Bulletin of the Seismological Society of America, Vol. 108, No. 2, pp. 929–945, April 2018, doi: 10.1785/0120170245

Ⓔ

Strong Correlation between Stress Drop and Peak Ground Acceleration for Recent **M** 1–4 Earthquakes in the San Francisco Bay Area

by Daniel T. Trugman* and Peter M. Shearer

Abstract Theoretical and observational studies suggest that between-event variability in the median ground motions of larger ($\mathbf{M} \geq 5$) earthquakes is controlled primarily by the dynamic properties of the earthquake source, such as Brune-type stress drop. Analogous results remain equivocal for smaller events due to the lack of comprehensive and overlapping ground-motion and source-parameter datasets in this regime. Here, we investigate the relationship between peak ground acceleration (PGA) and dynamic stress drop for a new dataset of 5297 earthquakes that occurred in the San Francisco Bay area from 2002 through 2016. For each event, we measure PGA on horizontal-component channels of stations within 100 km and estimate stress drop from *P*-wave spectra recorded on vertical-component channels of the same stations. We then develop a nonparametric ground-motion prediction equation (GMPE) applicable for the moderate (**M** 1–4) earthquakes in our study region, using a mixed-effects generalization of the Random Forest algorithm. We use the Random Forest GMPE to model the joint influence of magnitude, distance, and near-site effects on observed PGA. We observe a strong correlation between dynamic stress drop and the residual PGA of each event, with the events with higher-than-expected PGA associated with higher values of stress drop. The strength of this correlation increases as a function of magnitude but remains significant even for smaller magnitude events with corner frequencies that approach the observable bandwidth of the acceleration records. Mainshock events are characterized by systematically higher stress drop and PGA than aftershocks of equivalent magnitude. Coherent local variations in the distribution of dynamic stress drop provide observational constraints to support the future development of nonergodic GMPEs that account for variations in median stress drop at different source locations.

Electronic Supplement: Figures showing the relation between M_w and M_L, comparison of the ground-motion measurements from this study with the cross-listed records in the Next Generation Attenuation ground-motion database, the validation curve used to select the optimal tree depth for the Random Forest ground-motion prediction equation (GMPE) used in this study, the between-event ground-motion residual is plotted versus: (a) stress drop, (b) magnitude-adjusted stress drop, (c) depth, and (d) depth-adjusted stress drop, a table containing the ground-motion and stress-drop measurements associated with this study, and an example Python notebook.

Introduction

The intensity of earthquake-generated ground motion depends on a complex interaction of source, path, and site effects. Ground-motion prediction equations (GMPEs), in which observed ground-motion amplitudes are statistically modeled as a function of magnitude, source–site distance, and other auxiliary factors, have long served as an empirical basis for understanding the various features that most strongly influence ground-motion intensity. Modern GMPEs have been quite successful in modeling the first-order influence of magnitude scaling and distance-dependent geometrical spreading and attenuation (Douglas, 2003). However,

*Now at Los Alamos National Laboratory, P.O. Box 1663, Los Alamos, New Mexico 87545.

the modulating influence of source complexity on earthquake ground motion remains an area of active research that is of fundamental importance to the scientific understanding of earthquake rupture and the practical implementation of probabilistic seismic hazard assessment (Bozorgnia *et al.*, 2014; Douglas and Edwards, 2016; Baltay *et al.*, 2017). Probabilistic hazard forecasts are particularly sensitive to the between-event variability in the predicted ground motion, because the hazard curve for the time horizons of interest to the development of building codes is dominated by the occurrence of rare high-amplitude ground motions in the upper tail of the inferred probability distribution (Anderson and Brune, 1999; Bommer and Abrahamson, 2006; Cotton *et al.*, 2013; D'Amico *et al.*, 2017).

From a theoretical perspective, one of the most important features driving between-event variability in ground motion is variability in the dynamic stress drop (or stress parameter) of the events in question (e.g., Atkinson, 1990; Atkinson and Morrison, 2009; Baltay *et al.*, 2013; Baltay and Hanks, 2014; Yenier and Atkinson, 2014). Dynamic stress drop, as formulated by Brune (1970), can be estimated entirely from the seismically observed spectra of waveforms recording an earthquake and is typically used as a proxy for the relative proportion of high-frequency energy radiated by the earthquake source during rupture (Atkinson and Beresnev, 1997). Earthquakes with higher stress drops release their energy over shorter time scales, resulting in correspondingly higher peak-moment rates. For larger ($\mathbf{M} \geq 5$) earthquakes, finite-fault effects cause ground-motion amplitudes to saturate with increasing magnitude, and thus the near-source ground-motion amplitude is controlled primarily by stress drop (Baltay and Hanks, 2014).

Various observational and theoretical studies validated this basic correlation between stress drop and ground motion (e.g., Hanks, 1979; Hanks and McGuire, 1981; Boatwright, 1982; Boore, 1983, 2003; Baltay *et al.*, 2013; Yenier and Atkinson, 2015; Lior and Ziv, 2017; Oth *et al.*, 2017). However, the precise functional form of this relation remains poorly understood, in particular in the ways in which the relative influence of stress drop varies in response to other interacting factors, such as magnitude, distance, depth, and source region. Progress toward resolving these questions may provide a considerable step forward in reducing the epistemic uncertainties associated with ground-motion prediction (Anderson and Brune, 1999). However, exploring the influence of stress drop requires careful analyses of extensive joint datasets of ground-motion recordings and source-parameter estimates, which historically have been produced independently and for sparsely overlapping earthquake sequences, with few events in common.

In this study, we focus on the relationship between peak ground acceleration (PGA) and dynamic stress drop for a new dataset of more than 5000 **M** 1–4 earthquakes that occurred in the vicinity of the San Francisco Bay area, California, from 2002 through 2016 (Fig. 1). We use the spectral decomposition technique of Trugman and Shearer (2017) to derive dynamic stress drop estimates from *P*-wave spectra, which are less influenced by attenuation than *S*-wave spectral estimates and are best recorded on the vertical-component channels of near-source broadband and short-period stations. *P*-wave spectral decomposition has proven effective in analyzing the source spectral properties of large datasets of small-to-moderate magnitude earthquakes (Shearer *et al.*, 2006; Trugman and Shearer, 2017; Trugman *et al.*, 2017). We then measure PGA for these earthquakes using the full-waveform horizontal-component records of the same set of stations. Although smaller ($\mathbf{M} < 4$) earthquakes are often excluded from ground-motion analyses, including them here enables us to examine many more events, characterize local variations in median stress drop, and examine how the relation between stress drop and ground motion varies as a function of magnitude.

The analysis of between-event variability in ground motion requires a reference GMPE that models scaling ground-motion intensity as a function of magnitude and external features unrelated to the source, such as distance and local site effects. The set of existing GMPEs designed for earthquakes in California was derived from records of earthquakes that sample different source regions and magnitude ranges than those that comprise our dataset and hence would be inappropriate to extrapolate from for this purpose. We instead apply a novel, data-driven approach based on the Random Forest algorithm (Breiman, 2001) to derive a nonparametric GMPE that can be used to correct the observed ground-motion amplitudes for the interacting effects of geometrical spreading, attenuation, magnitude scaling, and near-site effects. This nonparametric approach has the advantage over techniques based on linear regression in that it can account for complex interactions between these features without risk of overfitting observational noise within the data or introducing systematic trends with magnitude or distance into the residuals from the model prediction (Derras *et al.*, 2012, 2014; Bindi, 2017). We use this framework to define a between-event residual term that quantifies the empirical contribution of each event to the observed variability in ground motion. The unbiased accuracy of the Random Forest GMPE is particularly useful in this context because it permits a detailed examination of the physical correspondence between PGA and the dynamic properties of the earthquake source. Further, whereas we focus in this study on PGA, the Random Forest GMPE framework can be easily extended to predict other metrics of ground motion, such as peak ground velocity, PGA, and response spectral acceleration that are of interest for engineering and hazard applications.

Study Region and Waveform Data

We examine earthquakes that occurred from 2002 through 2016 within a region encompassing the San Francisco Bay area, California (latitude and longitude bounds of [37.0°, 38.5°] and [−123.0°, −121.5°], respectively). This

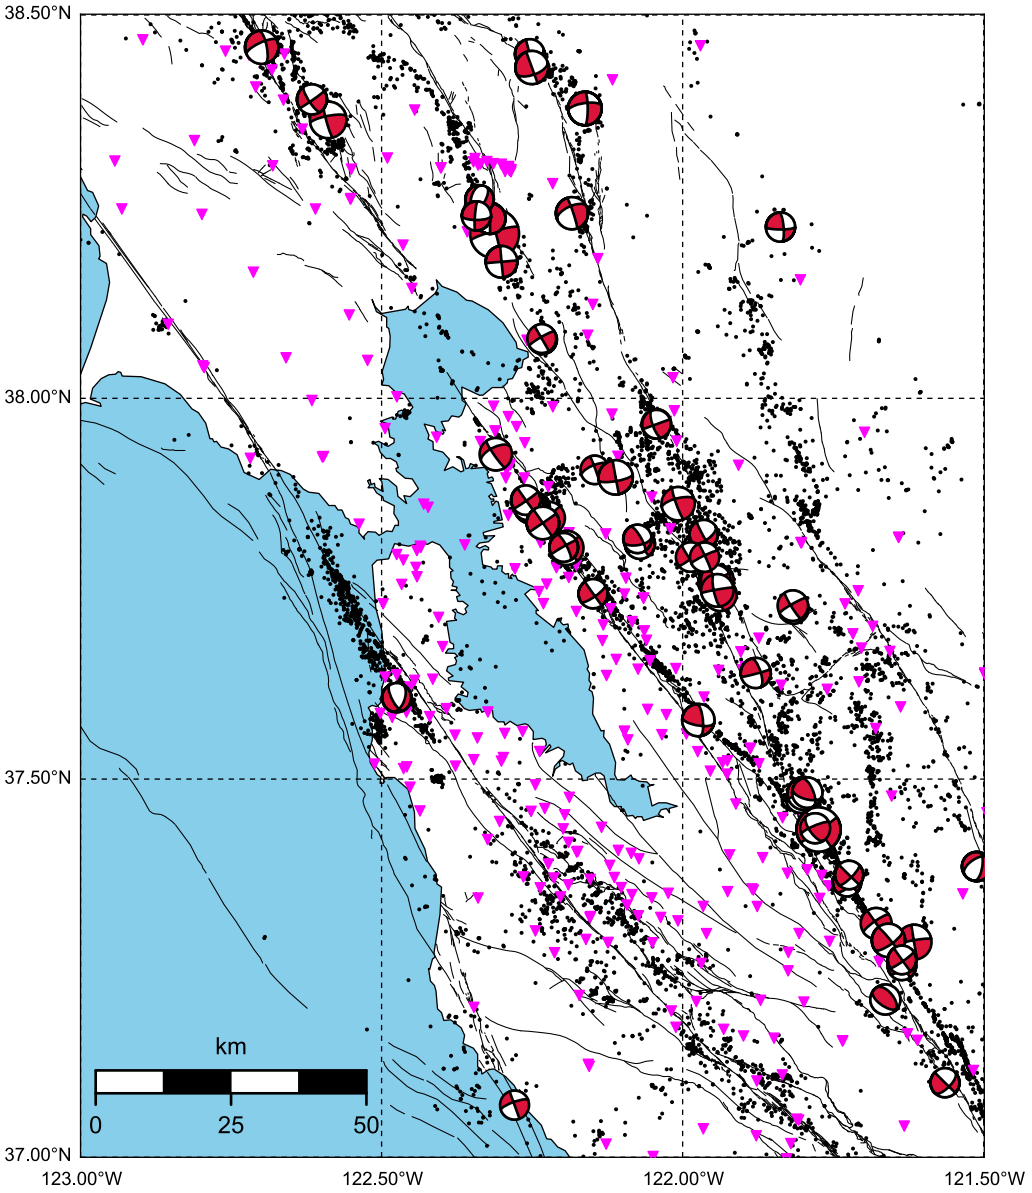

Figure 1. Map view of the San Francisco Bay area study region: **M** ≥ 1.0 seismicity from the relocated catalog of Waldhauser and Schaff (2008); focal mechanisms for **M** ≥ 3.5 events, Northern California Seismic Network station coverage (triangular symbols), and mapped fault structures (see Data and Resources) are shown for reference. The color version of this figure is available only in the electronic edition.

region is of particular interest for earthquake hazard analysis, due to its active seismicity and high population density. Major fault systems within this region include the San Andreas fault near the San Jose and San Francisco metro areas and the Hayward and Calaveras faults in the Oakland/East Bay metro area (Waldhauser and Ellsworth, 2002; Hardebeck *et al.*, 2007; Hardebeck and Aron, 2009). Numerous other subparallel faults of varying seismicity rate also strike through this region (Field *et al.*, 2014). The 24 August 2014 M_w 6.0 South Napa earthquake is the largest earthquake that occurred near the Bay area during our study period (Brocher *et al.*, 2015; Dreger *et al.*, 2015; Ji *et al.*, 2015; Wei *et al.*, 2015), and it triggered hundreds of moderate magnitude aftershocks throughout the West Napa fault system and surrounding region over the subsequent months (Hardebeck and Shelly, 2016; Llenos and Michael, 2017).

