature of the applied constraints, all aspects of the model are linked; hence, changing the data or parameterization for one particular fault can affect the solution elsewhere in the fault system. This can be problematic in that errors in the inputs can propagate spatially; however, the flexible, system-level UCERF3 approach allows for sensitivity to connectivity and other model assumptions to be explored. This leaves many exciting avenues for future work, such as investigating trade-offs in the model and testing new hypotheses of earthquake recurrence.

Data and Resources

Additional figures describing details of the UCERF3 model are available at http://pubs.usgs.gov/of/2013/1165/data/UCERF3_SupplementalFiles/UCERF3.3/index.html (last accessed January 2014). The inversion code implemented in OpenSHA, an open-source Java platform for seismic-hazard analysis, is available at www.opensha.org (last accessed June 2013).

Acknowledgments

Computation for the work described in this manuscript was supported by the University of Southern California Center for High-Performance Computing and Communications. The authors also acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin and the Extreme Science and Engineering Discovery Environment (XSEDE) for providing high-performance computing resources under XSEDE Grant Number TG-MCA03S012. We are grateful for helpful reviews from Tom Jordan, Art Frankel, Mark Stirling, Karen Felzer, and Sue Hough, as well as two anonymous reviewers. UCERF3 was sponsored by the California Earthquake Authority, the U.S. Geological Survey (USGS), the California Geological Survey (CGS), and the Southern California Earthquake Center

(SCEC). SCEC is supported in part by the National Science Foundation under Cooperative Agreement EAR-1033462 and the USGS under Cooperative Agreement G12AC20038. This document, with additions that describe Parkfield moment-rate reductions, is registered as USGS Open-File Report 2013-1165-N and CGS Special Report 1792N.

References

- Andrews, D. J., and E. Scherer (2000). Probability of rupture of multiple fault segments, *Bull. Seismol. Soc. Am.* **90**, no. 6, 1498–1506, doi: [10.1785/0119990163](https://doi.org/10.1785/0119990163).
- Arrowsmith, R., K. McNally, and J. Davis (1997). Potential for earthquake rupture and M 7 earthquakes along the Parkfield, Cholame, and Carrizo segments of the San Andreas fault, *Seismol. Res. Lett.* **68**, no. 6, 902–916, doi: [10.1785/gssrl.68.6.902](https://doi.org/10.1785/gssrl.68.6.902).
- Bakun, W. H., and A. G. Lindh (1985). The Parkfield, California, earthquake prediction experiment, *Science* **229**, 619–624.
- Biasi, G. P., R. J. Weldon II, and T. E. Dawson (2013). Appendix F: Distribution of slip in ruptures, *U.S. Geol. Surv. Open-File Rept. 2013-1165-F*, and *Calif. Geol. Surv. Spec. Rept. 1792F*, 41 pp.
- Custódio, S., and R. J. Archuleta (2007). Parkfield earthquakes: Characteristic or complementary? *J. Geophys. Res.* **112**, no. B05310, doi: [10.1029/2006JB004617](https://doi.org/10.1029/2006JB004617).
- Dawson, T. E. (2013). Appendix A: Updates to the California reference fault parameter database: UCERF3 fault models 3.1 and 3.2, *U.S. Geol. Surv. Open-File Rept. 2013-1165-A*, and *Calif. Geol. Surv. Spec. Rept. 1792A*, 18 pp.
- Dawson, T. E., and R. J. Weldon II (2013). Appendix B: Geologic slip-rate data and geologic deformation model, *U.S. Geol. Surv. Open-File Rept. 2013-1165-B*, and *Calif. Geol. Surv. Spec. Rept. 1792B*, 29 pp.
- Felzer, K. R. (2013). Appendix L: Observed magnitude frequency distributions, *U.S. Geol. Surv. Open-File Rept. 2013-1165-L*, and *Calif. Geol. Surv. Spec. Rept. 1792L*, 13 pp.
- Field, E. H., and M. T. Page (2011). Estimating earthquake-rupture rates on a fault or fault system, *Bull. Seismol. Soc. Am.* **101**, no. 1, 79–92, doi: [10.1785/0120100004](https://doi.org/10.1785/0120100004).
- Field, E. H., G. P. Biasi, P. Bird, T. E. Dawson, K. R. Felzer, D. D. Jackson, K. M. Johnson, T. H. Jordan, C. Madden, A. J. Michael, K. R. Milner, M. T. Page, T. Parsons, P. M. Powers, B. E. Shaw, W. R. Thatcher, R. J. Weldon II, and Y. Zeng (2013). Uniform California earthquake rupture forecast, version 3 (UCERF3)—The time-independent model, *U.S. Geol. Surv. Open-File Rept. 2013-1165*, *Calif. Geol. Surv. Spec. Rept. 228*, and *Southern Calif. Earthq. Center Publication 1792*, 97 pp., <http://pubs.usgs.gov/of/2013/1165/> (last accessed March 2014).
- 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).
- Geman, S., and D. Geman (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images, *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-6**, no. 6, 721–741, doi: [10.1109/TPAMI.1984.4767596](https://doi.org/10.1109/TPAMI.1984.4767596).
- Granville, V., M. Krivanek, and J.-P. Rasson (1994). Simulated annealing: A proof of convergence, *IEEE Trans. Pattern Anal. Mach. Intell.* **16**, no. 6, 652–656, doi: [10.1109/34.295910](https://doi.org/10.1109/34.295910).
- 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 pp.
- Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, University of Michigan Press, Ann Arbor, Michigan.
- Ingber, L. (1989). Very fast simulated re-annealing, *Math. Comput. Model.* **12**, no. 8, 967–973, doi: [10.1016/0895-7177\(89\)90202-1](https://doi.org/10.1016/0895-7177(89)90202-1).
- Kirkpatrick, S., C. D. Gelatt, and M. P. Vecchi (1983). Optimization by simulated annealing, *Science* **220**, 671–680.