Earthquakes that occurred in the San Francisco Bay area during our study period were well recorded by the Northern California Seismic Network (NCSN). We obtain NCSN waveforms archived by the Northern California Earthquake Data Center using the event-data query system (see Data and Resources). We restrict our analysis to earthquakes within our study region with catalog magnitude (typically M_D) ≥ 1.5 and that are listed in the waveform-relocated catalog of Waldhauser and Schaff (2008). *P*-wave spectral estimates are derived from the vertical-component HN, EH, and HH channels at stations with an epicentral distance from the source of less than 100 km. PGA estimates are

derived from the geometric mean of the horizontal components of the same set of stations. The event-data query system provides waveform data in miniSEED format, which we then convert to Seismic Analysis Code (SAC) format to use the SAC waveform processing subroutines (see Data and Resources).

Methods

Source Spectral Analysis and Dynamic Stress Drop

Measurements of the dynamic properties of the earthquake source, including the dynamic stress drop $\Delta\sigma$, are derived from spectra of earthquake waveforms. We obtain estimates of $\Delta\sigma$ for the earthquakes in our study region using the spectral decomposition approach described in detail by Trugman and Shearer (2017) and summarize only the essential points of the algorithm here. As discussed by Atkinson and Beresnev (1997) and others, it is important to understand that the dynamically measured $\Delta\sigma$ based on the Brune (1970) model of earthquake source spectra may differ from the true static stress release on the fault plane, and hence the dynamic $\Delta\sigma$ is more precisely considered a measure of the frequency content of the earthquake source. For this reason, $\Delta\sigma$ is often referred to as the stress parameter, rather than the stress drop, in the ground motion and earthquake hazard literature, but we continue to use the term stress drop for consistency with previous earthquake-source studies.

The conceptual framework underlying the spectral decomposition method can be understood by noting that for sufficiently large and well-recorded waveform datasets, such as that of our Bay area study region, each earthquake is recorded by many stations, each station records many earthquakes, and each source–receiver ray path is sampled (to good approximation) many times. Thus, each of the three main contributions to the observed waveform spectra (the source, the path, and the site) have sufficient observational constraints to be resolved as part of an overdetermined inverse problem. Consider the waveform spectra $d_{ij}(f)$ of event i recorded at station j, which in general is a convolutional product of source, site, and path effects. In the log frequency domain, these three contributions are additive, and thus the d_{ij} at a given frequency f can be decomposed into a linear equation of the form

$$d_{ij} = e_i + st_j + tt_{k(i,j)} + r_{ij}, \tag{1}$$

in which e_i and st_j denote the relative contribution of the source and station, $tt_{k(i,j)}$ is a travel-time-dependent path term that we approximate as isotropic and parameterize into k discrete travel-time bins of 0.5 s width, and r_{ij} is a residual error term.

The spectral decomposition method comprises four basic steps:

1. compute P-wave displacement spectra $d_{ij}(f)$ for each waveform record;
2. decompose these waveform spectra into relative source, station, path, and residual terms by solving equation (1) at each frequency point;
3. estimate the empirical correction spectrum that best captures path effects common to all sources, such as average near-source and near-receiver attenuation;
4. use the corrected source spectra to obtain stress-drop estimates and uncertainties.

We proceed as follows to apply the spectral decomposition workflow summarized above to the earthquakes in our Bay area dataset. For each event, we compute the P-wave spectra from the vertical component of high-broadband and short-period channels (HN, HH, and EH, with a preference for HN when available) of all NCSN stations within 100 km. These spectral estimates are obtained using a multitaper algorithm (Park *et al.*, 1987; Prieto *et al.*, 2009) and a temporal window that begins 0.05 s before the catalog listed P-phase arrival time. Each temporal window bracketing the P-phase arrival has a window length that ranges from a minimum of 1.5 s to a maximum of 4.5 s, with longer windows applied to larger-magnitude events (Ross *et al.*, 2016; Abercrombie *et al.*, 2017; Trugman *et al.*, 2017). We discard clipped waveforms and those with obvious noise spikes, using an automatic detection algorithm based on the observed signal amplitude distribution (Trugman and Shearer, 2017; Trugman *et al.*, 2017). We ensure adequate data quality and station coverage by only considering events recorded by at least six stations in which the average signal-to-noise amplitude is greater than 3 in each of five frequency bands: 2.5–6, 6–10, 10–15, 15–20, and 20–25 Hz. We also limit our analyses to NCSN stations that have recorded at least 50 different events.

We next convert the remaining, quality-controlled spectra d_{ij} to units of displacement and solve equation (1) using an iterative, robust, least-squares inversion algorithm with outlier suppression. This spectral decomposition is effective in isolating the relative contributions of source, path, and site at each frequency point, not their absolute values. To yield source parameter estimates compatible with the Brune (1970) model, it is therefore necessary to apply an empirical correction term that represents path and site effects that are common to all sources, including the average near-source and near-receiver attenuation (relative differences in average site attenuation are incorporated into the station terms st_j themselves). To do so, we use the approach described by Trugman and Shearer (2017) that infers the optimal empirical correction by fitting stacked relative source spectra (Fig. 2a), averaged in bins of spectral moment, to a Brune-type theoretical spectrum of the form

$$\hat{s}(f|\Omega_0, f_c) = \frac{\Omega_0}{1 + (f/f_c)^2}, \tag{2}$$

in which Ω_0 is the long-period spectral moment of each stack (and is proportional to seismic moment) and f_c is the best-fitting corner frequency. We fix the high-frequency spectral fall-off rate to a value of 2 per the canonical ω^{-2} spectral

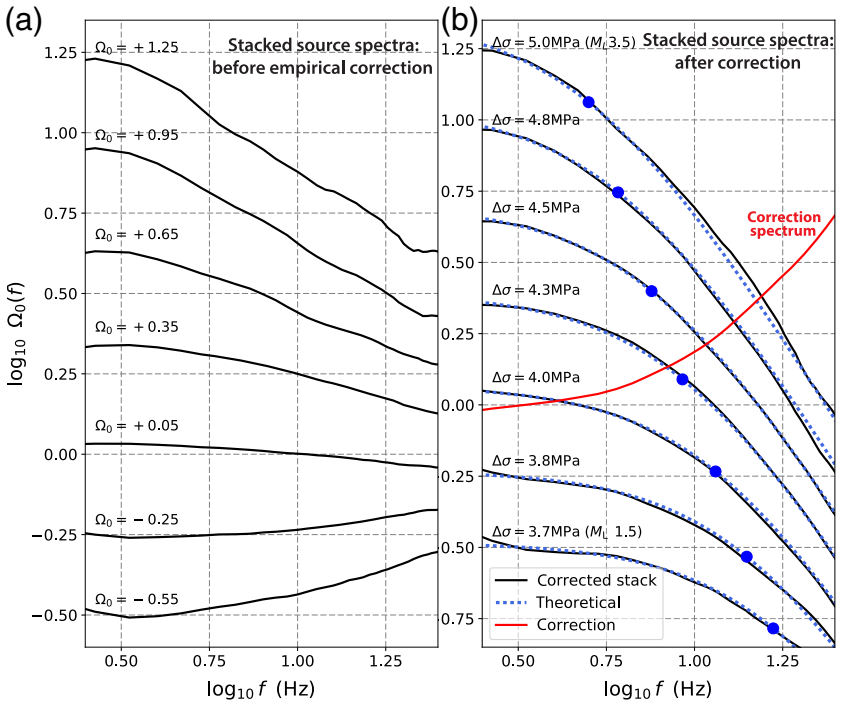

Figure 2. Stacked relative and corrected source spectra from earthquakes in the San Francisco Bay area. (a) Stacked relative source spectra, binned by spectral moment Ω_0 (logarithmic scale), prior to the application of the empirical correction term that accounts for common path effects. (b) Corrected source spectra (solid lines) and comparison to Brune-type theoretical spectra (dashed lines). The implied corner frequency of each stack and the correction spectrum (upward-trending solid line; subtracted from the relative spectra to obtain the corrected spectra) are marked for reference. The color version of this figure is available only in the electronic edition.

model (Aki, 1967; Brune, 1970). Our technique differs from the stacking approach of Shearer *et al.* (2006) in that it does not presume self-similar scaling of the stacked spectra (Fig. 2b). A notable advantage of spectral decomposition is that the inference of the empirical correction term is based entirely upon the shape of the stacked spectra within the 2.5–25 Hz frequency band (for which we verified that the signal-to-noise amplitude is greater than 3) and thus does not require resolving corner frequencies of individual events. We assume a spatially uniform correction term for the results presented in this study, which is justified based on sensitivity tests in which we found that permitting smooth lateral variations in the inferred correction term produces statistically insignificant differences in the final results.

The corrected source spectra are used to estimate two source parameters, seismic moment M_0 and Brune corner frequency f_c, which are required to compute the dynamic stress drop

$$\Delta\sigma = \frac{7}{16} M_0 \left(\frac{f_c}{k\beta}\right)^3 \qquad (3)$$

(Brune, 1970; Madariaga, 1976), in which β is the shear wavespeed, and k is a numerical factor that we set to 0.38, following Kaneko and Shearer (2014). Although the choice of k is a modeling assumption that influences the absolute values of $\Delta\sigma$, the relative values and variations in $\Delta\sigma$ that are of direct relevance to this study are unaffected by this choice.

Seismic moment M_0 and moment magnitude M_w cannot in general be obtained directly from the NCSN catalog, because listed magnitudes are of mixed type but are typically duration magnitude M_D for smaller earthquakes. For stress-drop and GMPE analyses, M_w is more appropriate. To obtain a uniform set of M_w estimates for our dataset, we follow the approach of Shearer *et al.* (2006) that assumes on average seismic moment should be proportional to the long-period spectral moment Ω_0 but this scale factor is unknown. The regression between Ω_0 and M_D is used find the expected value of Ω_0 as a function of M_D. Because M_D and M_w converge at $\mathbf{M} \sim 3.5$ for earthquakes in California (e.g., Hawthorne *et al.*, 2016), we can use this relationship to solve for the scaling factor and thus convert spectral moment to seismic-moment estimates for all earthquakes, and in turn, calibrate a relation between M_w and M_D (Ⓔ Fig. S1, available in the electronic supplement to this article). The smallest earthquakes in our dataset correspond to M_w 1.9.

We obtain f_c estimates for each event using a bounded optimization algorithm that minimizes the root mean square (rms) residual between the corrected relative source spectrum $s_i(f)$ and the Brune theoretical spectrum $\hat{s}_i(f|f_c)$ in the 2.5–25 Hz band in which we measured adequate signal-to-noise. To compute $\Delta\sigma$ given M_0 and f_c, we use a fixed shear wavespeed β in equation (3) for consistency with previous studies of the influence of $\Delta\sigma$ on ground motion (Atkinson and Morrison, 2009; Bozorgnia *et al.*, 2014; Baltay *et al.*, 2015), and, as such, $\Delta\sigma$ should be interpreted as a dynamic source parameter rather than an unbiased estimator of the static stress release. Lastly, we derive uncertainty estimates for M_0, f_c, and $\Delta\sigma$ using a bootstrap-resampling procedure (Trugman and Shearer, 2017) that assesses the variability in the measured source spectra at each station.

Peak Ground Acceleration and the Random Forest GMPE

Because our objective is to analyze the extent to which $\Delta\sigma$ influences ground-motion intensity (PGA), we require (1) measurements of PGA at the stations recording the events comprising our dataset and (2) a model of how median PGA values should be expected to vary as a function of distance, magnitude, site, and other effects unrelated to the dynamic

properties of the earthquake source. In this section, we describe both our procedure to obtain PGA measurements for this dataset and our approach to model PGA using a data-driven GMPE based on a Random Forest regression algorithm.

Although the Next Generation Attenuation-West 2 (NGA-W2) database (Bozorgnia *et al.*, 2014) includes high-quality ground-motion records of a small subset of the events in our dataset, most of the events we consider are unlisted. To compile a complete and self-consistent PGA database, we compute PGA as follows. For each event, we consider horizontal-component records of stations within 100-km source–station distance (consistent with the data selection for the $\Delta\sigma$ computations described in the Source Spectral Analysis and Dynamic Stress Drop section). Using the SAC waveform analysis software, we then demean and detrend each record, remove the instrument response, and convert to units of acceleration where necessary. Our implementation of the SAC transfer command to remove the instrument response applies a band-pass filter with a low-pass transition band from 0.3 and 1.5 Hz and a high-pass transition band from 30 and 40 Hz. The latter corresponds to a significantly higher upper corner than is used for the NGA-W2 database (Bozorgnia *et al.*, 2014), but this modification is necessary to analyze the lower magnitude events that comprise the majority of our dataset. We then use an automated procedure to discard records that have high-amplitude noise spikes or low signal-to-noise rms amplitudes relative to a pre-event time window. We also exclude a small subset of events with origin times that are nearly overlapping (within 45 s spacing of each other), because in these cases the attribution of peak ground motions to one event or another becomes ambiguous. For the remaining quality-controlled records, we compute PGA from the geometric mean of both horizontal components. For most events, the peak amplitude is from the *S* wave or surface wave, whereas our dynamic $\Delta\sigma$ estimates are derived from *P*-wave spectra, which are less affected by attenuation and most clearly recorded on the vertical component. Our PGA estimates are in good agreement with respect to records that are cross-listed in the NGA-W2 database (Ⓔ Fig. S2). Although the nonparametric ground-motion modeling procedure described below permits the use of seismic stations without published metadata, such as site classification or V_{S30}, we note for reference that most of the NCSN stations considered in this study are of National Earthquake Hazard Reduction Program site class B or C, with typical V_{S30} values ranging from 200 to 800 m/s.

Ground-motion intensity measures such as PGA are typically modeled using GMPEs based on linear regression. Such GMPEs predict the expected level of ground motion $\hat{y}$ (logarithmic units) as a linear combination w of input features $\mathbf{X}$ that are thought to influence ground-motion amplitudes:

$$\hat{y} = \mathbf{X}w. \tag{4}$$

The feature matrix $\mathbf{X}$ typically includes (but is not limited to) the magnitude (usually M_w) of the earthquake and a measure of the distance R from the station to the source (usually epicentral distance, hypocentral distance, or finite-fault generalizations, such as the Joyner–Boore distance). In recent years, GMPEs based on linear regression tended to become more complex to accommodate a larger selection of potential input features, higher-order (e.g., quadratic) terms, and interaction terms among predictor features, such as magnitude and distance (e.g., Campbell and Bozorgnia, 2008; Boore *et al.*, 2013; Yenier and Atkinson, 2015).

Typically, the performance of the GMPE is measured in terms of its uncertainty, often denoted as σ and defined as the standard deviation of the distribution of residuals between the observed and GMPE-predicted ground-motion intensity (both measured in logarithmic units). Reduction in the uncertainty of the GMPE is appealing because it leads to improved accuracy in seismic hazard forecasts, especially over long time horizons. However, because σ is measured with respect to the dataset used to fit the model, it is not a measure of the uncertainty with respect to future predictions (e.g., Hastie *et al.*, 2009; Murphy, 2012). Thus, a reduction in the total σ through increased model complexity is only beneficial in cases in which the performance of a model that is developed on the input (or training) dataset is validated by a corresponding decrease in σ with respect to an independent testing dataset (Bindi, 2017; Mak *et al.*, 2017). Further, GMPEs based on standard linear regression techniques may produce systematic residuals as a function of M and R if the true interaction between features deviates from the assumed linear model (Bishop, 2006) or if the number of observations varies significantly with M or R, as is often the case in circumstances with limited data availability.

To study the relative influence of $\Delta\sigma$ on PGA for our study, we require a GMPE with the following properties:

1. it is applicable to the ground motions of moderate magnitude earthquakes within the Bay area and is valid to (at least) 100-km hypocentral distance;
2. it can be used to correct the observed PGA for the potentially nonlinear influence of and interactions between magnitude, distance, and site, which we postulate are the primary features controlling ground motion within this regime unrelated to dynamic earthquake source properties;
3. it does not produce systematic misfits as a function of magnitude or distance that may bias our interpretation of residual PGA in relation to $\Delta\sigma$;
4. it provides a comparable or better GMPE uncertainty σ as compared to established linear regression techniques, as well as a means to assess whether this improved performance is generalizable to an independent dataset;
5. it is robust with respect to outlier data points, which are more common for the recordings of smaller-magnitude events that are an important part of our dataset.

With these requirements in mind, we develop a data-driven GMPE based on a statistical modeling technique known as a Random Forest (Breiman, 2001). Although application of

Random Forests to regression problems is well established within the pattern-recognition and machine-learning community (Bishop, 2006; Geurts *et al.*, 2006; Hastie *et al.*, 2009), to our knowledge this is the first application to the modeling of earthquake ground motions. Random Forest regression models are particularly suitable for our purposes because of their capability to handle nonlinear feature influence and interaction, their robustness with respect to data outliers and co-dependent features, and their ability to automatically assess the predictive performance of the model (Fernandez-Delgado *et al.*, 2014; Louppe, 2014).

Random Forests use an ensemble learning approach in which a set of randomized and decorrelated decision tree regressors are trained independently to predict ground-motion amplitudes, with the final prediction of the Random Forest based on an ensemble average of the individual trees. Individual decision trees are trained through an iterative algorithm that recursively partitions the feature space of predictor variables (Bishop, 2006; Hastie *et al.*, 2009). Although decision trees can be effectively employed in certain classification or regression problems, they are sensitive to noise in the training data and may generalize poorly to independent-testing data (Murphy, 2012). In a Random Forest, each decision tree in the ensemble forest takes a bootstrap resampling of the original training dataset as input, and at each node in the tree a random subset of the features is considered for partitioning. This multiscale randomization and decorrelation of individual tree-based regressors helps reduce the variance in an ensemble prediction through averaging and mitigates the tendency of individual trees to overfit their training dataset and hence provide poor generalization to independent-testing data (Breiman, 2001).

Our implementation of the Random Forest GMPE makes use of the publicly available scikit-learn Python programming package (Pedregosa *et al.*, 2011; see Data and Resources), with some GMPE-specific algorithm modifications that we describe below. Specifically, our input dataset consists of the measured PGA for the subset of well-recorded events for which we obtained estimates of M_w and $\Delta\sigma$ (see the Methods section). We use Random Forest to model the measured PGA values as a nonparametric function G^{rf} of moment magnitude M_w and hypocentral distance $\log R_{hyp}$, and treat the contribution of each event ΔPGA_i and station ΔS_j as random effect terms

$$\log\text{PGA} = \frac{1}{B}\sum_{b=1}^{B} G_b^{rf}(M_w, \log R_{hyp}) + \Delta\text{PGA}_i + \Delta S_j \quad (5)$$

(e.g., Abrahamson and Youngs, 1992; Stafford, 2014). The summation in equation (5) is an ensemble average of B randomized regression trees that are fit to bootstrap-resampled datasets. We apply an iterative procedure to estimate the event and station terms ΔPGA_i and ΔS_j from the mean model residuals, and in practice the algorithm converges after only 2–3 iterations. Defined in this way, the ΔPGA_i values measure the empirical contribution of each event i to the observed ground-motion amplitude (after controlling for magnitude, source–receiver distance, and site), and the distribution of ΔPGA is a measure of between-event variability. It is important to note that to apply the Random Forest GMPE to predict ground motions for a new set of earthquakes, the station terms ΔS_j will be known, but in general the event terms ΔPGA_i will be unknown *a priori*. In this case, the expected total variability in the ground motion from the model prediction σ^2 can be partitioned into between-event variability, defined by the distribution of ΔPGA, and within-event variability, defined by the distribution of GMPE model residuals (e.g., Atik *et al.*, 2010; Baltay *et al.*, 2017).

The fitting procedure of each tree involves repeatedly partitioning the input data in feature space. We use scikit-learn's ExtraTreesRegressor algorithm (Geurts *et al.*, 2006) for this purpose to provide an additional layer of randomization of the partitioning level (rather than greedy optimization), which further helps decorrelate the individual trees. At each node in a given tree, the algorithm randomly selects one of the two features (M_w or $\log R_{hyp}$) for partitioning. Rather than optimizing performance or misfitting to the training data itself, we can use the Random Forest to synthesize an independent-testing dataset without explicitly withholding a subset of input data. To do this, we take advantage of the fact that, for each tree in the Random Forest, bootstrap resampling leaves out approximately one-third of the input data. These excluded data points are known as out-of-bag samples, and each tree has a different subset of in-bag and out-of-bag samples. For a sufficiently large number of trees (we use $B = 200$ in this study), each data point will be left out-of-bag by multiple trees, where it can then be leveraged as part of an independent-validation dataset with minimal computational effort. We can use this property to optimize predictive performance by constraining the maximum depth of the individual trees that comprise the Random Forest. Minimization of the misfit with respect to the set out-of-bag samples yields an optimal tree depth of 18 (Ⓔ Fig. S3), which is itself a conservative choice that errs on the side of underfitting rather than overfitting, due to the slight bias in the out-of-bag misfit-error estimate compared to the true prediction error (Breiman, 2001; Louppe, 2014).

Results

In total, 5297 events within our study region are sufficiently well recorded to meet our quality control criteria for measurements of $\Delta\sigma$ and PGA. Overall, we found that the $\Delta\sigma$ estimates of these events closely follow a lognormal distribution with a median value of 2.95 MPa and a $\log_{10}$ standard deviation of 0.40. We observe a mild increase of $\Delta\sigma$ with M_0 (Fig. 3), which we can quantify in terms of a scaling parameter ϵ_1 that measures the slope of the increase in $\log_{10}\Delta\sigma$ with $\log_{10} M_0$:

$$\log_{10}\Delta\sigma = \epsilon_0 + \epsilon_1 \log_{10} M_0. \quad (6)$$

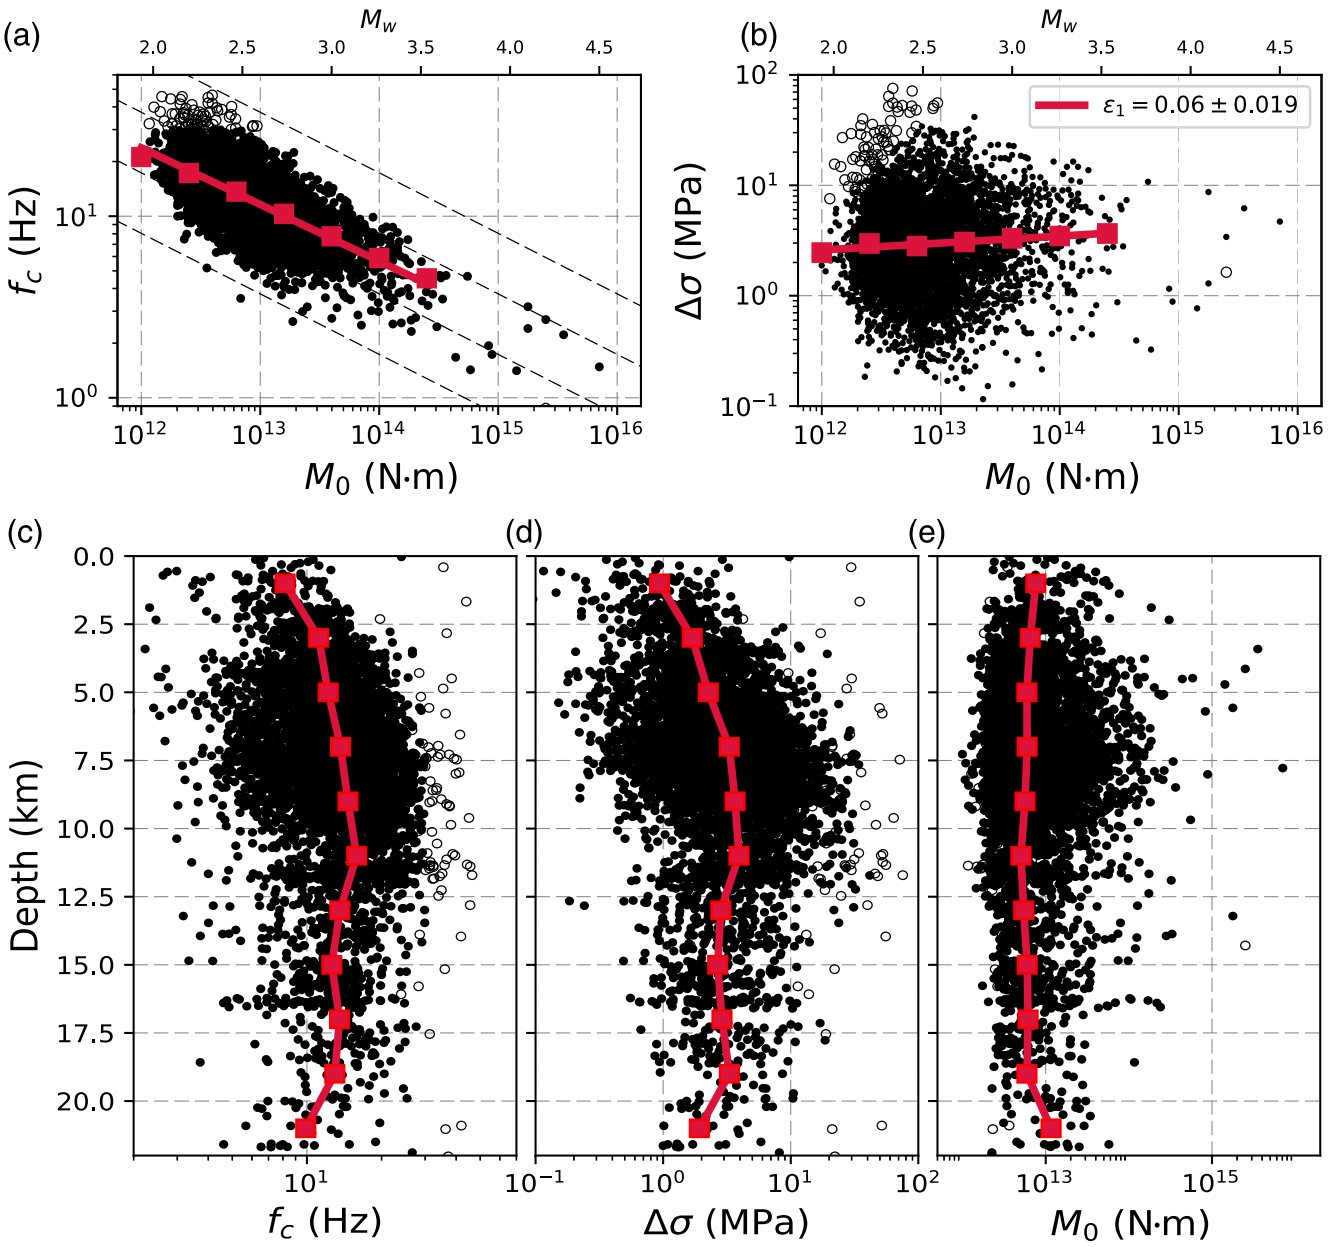

Figure 3. Distribution, scaling, and depth dependence of M_0, f_c, and $\Delta\sigma$. (a,b) Corner frequency f_c and stress drop $\Delta\sigma$ are plotted versus seismic moment M_0. Point markers correspond to measurements of source properties for individual events and the median f_c and $\Delta\sigma$ in M_0 bins of width 0.4 ($\log_{10}$ N·m units) are marked with squares. The best-fitting scaling parameter ϵ_1 for the binned data is plotted with a solid line, with its numerical value and 2σ uncertainty listed in the legend of (b). The dashed lines in (b) correspond to constant $\Delta\sigma$ contours of 0.1, 1, 10, and 100 MPa. Open circles denote events with poorly resolved corner frequencies due to bandwidth limitations ($f_c > 25$ Hz or bootstrap interquartile uncertainty > 5 Hz). (c–e) The depth dependence of f_c, $\Delta\sigma$, and M_0, with the trend of median values in bins of 2-km width marked with solid lines. The color version of this figure is available only in the electronic edition.

A weighted regression analysis based on median values in bins of width 0.4 in $\log_{10} M_0$ yields a scaling parameter $\epsilon_1 = 0.06$, which is consistent with the scaling inferred from the shape of the stacked spectra (Fig. 2) but shallower in slope and hence closer to the self-similar value of $\epsilon_1 = 0$ than any of the five southern California study regions analyzed by Trugman and Shearer (2017).

Both f_c and $\Delta\sigma$ increase slightly as a function of depth within the upper 8 km of the crust (Fig. 3c,d), but we do not observe an analogous depth dependence in M_0 (Fig. 3e). These trends would be consistent with an increase in average rupture velocity with depth, although other systematic variations in rupture characteristics or fault geometry may also contribute. If we examine the spatial patterns of $\Delta\sigma$ in map view (Fig. 4), we observe variations in median $\Delta\sigma$ on regional-length scales, but significant coherence in $\Delta\sigma$ within more localized and spatially compact clusters of events. For example, $\Delta\sigma$ values are relatively high along the peninsular San Andreas fault and the

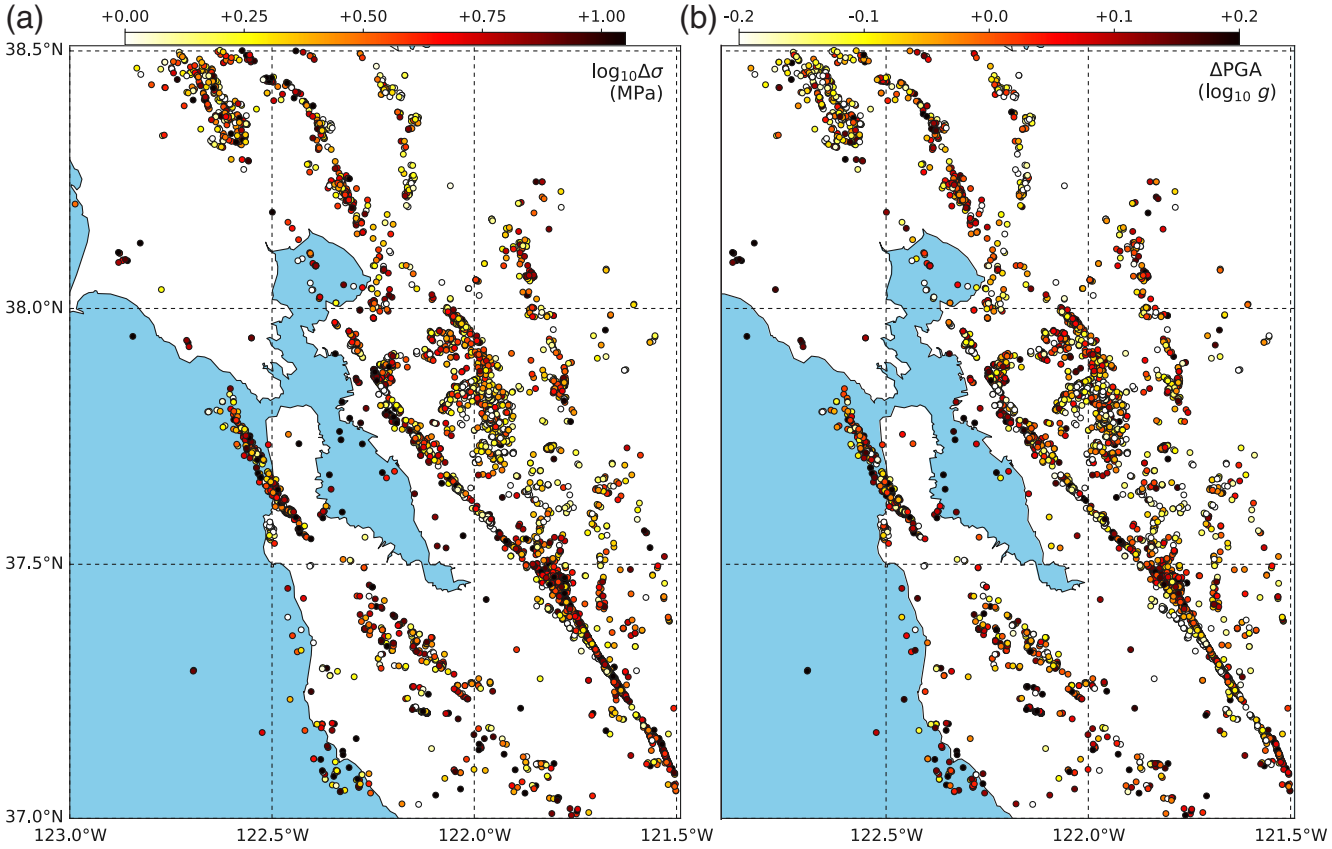

Figure 4. Spatial variations in (a) stress drop $\Delta\sigma$ and (b) between-event ground-motion residual ΔPGA (peak ground acceleration), for events within the San Francisco Bay area study region. The color version of this figure is available only in the electronic edition.

southern portion of the Calaveras fault to the south of its junction with the Hayward fault. Median $\Delta\sigma$ values are markedly lower along the northern segment of the Calaveras fault near the Mt. Diablo thrust and the stepover with the Concord-Green Valley fault. Seismicity along the Hayward and Rodgers Creek faults tends to exhibit intermediate $\Delta\sigma$ values, though with some systematic along-strike variations.

The GMPE modeling procedure described in the Methods section provides estimates of the between-event residual ground motions (ΔPGA) that measure the deviation between the observed and model-predicted ground-motion amplitude. We observe a strong correlation between $\Delta\sigma$ and ΔPGA (Fig. 5), which suggests that variations in dynamic source properties have a significant influence on the between-event variability in ground motion. For example, an increase in stress drop by a factor of 10 would be expected to produce a corresponding increase in PGA by factor of 2.5. The strength of this correlation increases as a function of magnitude, with a correlation coefficient of 0.46 for $M_w < 2.25$, compared to 0.62 for $2.25 < M_w < 2.50$ and 0.74 for $M_w > 2.50$. The magnitude dependence of this correlation is expected, because the effects of attenuation on PGA are more severe for smaller-magnitude earthquakes (e.g, Baltay and Hanks, 2014). It is also notable that the variability in ΔPGA increases significantly with magnitude, with a $\log_{10}$ standard deviation in ΔPGA of 0.11 for $M_w < 2.25$ and 0.13 for $2.25 < M_w < 2.50$, compared to 0.18 for $2.50 < M_w < 3.00$ and 0.26 for $3.00 < M_w < 4.00$. Meanwhile, the variability in $\Delta\sigma$ is nearly constant with magnitude. We consider this effect further in the Discussion section.

In map view, the spatial patterns of ΔPGA mirror those of $\Delta\sigma$ (Fig. 4), with spatial coherence on local length scales but significant variations in median values across the study region. The strength of the correlation remains unchanged, whether one considers the absolute value of $\Delta\sigma$ or a magnitude-adjusted, relative stress drop (Trugman and Shearer, 2017) that corrects for the mild increase in $\Delta\sigma$ with M_0 (Ⓔ Fig. S4b). In contrast, the event term ΔPGA exhibits a rather weak relationship with source depth (Ⓔ Fig. S4c), except perhaps for the shallowest of source depths (0–3 km) where rupture of weaker lithology may tend to reduce ground-motion amplitudes. Likewise, the strong observed correlation between $\Delta\sigma$ and ΔPGA is not simply a result of the observed depth dependence of $\Delta\sigma$, because the correlation remains nearly as strong (Ⓔ Fig. S4d) if we compute a depth adjustment to our $\Delta\sigma$ estimates based on the median trend with depth shown in Figure 3d.