- Madden, C., D. E. Haddad, J. B. Salisbury, O. Zielke, J. R. Arrowsmith, R. J. Weldon II, and J. Colunga (2013). Appendix R: Compilation of slip in the last event data and analysis of last event, repeated slip, and average displacement for recent and prehistoric ruptures, *U.S. Geol. Surv. Open-File Rept. 2013-1165-R*, and *Calif. Geol. Surv. Spec. Rept. 1792R*, 65 pp.
- Milner, K. R., M. T. Page, E. H. Field, T. Parsons, G. P. Biasi, and B. E. Shaw (2013). Appendix T: Defining the inversion rupture set via plausibility filters, *U.S. Geol. Surv. Open-File Rept. 2013-1165-T*, and *Calif. Geol. Surv. Spec. Rept. 1792T*, 14 pp.
- Parsons, T., K. M. Johnson, P. Bird, J. M. Bormann, T. E. Dawson, E. H. Field, W. C. Hammond, T. A. Herring, R. McCaffrey, Z.-K. Shen, W. R. Thatcher, R. J. Weldon II, and Y. Zeng (2013). Appendix C: Deformation models for UCERF3.3, *U.S. Geol. Surv. Open-File Rept. 2013-1165-C*, and *Calif. Geol. Surv. Spec. Rept. 1792C*, 66 pp.
- Schwartz, D. P., and K. J. Coppersmith (1984). Fault behavior and characteristic earthquakes: Examples from the Wasatch and San Andreas fault zones, *J. Geophys. Res.* **89**, no. 5, 681–5698.
- 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/B0120080006](https://doi.org/10.1785/B0120080006).
- Shaw, B. E. (2013). Appendix E: Evaluation of magnitude-scaling relationships and depth of rupture, *U.S. Geol. Surv. Open-File Rept. 2013-1165-E*, and *Calif. Geol. Surv. Spec. Rept. 1792-E*, 23 pp.
- Szu, H., and R. Hartley (1987). Fast simulated annealing, *Phys. Lett. A* **122**, nos. 3/4, 157–162, doi: [10.1016/0375-9601\(87\)90796-1](https://doi.org/10.1016/0375-9601(87)90796-1).
- Weldon, R. J., II, and G. P. Biasi (2013). Appendix I: Probability of detection of ground rupture at Paleoseismic sites, *U.S. Geol. Surv. Open-File Rept. 2013-1165-I*, and *Calif. Geol. Surv. Spec. Rept. 1792I*, 26 pp.
- Weldon, R. J., II, T. E. Dawson, and C. Madden (2013). Appendix G: Paleoseismic sites recurrence database, 73 pp.
- Wesnousky, S. G. (2008). Displacement and geometrical characteristics of earthquake surface ruptures: Issues and implications for seismic-hazard analysis and the process of earthquake rupture, *Bull. Seismol. Soc. Am.* **98**, no. 4, 1609–1632, doi: [10.1785/B0120070111](https://doi.org/10.1785/B0120070111).
- Wesnousky, S. G., C. H. Scholz, K. Shimazaki, and T. Matsuda (1983). Earthquake frequency distribution and the mechanics of faulting, *J. Geophys. Res.* **88**, 9331–9340.
- 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.* 88-398, 62 pp.
- Working Group on California Earthquake Probabilities (WGCEP) (2003). Earthquake probabilities in the San Francisco Bay region: 2002–2031, *USGS Open-File Rept. 2003-214*, 235 pp.
- U.S. Geological Survey
525 S. Wilson Avenue
Pasadena, California 91106
pagem@caltech.edu
(M.T.P.)
- U.S. Geological Survey
1711 Illinois Street
Golden, Colorado 80401
field@usgs.gov
pmpowers@usgs.gov
(E.H.F., P.M.P.)
- University of Southern California
Department of Earth Sciences
3651 Trousdale Pkwy
Los Angeles, California 90089
kmilner@usc.edu
(K.R.M.)

Manuscript received 28 June 2013;
Published Online 3 June 2014