The Random Forest GMPE provides a consistently good fit to the observed distribution of PGA and does not exhibit any significant biases with magnitude, distance, or site (Fig. 6a). To get a sense for the performance of this model

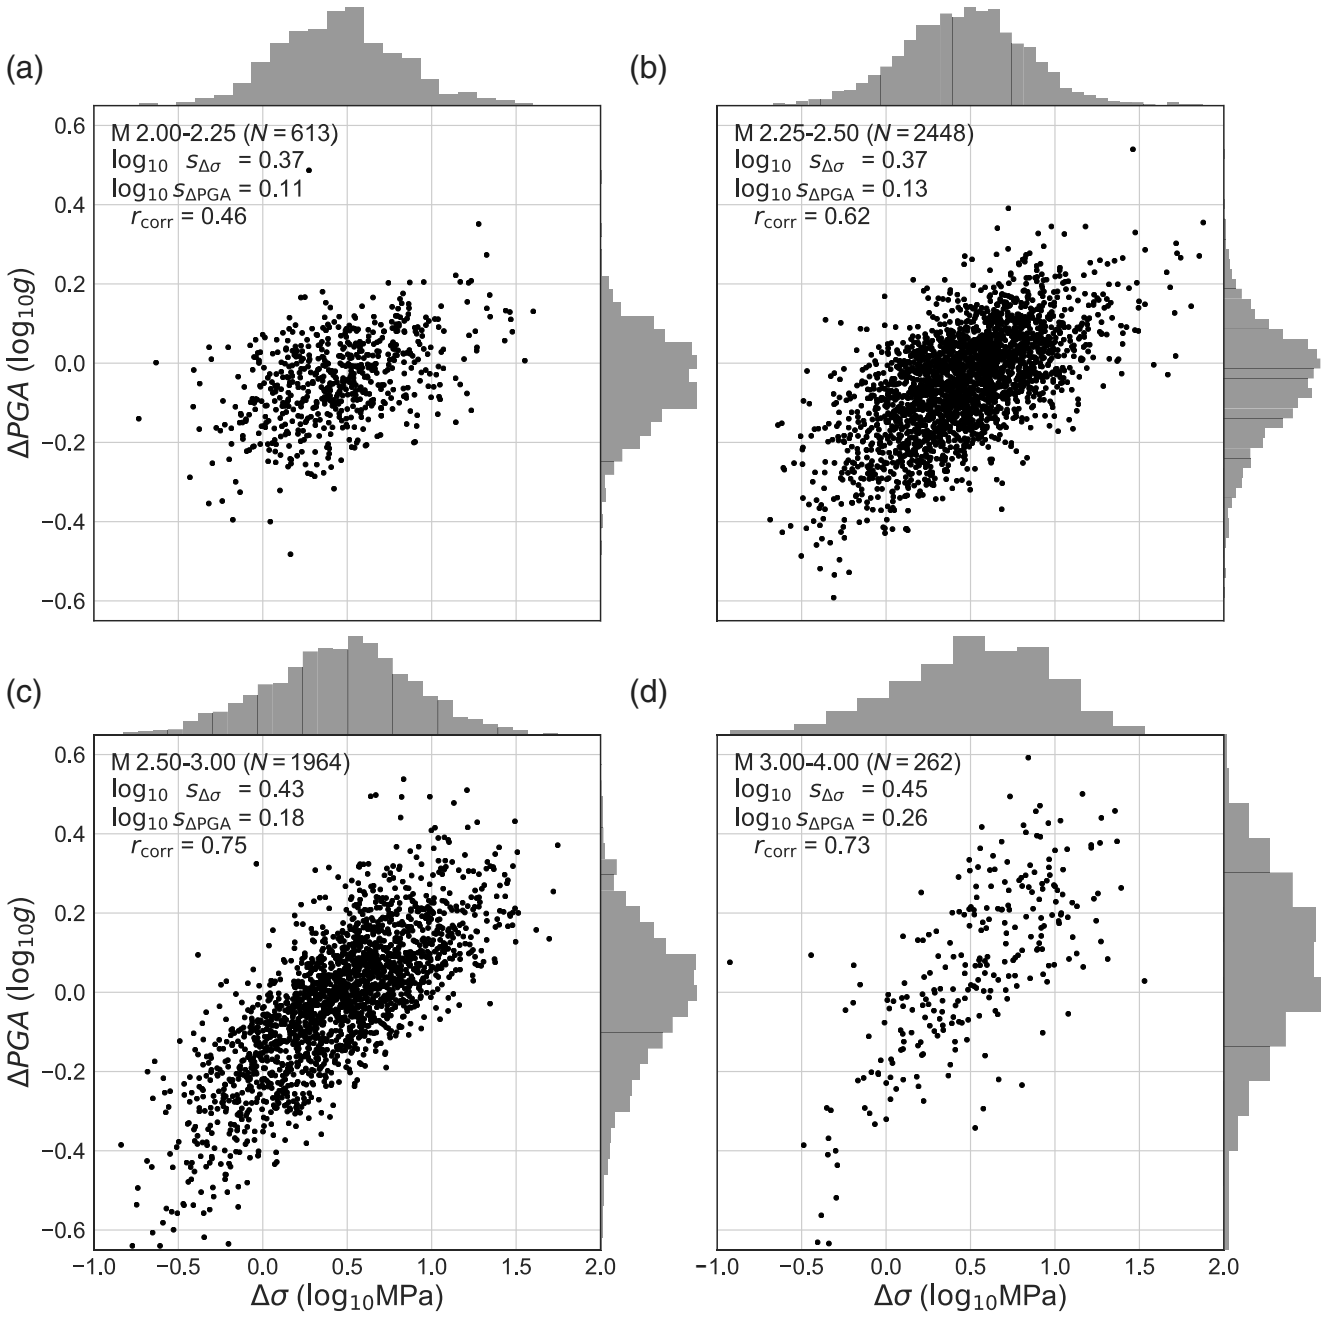

Figure 5. Correlation between stress drop $\Delta\sigma$ and between-event residual ΔPGA for events within the San Francisco Bay area study region. Each panel corresponds to events within four distinct magnitude ranges: (a) $2.00 \le M_w \le 2.25$, (b) $2.25 \le M_w \le 2.50$, (c) $2.50 \le M_w \le 3.00$, and (d) $3.00 \le M_w \le 4.00$. The number of events N, the $\log_{10}$ standard deviations $s_{\Delta\sigma}$ and $s_{\Delta PGA}$, and the correlation coefficient r_{corr} are labeled in the legends.

relative to linear GMPEs, we compare the Random Forest GMPE to two analogous mixed-effect linear regression models (e.g., Campbell and Bozorgnia, 2008; Bindi *et al.*, 2011; Kurzon *et al.*, 2014) of the form

$$\log \hat{y} = a_0 + a_1 M + a_2 \log R + \Delta\mathrm{PGA}_i + \Delta S_j \qquad (7)$$

$$\log \hat{y} = b_0 + (b_1 + b_2 M) \log R + b_3 R + b_4 M + b_5 M^2 + \Delta\mathrm{PGA}_i + \Delta S_j, \qquad (8)$$

in which $\{a_k\}$ and $\{b_k\}$ are linear regression model coefficients and $\Delta\mathrm{PGA}_i$ and ΔS_j are event and station random-effect terms analogous to those in the Random Forest GMPE. Equation (7) is a relatively simple first-order linear regression model (Fig. 6b) that accounts for the first-order influence of magnitude, geometrical spreading, and site, whereas the model described by equation (8) allows for more complexity (Fig. 6c) through higher-order features, interaction terms, and a linear term in R to account for attenuation. In both cases, the

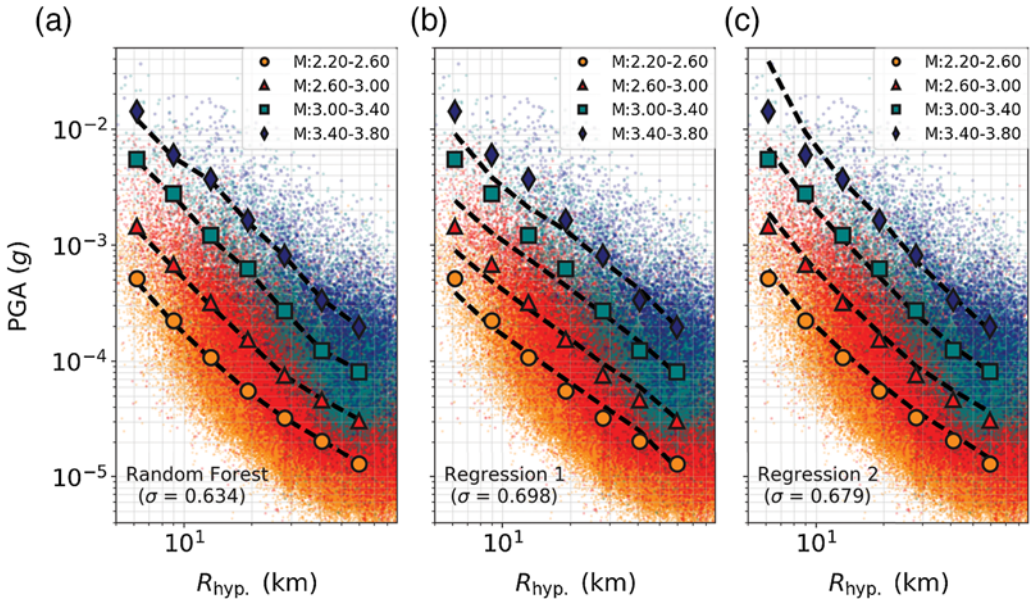

Figure 6. Random Forest ground-motion prediction equation (GMPE) model fit and comparison to mixed-effect linear regression. In each panel, measured PGA is plotted as a function of R_{hypo} (log–log scale). Median values of measured PGA are marked for reference for M_w 4 ranges: **M** 2.2–2.6, 2.6–3.0, 3.0–3.4, and 3.4–3.8. Corresponding model predictions for these magnitude ranges are shown with dashed lines. Both the observations and model predictions have been corrected for station effects based on the station terms inferred using the iterative procedure described in the Methods section. (a) The model fit for the Random Forest GMPE. (b,c) Model fits for the first-order and the second-order linear regression models defined by equations (7) and (8). Note that station-term correction causes the regression models in (b) and (c) to deviate slightly from expected linear and quadratic trends. The root mean square misfit σ is significantly lower for the Random Forest model than for the regression models (0.634, compared to 0.698 and 0.679). The color version of this figure is available only in the electronic edition.

Random Forest GMPE provides a better fit to the data, both in terms of the total misfit σ (0.634 in logarithmic units for the Random Forest, compared to 0.698 and 0.679 for the regression models) and its lack of systematic trends in misfit as a function of M and R (Fig. 6). For this comparison, it is important to note that we guarded against the potential that the Random Forest GMPE overfit the data using the out-of-bag sampling procedure described in the Methods section. Although it may be possible to devise a more complex linear regression model to match its performance, this simple experiment is sufficient to motivate the use of the Random Forest GMPE for our purposes.

From a practical perspective, one question of interest is whether or not there are systematic differences in the distributions of $\Delta\sigma$ or ΔPGA between mainshock and aftershock events. Many established GMPE models either exclude aftershock events entirely from the model fit or explicitly include an adjustment factor that quantifies the difference in expected ground-motion intensity of mainshock and aftershock classes of events (Douglas and Edwards, 2016). However, distinguishing between mainshock events and their associated aftershocks is often difficult in practice, as exemplified by the wide range of plausible declustering algorithms designed to accomplish this task (van Stiphout *et al.*, 2012). The classical space–time windowing method of Gardner and Knopoff (1974) is widely used for this purpose in studies of large-magnitude mainshock–aftershock sequences, in part due to its simplicity and insensitivity to the absence of smaller events in a given ground-motion database (Wooddell and Abrahamson, 2014). However, the Gardner and Knopoff (1974) declustering algorithm is not optimal for use in this study for two reasons. First, the formulas for the space–time windows are designed for larger magnitudes than those that comprise much of our dataset, and an extrapolation leads to poorly defined mainshock–aftershock classifications. Second, the Gardner and Knopoff (1974) algorithm cannot be used directly to group events into individual mainshock–aftershock sequences, due to the nonuniqueness that occurs in cases of overlapping space–time windows.

We instead apply the magnitude–space–time nearest-neighbor method of Zaliapin and Ben-Zion (2013), which has been used effectively to characterize clustered seismicity in numerous previous studies. The nearest-neighbor algorithm is more complicated and has more free parameters to consider than Gardner and Knopoff (1974) declustering, but is appropriate for the magnitude of our dataset and has the additional advantage that it can be used to define individual sequences of events in which the nearest neighbors are closely linked in distance–time–magnitude space. For the purposes of this study, we denote the largest event in each sequence as a mainshock (including those events within singleton sequences) and denote the remaining events as aftershocks. In applying the Zaliapin and Ben-Zion (2013) algorithm, we use the full set of **M** $\geq$ 1.1 earthquakes located within 0.2° latitude and

longitude of our study region that are listed in the relocated catalog of Waldhauser and Schaff (2008) (80,761 events in total). Doing so helps mitigate potential clustering artifacts related to the edge effects of events outside of but adjacent to our study region in space and time, as well as those within our study region but beneath the nominal completeness of our GMPE database.

If one compares the distribution of $\Delta\sigma$ for events classified as mainshocks and aftershocks (Fig. 7a), it is apparent that the mainshocks exhibit systematically higher values of $\Delta\sigma$ than do aftershocks of equivalent moment. The difference between the median $\Delta\sigma$ values of mainshocks and aftershocks is slightly larger for high M_0 events than for low M_0 events, and indeed the two classes are virtually indistinguishable for the smallest events in our dataset. An analogous trend is observed for the between-event residual ΔPGA (Fig. 7b), which is to be expected given the strong correspondence between $\Delta\sigma$ and ΔPGA. These results are broadly consistent with previous studies examining the short-period ground-motion amplitudes of larger-magnitude mainshocks and aftershocks (Campbell and Bozorgnia, 2008; Boore *et al.*, 2013; Wooddell and Abrahamson, 2014; Yenier *et al.*, 2017). The existing literature on possible mainshock-versus-aftershock differences in $\Delta\sigma$ is more equivocal, with various studies finding similar results to those presented here (e.g., Izutani, 2005; Mayeda *et al.*, 2007; Boyd *et al.*, 2017; Cramer, 2017; Sumy *et al.*, 2017), whereas others observe no such trend (e.g., Allmann and Shearer, 2009; Viegas *et al.*, 2010; Baltay *et al.*, 2013; Abercrombie *et al.*, 2017).

As noted above, the nearest-neighbor method of Zaliapin and Ben-Zion (2013) can be used to group events into individual earthquake sequences that are clustered in space and time. We use this framework to examine $\Delta\sigma$ and ΔPGA for the most prominent sequences in our dataset. We limit our analysis to sequences in which the largest event (the mainshock) has $M_w \geq 3.0$ and contain at least 10 other (aftershock) events with measured values of $\Delta\sigma$ and ΔPGA. Mainshock $\Delta\sigma$ estimates are higher than the median value for all 10 sequences (Fig. 7c) that meet this criteria, whereas the same is true for ΔPGA for 9 of the 10 sequences (Fig. 7d). Median sequence values of $\Delta\sigma$ and ΔPGA are strongly correlated, whereas the correlation for within-sequence variability is somewhat weaker but still significant. These observations suggest that analysis of the relation between source parameters and sequence-specific ground-motion data may warrant further consideration for operational earthquake forecasting and real-time hazard assessment (e.g., Jordan *et al.*, 2011; Page *et al.*, 2016).

Discussion

In this study, we present evidence for a strong correlation between dynamic stress drop ($\Delta\sigma$) and residual ground motion (ΔPGA). Although this correspondence has been previously documented for large earthquakes, our study brings to light several new aspects of this relation that have important implications for our understanding of how variability in earthquake source properties can influence earthquake hazard. First, we augment the existing literature with a new, jointly derived dataset of source-parameter and ground-motion estimates for contemporary seismicity in the San Francisco Bay area (Ⓔ Table S1). The events included within this dataset sample a wide magnitude range and include smaller-magnitude events than is typical for ground-motion studies, allowing us to examine how the relationship between $\Delta\sigma$ and between-event variability in PGA varies as a function of magnitude. In so doing, we find that the observed correlation between $\Delta\sigma$ and PGA weakens with decreasing magnitude but is still apparent for the smallest events in our dataset.

The relation between $\Delta\sigma$ and PGA has been studied extensively using theoretical point-source models and random-vibration theory (Boore, 1983, 2003; Baltay and Hanks, 2014; Yenier and Atkinson, 2014), in which the acceleration spectra are typically assumed to be Brune type and thus flat within the frequency band $f_c < f < f_{max}$. In this model, the parameter f_{max} is used as a proxy for the effects of attenuation, especially in the shallow subsurface near the site (Hanks, 1982; Anderson and Hough, 1984). The observed increase in the strength of the correlation between $\Delta\sigma$ and ΔPGA with increasing magnitude is consistent with the predictions of this model, because smaller magnitude events have higher f_c on average and, thus, less bandwidth between f_c and f_{max}. However, the fact that we still observe a significant correlation between $\Delta\sigma$ and ΔPGA, even for the smallest events in our dataset, is unexpected in this context (Baltay and Hanks, 2014) and is perhaps suggestive of the limitations in validity of these theoretical models of acceleration spectra within this regime.

Studies of smaller-magnitude events are also useful in the sense that they can provide a sufficient density of events to examine spatial variations in source properties or ground-motion amplitudes. Earthquakes in our study region exhibit both local coherence in median $\Delta\sigma$ and ΔPGA, as well as systematic variations in these median values over larger regional-length scales throughout the Bay area. If this finding, that the distribution of $\Delta\sigma$ and hence ΔPGA, varies as a function of source region is robust, it would provide an important observational constraint for the reduction of epistemic uncertainty in ground-motion models through the removal of the ergodic assumption (Anderson and Brune, 1999; Atik *et al.*, 2010; Lin *et al.*, 2011; Stafford, 2014; Ameri *et al.*, 2017; Baltay *et al.*, 2017). One relatively simple implementation of nonergodicity would be to allow predicted ground motions to vary as a function of source region, guided by the spatial statistics and correlation-length scales of the $\Delta\sigma$ estimates. This model would effectively assume time stationarity in average source properties and that the spatial distribution of stress drop inferred from moderate seismicity is generalized to the larger events that drive regional earthquake hazard. Both of these topics warrant further study.

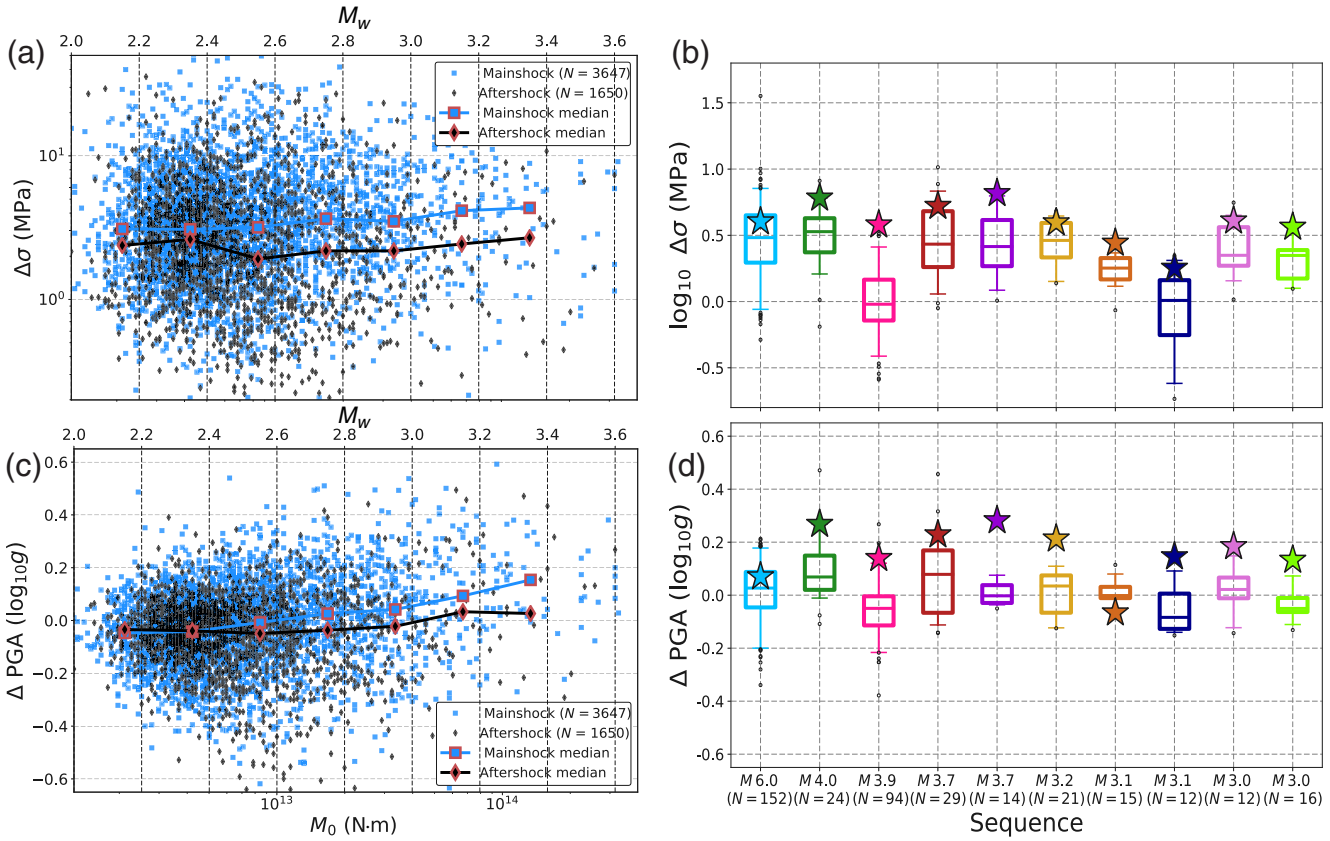

Figure 7. Differences in the distributions of stress drop and ground-motion amplitudes of mainshock and aftershock events. (a) Stress drop $\Delta\sigma$ and (b) between-event residual ΔPGA versus seismic moment M_0. As detailed in the Results section, the neighbors clustering algorithm of Zaliapin and Ben-Zion (2013) is used to classify events as mainshocks (events that are the largest within their respective cluster, marked as square symbols), and aftershocks (nonmainshock events, marked as diamond symbols). Median values of $\Delta\sigma$, binned by M_0, for mainshock and aftershock events are plotted as solid lines. Variability in (c) $\Delta\sigma$ and (d) ΔPGA for sequences with mainshock magnitude $M_w > 3.0$ with at least 10 events. Sequence medians are denoted with a solid horizontal line, whereas the box and whiskers denote the interquartile range (50% confidence interval) and 90% confidence interval, respectively. Mainshock values of (c) $\Delta\sigma$ and (d) ΔPGA are marked with stars, and their corresponding magnitudes are listed along the x axis. The color version of this figure is available only in the electronic edition.

The between-event variability in ΔPGA for events in our dataset is significantly less ($\log_{10}$ standard deviation of 0.16) than would be implied by the between-event variability in $\Delta\sigma$ ($\log_{10}$ standard deviation of 0.40). This paradox has been discussed in many previous studies (e.g., Baltay *et al.*, 2013; Cotton *et al.*, 2013; Oth *et al.*, 2017), and we do not presume to offer a satisfactory resolution, only additional observational constraints. Measurements of ground-motion amplitudes are relatively simple and require fewer modeling assumptions than analogous source–spectral estimates. It is therefore possible that the combined effects of measurement uncertainties (e.g., precision in the resolution of the corner frequency) and uncertainties associated with the model parameterization (e.g., the deviation of the spectra of real earthquakes from the idealized Brune model) may account for a significant fraction of this disparity. The fact that the nominal measurement uncertainties in $\Delta\sigma$, as computed from a bootstrap analysis of the apparent source spectra (Trugman and Shearer, 2017), are comparable to the within-event variability in PGA (median $\log_{10}$ values of 0.219 and 0.214, respectively) suggests that modeling parameterization uncertainty may indeed be an important consideration. Another key difference is that measurements of $\Delta\sigma$ that are derived from source spectra must explicitly account for ray path and near-site attenuation, whereas measurements of ground-motion intensities are effectively dampened because they apply no such correction for attenuation and are further bandlimited through the application of bandpass filtering. Damping PGA variability from attenuation effects should be stronger for smaller earthquakes because of their higher concentration of spectral energy at high frequencies. Our results support the importance of this damping effect, because the variability in ΔPGA is observed to increase significantly with magnitude (Fig. 5), whereas the variability in $\Delta\sigma$ does not. For M_w 3–4 earthquakes, the ratio of log ΔPGA to $\Delta\sigma$ variability is 0.58, only 30% less than the theoretically expected value of 5/6 (Cotton *et al.*, 2013) in the absence of attenuation.

One further issue we can address using the Random Forest GMPE is the possibility that the measured between-event variability in PGA is biased downward, which could

potentially occur if true dynamic source effects are absorbed into the modeled ground motions. To test for this possibility, we generate a synthetic PGA dataset in which we replace the measured between-event residual ΔPGA_i with a synthetic residual δPGA_i equal to the deviation in PGA predicted by the deviation $\delta\sigma_i$ in stress drop from the average value

$$\delta\text{PGA}_i = \frac{5}{6}\delta\sigma_i \tag{9}$$

(Cotton *et al.*, 2013), in which δPGA and $\delta\sigma$ are both measured in consistent logarithmic units. We then apply the algorithm described in the Methods section to fit the synthetic dataset with a Random Forest GMPE and hence derive model estimates of the between-event residuals of the synthetic dataset. The estimated synthetic residuals closely match the input synthetic distribution and are thus consistent with the observed $\Delta\sigma$ variability (Ⓔ Fig. S5), which suggests that we can rule out GMPE modeling artifacts of this type as the cause of the disparity between the variability in ground motion and stress drop.

A final point worthy of discussion is our new approach to GMPE modeling based upon a mixed-effect generalization of the Random Forest algorithm. There are advantages and disadvantages inherent in using this technique, but we argue that the former outweigh the latter for the purposes of this study and perhaps for several other potential applications of ground-motion modeling and estimation. The most appealing aspect of the Random Forest in the context of GMPEs is the ability to use an arbitrarily complex set of features to make robust predictions without the need to specify a parametric form for the relationship between these predictive features and PGA. The Random Forest GMPE described in this study is simplistic in that it only uses two such features, magnitude and source–site distance, to predict PGA, but the same basic framework could easily be extended to include a much wider range of potential predictor features and to make predictions for other ground-motion intensity measures. This would be particularly useful for GMPEs designed to predict ground-motion intensity of large-magnitude events, in which predictive features related to finite-fault rupture, such as hanging-wall effects and directivity or those related to nonlinear site response, could be included in the model without defining a functional form for their influence on ground-motion intensity. One could then easily evaluate which of these input features provides the most fundamental constraints on predicted ground motions using the metrics of feature importance derived from the Random Forest model fit (Breiman, 2001; Pedregosa *et al.*, 2011).

The primary disadvantage to the Random Forest approach is that its lack of a parametric form makes it more difficult to interpret and to apply by external-user groups that are not involved in the model-development process. However, this drawback may be overcome with only modest knowledge of computer programming and statistical theory (see Ⓔ Script S1 for an example). Indeed, nonparametric GMPEs using artificial neural networks have already demonstrated significant promise in ground-motion modeling and hazard applications in Japan (Derras *et al.*, 2012) and Europe (Derras *et al.*, 2014). As is the case for such neural network GMPEs, the nonparametric form of the Random Forest is critical to its performance, because it mitigates the tendency for parametric models to introduce biased predictions or model residuals that vary systematically as a function of magnitude and distance. Although Random Forest regressors are slightly less computationally efficient than neural networks, they do have some powerful advantages in the context of ground motion and hazard estimation. Random Forest models benefit from their simplicity of use and design. Unlike a neural network, there is no need to choose the number of hidden layers or the distribution of neurons within these layers, adjust the layer activation function or regularization, or normalize the input and output data. Instead, the algorithm used to train each tree in the Random Forest automatically searches for and iteratively refines an optimal partitioning in feature space to minimize the predictive misfit. This procedure, when combined with ensemble averaging of randomized and hence decorrelated individual trees, helps ensure that the Random Forest predictions are both robust to data outliers and are valid locally within feature space. Lastly, Random Forest regressors automatically provide a means of predictive validation through the evaluation of the out-of-bag σ, and in this way allow the user to assess its expected performance and uncertainty with respect to future predictions.

Summary

We examine the relation between Brune-type stress drop ($\Delta\sigma$) and ground-motion amplitudes (PGA) for **M** 1–4 earthquakes that occurred near the San Francisco Bay area, California, from 2002 through 2016. We estimate $\Delta\sigma$ for each event using a spectral decomposition method applied to vertical-component *P*-wave spectra recorded by NCSN stations within 100 km. We then measure PGA from horizontal-component channels of the same set of stations and develop a nonparametric GMPE using a mixed-effect implementation of the Random Forest algorithm. We use the Random Forest GMPE to examine between-event variability in ground motion by defining event terms ΔPGA as the average residual between the observed and model-predicted ground motion, given magnitude, distance, and station. The between-event residual ΔPGA exhibits a strong correlation with $\Delta\sigma$, especially for earthquakes with $M_w \geq 2.5$. Estimated values of $\Delta\sigma$ and ΔPGA are slightly higher for mainshocks than for aftershocks in the dataset as whole and for individual earthquake sequences in particular. Local coherence in the spatial patterns of $\Delta\sigma$ and ΔPGA supports future research into nonergodic GMPEs in which the distribution of expected ground motion depends on the location of the source region.

Data and Resources

Waveform data, earthquake catalogs, and station metadata for this study were accessed through the Northern California Earthquake Data Center (NCEDC) and are publicly available from http://service.ncedc.org/ (last accessed August 2017). We use waveform analysis software for data in miniSEED and Seismic Analysis Code (SAC) format that is publicly available as part of the Incorporated Research Institutions for Seismology (IRIS) consortium (http://ds.iris.edu/ds/nodes/dmc/software/, last accessed July 2017). The mapped faulting structures shown in Figure 1 were obtained from the U.S. Geological Survey (USGS) Quaternary Fault and Fold Database for the United States (https://earthquakes.usgs.gov/hazards/qfaults/, last accessed June 2017). The ground-motion data associated with this study are available as Ⓔ Table S1, available in the electronic supplement to this article. We use functions from the scikit-learn Python programming package (Pedregosa *et al.*, 2011) for our analyses in this study, and we provide an example script of our implementation of the Random Forest GMPE in Ⓔ Script S1.

Acknowledgments

This article is based on work supported by the National Science Foundation Graduate Research Fellowship Program (NSFGRFP) under Grant Number DGE-1144086. Additional support was provided by the Southern California Earthquake Center (SCEC) under Grant Number 16020. The authors are grateful to Men Andrin-Meier, Fabrice Cotton, and Associate Editor Adrian Rodriguez-Marek for their constructive and thorough reviews that helped improve and clarify this article. The authors also thank Annemarie Baltay, Peter Gerstoft, and Tara Hutchinson for stimulating discussions that helped guide and focus their initial work.

References

Abercrombie, R. E., S. Bannister, J. Ristau, and D. Doser (2017). Variability of earthquake stress drop in a subduction setting, the Hikurangi Margin, New Zealand, *Geophys. J. Int.* **208,** no. 1, 306–320, doi: 10.1093/gji/ggw393.

Abrahamson, N. A., and R. R. Youngs (1992). A stable algorithm for regression analyses using the random effects model, *Bull. Seismol. Soc. Am.* **82,** no. 1, 505–510.

Aki, K. (1967). Scaling law of seismic spectrum, *J. Geophys. Res.* **72,** no. 4, 1217–1231, doi: 10.1029/JZ072i004p01217.

Allmann, B. P., and P. M. Shearer (2009). Global variations of stress drop for moderate to large earthquakes, *J. Geophys. Res.* **114,** no. B01310, doi: 10.1029/2008JB005821.

Ameri, G., S. Drouet, P. Traversa, D. Bindi, and F. Cotton (2017). Toward an empirical ground motion prediction equation for France: Accounting for regional differences in the source stress parameter, *Bull. Earthq. Eng.* **15,** no. 11, 4681–4717, doi: 10.1007/s10518-017-0171-1.

Anderson, J. G., and J. N. Brune (1999). Probabilistic seismic hazard analysis without the ergodic assumption, *Seismol. Res. Lett.* **70,** no. 1, 19–28, doi: 10.1785/gssrl.70.1.19.

Anderson, J. G., and S. E. Hough (1984). A model for the shape of the Fourier amplitude spectrum of acceleration at high frequencies, *Bull. Seismol. Soc. Am.* **74,** no. 5, 1969–1993.

Atik, L. A., N. Abrahamson, J. J. Bommer, F. Scherbaum, F. Cotton, and N. Kuehn (2010). The variability of ground-motion prediction models and its components, *Seismol. Res. Lett.* **81,** no. 5, 794–801, doi: 10.1785/gssrl.81.5.794.

Atkinson, G. M. (1990). A comparison of eastern North American ground motion observations with theoretical predictions, *Seismol. Res. Lett.* **61,** nos. 3/4, 171–180.

Atkinson, G. M., and I. Beresnev (1997). Don't call it stress drop, *Seismol. Res. Lett.* **68,** no. 1, 3–4, doi: 10.1785/gssrl.68.1.3.

Atkinson, G. M., and M. Morrison (2009). Observations on regional variability in ground-motion amplitudes for small-to-moderate earthquakes in North America, *Bull. Seismol. Soc. Am.* **99,** no. 4, 2393–2409, doi: 10.1785/0120080223.

Baltay, A., J. L. Rubinstein, F. M. Terra, T. C. Hanks, and R. B. Herrmann (2015). Stress drop and depth controls on ground motion from induced earthquakes, *AGU Fall Meeting Abstracts*, San Francisco, California, 14–18 December 2015.

Baltay, A. S., and T. C. Hanks (2014). Understanding the magnitude dependence of PGA and PGV in NGA-West 2 data, *Bull. Seismol. Soc. Am.* **104,** no. 6, 2851–2865, doi: 10.1785/0120130283.

Baltay, A. S., T. C. Hanks, and N. A. Abrahamson (2017). Uncertainty, variability, and earthquake physics in ground-motion prediction equations, *Bull. Seismol. Soc. Am.* **107,** no. 4, 1754–1772, doi: 10.1785/0120160164.

Baltay, A. S., T. C. Hanks, and G. C. Beroza (2013). Stable stress-drop measurements and their variability: Implications for ground-motion prediction, *Bull. Seismol. Soc. Am.* **103,** no. 1, 211–222, doi: 10.1785/0120120161.

Bindi, D. (2017). The predictive power of ground-motion prediction equations, *Bull. Seismol. Soc. Am.* **107,** no. 2, 1005–1011, doi: 10.1785/0120160224.

Bindi, D., F. Pacor, L. Luzi, R. Puglia, M. Massa, G. Ameri, and R. Paolucci (2011). Ground motion prediction equations derived from the Italian strong motion database, *Bull. Earthq. Eng.* **9,** no. 6, 1899–1920, doi: 10.1007/s10518-011-9313-z.

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning, Information Science and Statistics*, Springer, New York, New York.

Boatwright, J. (1982). A dynamic model for far-field acceleration, *Bull. Seismol. Soc. Am.* **72,** no. 4, 1049–1068.

Bommer, J. J., and N. A. Abrahamson (2006). Why do modern probabilistic seismic-hazard analyses often lead to increased hazard estimates? *Bull. Seismol. Soc. Am.* **96,** no. 6, 1967–1977, doi: 10.1785/0120060043.

Boore, D. M. (1983). Stochastic simulation of high-frequency ground motions based on seismological models of the radiated spectra, *Bull. Seismol. Soc. Am.* **73,** no. 6A, 1865–1894.

Boore, D. M. (2003). Simulation of ground motion using the stochastic method, in *Seismic Motion, Lithospheric Structures, Earthquake and Volcanic Sources: The Keiiti Aki Volume*, Y. Ben-Zion (Editor), Pageoph Topical Volumes, Birkhauser, Basel, Switzerland, 635–676, doi: 10.1007/978-3-0348-8010-7_10.

Boore, D. M., J. P. Stewart, E. Seyhan, and G. M. Atkinson (2013). NGA-West2 equations for predicting PGA, PGV, and 5% damped PSA for shallow crustal earthquakes, *Earthq. Spectra* **30,** no. 3, 1057–1085, doi: 10.1193/070113EQS184M.

Boyd, O. S., D. E. McNamara, S. Hartzell, and G. Choy (2017). Influence of lithostatic stress on earthquake stress drops in North America, *Bull. Seismol. Soc. Am.* **107,** no. 2, 856–868, doi: 10.1785/0120160219.

Bozorgnia, Y., N. A. Abrahamson, L. A. Atik, T. D. Ancheta, G. M. Atkinson, J. W. Baker, A. Baltay, D. M. Boore, K. W. Campbell, B. S.-J. Chiou, *et al.* (2014). NGA-West2 Research Project, *Earthq. Spectra* **30,** no. 3, 973–987, doi: 10.1193/072113EQS209M.

Breiman, L. (2001). Random Forests, *Mach. Learn.* **45,** no. 1, 5–32, doi: 10.1023/A:1010933404324.

Brocher, T. M., A. S. Baltay, J. L. Hardebeck, F. F. Pollitz, J. R. Murray, A. L. Llenos, D. P. Schwartz, J. L. Blair, D. J. Ponti, J. J. Lienkaemper, *et al.* (2015). The M_w 6.0 24 August 2014 South Napa earthquake, *Seismol. Res. Lett.* **86,** no. 2A, 309–326, doi: 10.1785/0220150004.

Brune, J. N. (1970). Tectonic stress and the spectra of seismic shear waves from earthquakes, *J. Geophys. Res.* **75,** no. 26, 4997–5009, doi: 10.1029/JB075i026p04997.

Campbell, K. W., and Y. Bozorgnia (2008). NGA ground motion model for the geometric mean horizontal component of PGA, PGV, PGD and 5% damped linear elastic response spectra for periods ranging from 0.01 to 10 s, *Earthq. Spectra* **24,** no. 1, 139–171, doi: 10.1193/1.2857546.

Cotton, F., R. Archuleta, and M. Causse (2013). What is sigma of the stress drop? *Seismol. Res. Lett.* **84,** no. 1, 42–48, doi: 10.1785/0220120087.

Cramer, C. H. (2017). Brune stress parameter estimates for the 2016 M_w 5.8 Pawnee and other Oklahoma earthquakes, *Seismol. Res. Lett.* **88,** no. 4, 1005–1016, doi: 10.1785/0220160224.

D'Amico, M., M. M. Tiberti, E. Russo, F. Pacor, and R. Basili (2017). Ground-motion variability for single site and single source through deterministic stochastic method simulations: Implications for PSHA, *Bull. Seismol. Soc. Am.* **107,** no. 2, 966–983, doi: 10.1785/0120150377.

Derras, B., P. Y. Bard, and F. Cotton (2014). Towards fully data driven ground-motion prediction models for Europe, *Bull. Earthq. Eng.* **12,** no. 1, 495–516, doi: 10.1007/s10518-013-9481-0.

Derras, B., P.-Y. Bard, F. Cotton, and A. Bekkouche (2012). Adapting the neural network approach to PGA prediction: An example based on the KiK-net data, *Bull. Seismol. Soc. Am.* **102,** no. 4, 1446–1461, doi: 10.1785/0120110088.

Douglas, J. (2003). Earthquake ground motion estimation using strong-motion records: A review of equations for the estimation of peak ground acceleration and response spectral ordinates, *Earth Sci. Rev.* **61,** no. 1, 43–104, doi: 10.1016/S0012-8252(02)00112-5.

Douglas, J., and B. Edwards (2016). Recent and future developments in earthquake ground motion estimation, *Earth Sci. Rev.* **160,** 203–219, doi: 10.1016/j.earscirev.2016.07.005.

Dreger, D. S., M.-H. Huang, A. Rodgers, T. Taira, and K. Wooddell (2015). Kinematic finite-source model for the 24 August 2014 South Napa, California, earthquake from joint inversion of seismic, GPS, and InSAR data, *Seismol. Res. Lett.* **86,** no. 2A, 327–334, doi: 10.1785/0220140244.

Fernandez-Delgado, M., E. Cernadas, S. Barro, and D. Amorim (2014). Do we need hundreds of classifiers to solve real world classification problems? *J. Mach. Learn. Res.* **15,** 3133–3181.

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,** no. 3, 1122–1180, doi: 10.1785/0120130164.

Gardner, J. K., and L. Knopoff (1974). Is the sequence of earthquakes in Southern California, with aftershocks removed, Poissonian?, *Bull. Seismol. Soc. Am.* **64,** no. 5, 1363–1367.

Geurts, P., D. Ernst, and L. Wehenkel (2006). Extremely randomized trees, *Mach. Learn.* **63,** no. 1, 3–42, doi: 10.1007/s10994-006-6226-1.

Hanks, T. C. (1979). b values and $\omega^{-\gamma}$ seismic source models: Implications for tectonic stress variations along active crustal fault zones and the estimation of high-frequency strong ground motion, *J. Geophys. Res.* **84,** no. B5, 2235–2242, doi: 10.1029/JB084iB05p02235.

Hanks, T. C. (1982). f_{max}, *Bull. Seismol. Soc. Am.* **72,** no. 6A, 1867–1879.

Hanks, T. C., and R. K. McGuire (1981). The character of high-frequency strong ground motion, *Bull. Seismol. Soc. Am.* **71,** no. 6, 2071–2095.

Hardebeck, J. L., and A. Aron (2009). Earthquake stress drops and inferred fault strength on the Hayward fault, east San Francisco Bay, California, *Bull. Seismol. Soc. Am.* **99,** no. 3, 1801–1814, doi: 10.1785/0120080242.

Hardebeck, J. L., and D. R. Shelly (2016). Aftershocks of the 2014 South Napa, California, earthquake: Complex faulting on secondary faults, *Bull. Seismol. Soc. Am.* **106,** no. 3, 1100–1109, doi: 10.1785/0120150169.

Hardebeck, J. L., A. J. Michael, and T. M. Brocher (2007). Seismic velocity structure and seismotectonics of the eastern San Francisco Bay region, California, *Bull. Seismol. Soc. Am.* **97,** no. 3, 826–842, doi: 10.1785/0120060032.

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning*, Springer Series in Statistics, Springer Publishing, New York, New York, doi: 10.1007/978-0-387-84858-7.

Hawthorne, J. C., J.-P. Ampuero, and M. Simons (2016). A method for calibration of the local magnitude scale based on relative spectral amplitudes, and application to the San Juan Bautista, California, area, *Bull. Seismol. Soc. Am.* **107,** no. 1, doi: 10.1785/0120160141.

Izutani, Y. (2005). Radiated energy from the mid Niigata, Japan, earthquake of October 23, 2004, and its aftershocks, *Geophys. Res. Lett.* **32,** no. 21, doi: 10.1029/2005GL024116.

Ji, C., R. J. Archuleta, and C. Twardzik (2015). Rupture history of 2014 M_w 6.0 South Napa earthquake inferred from near-fault strong motion data and its impact to the practice of ground strong motion prediction, *Geophys. Res. Lett.* **42,** no. 7, 2149–2156, doi: 10.1002/2015GL063335.

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 utilization, *Ann. Geophys.* **54,** no. 4, doi: 10.4401/ag-5350.

Kaneko, Y., and P. M. Shearer (2014). Seismic source spectra and estimated stress drop derived from cohesive-zone models of circular subshear rupture, *Geophys. J. Int.* **197,** no. 2, 1002–1015, doi: 10.1093/gji/ggu030.

Kurzon, I., F. L. Vernon, Y. Ben-Zion, and G. Atkinson (2014). Ground motion prediction equations in the San Jacinto fault zone: Significant effects of rupture directivity and fault zone amplification, *Pure Appl. Geophys.* **171,** no. 11, 3045–3081, doi: 10.1007/s00024-014-0855-2.

Lin, P.-S., B. Chiou, N. Abrahamson, M. Walling, C.-T. Lee, and C.-T. Cheng (2011). Repeatable source, site, and path effects on the standard deviation for empirical ground-motion prediction models, *Bull. Seismol. Soc. Am.* **101,** no. 5, 2281–2295, doi: 10.1785/0120090312.

Lior, I., and A. Ziv (2017). The relation between ground acceleration and earthquake source parameters: Theory and observations, *Bull. Seismol. Soc. Am.* **107,** no. 2, doi: 10.1785/0120160251.

Llenos, A. L., and A. J. Michael (2017). Forecasting the (un)productivity of the 2014 M 6.0 South Napa aftershock sequence, *Seismol. Res. Lett.* **88,** no. 5, 1241–1251, doi: 10.1785/0220170050.

Louppe, G. (2014). Understanding Random Forests: From theory to practice, *Ph.D. Thesis*, Universite de Liege, Liege, Belgium.

Madariaga, R. (1976). Dynamics of an expanding circular fault, *Bull. Seismol. Soc. Am.* **66,** no. 3, 639–666.

Mak, S., F. Cotton, and D. Schorlemmer (2017). Measuring the performance of ground-motion models: The importance of being independent, *Seismol. Res. Lett.* **88,** no. 5, 1212–1217, doi: 10.1785/0220170097.

Mayeda, K., L. Malagnini, and W. R. Walter (2007). A new spectral ratio method using narrow band coda envelopes: Evidence for non-self-similarity in the Hector Mine sequence, *Geophys. Res. Lett.* **34,** L11303, doi: 10.1029/2007GL030041.

Murphy, K. P. (2012). *Machine Learning: A Probabilistic Perspective*, First Ed., MIT Press, Cambridge, Massachusetts.

Oth, A., H. Miyake, and D. Bindi (2017). On the relation of earthquake stress drop and ground motion variability, *J. Geophys. Res.* **122,** no. 7, 2017JB014, 026, doi: 10.1002/2017JB014026.

Page, M. T., N. van der Elst, J. Hardebeck, K. Felzer, and A. J. Michael (2016). Three ingredients for improved global aftershock forecasts: Tectonic region, time-dependent catalog incompleteness, and intersequence variability, *Bull. Seismol. Soc. Am.* **106,** no. 5, 2290–2301, doi: 10.1785/0120160073.

Park, J., C. R. Lindberg, and F. L. Vernon (1987). Multitaper spectral analysis of high-frequency seismograms, *J. Geophys. Res.* **92,** no. B12, 12,675, doi: 10.1029/JB092iB12p12675.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.* (2011). Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12,** 2825–2830.

Prieto, G., R. Parker, and F. Vernon III (2009). A Fortran 90 library for multitaper spectrum analysis, *Comput. Geosci.* **35,** no. 8, 1701–1710, doi: 10.1016/j.cageo.2008.06.007.

Ross, Z. E., Y. Ben-Zion, M. C. White, and F. L. Vernon (2016). Analysis of earthquake body wave spectra for potency and magnitude values:

Implications for magnitude scaling relations, *Geophys. J. Int.* **207,** no. 2, 1158–1164, doi: 10.1093/gji/ggw327.

Shearer, P. M., G. A. Prieto, and E. Hauksson (2006). Comprehensive analysis of earthquake source spectra in southern California, *J. Geophys. Res.* **111,** no. B6, doi: 10.1029/2005JB003979.

Stafford, P. J. (2014). Crossed and nested mixed-effects approaches for enhanced model development and removal of the ergodic assumption in empirical ground-motion models, *Bull. Seismol. Soc. Am.* **104,** no. 2, 702–719, doi: 10.1785/0120130145.

Sumy, D. F., C. J. Neighbors, E. S. Cochran, and K. M. Keranen (2017). Low stress drops observed for aftershocks of the 2011 M_w 5.7 Prague, Oklahoma, earthquake, *J. Geophys. Res.* **122,** no. 5, 3813–3834, doi: 10.1002/2016JB013153.

Trugman, D. T., and P. M. Shearer (2017). Application of an improved spectral decomposition method to examine earthquake source scaling in southern California, *J. Geophys. Res.* **122,** no. 4, 2890–2910, doi: 10.1002/2017JB013971.

Trugman, D. T., S. L. Dougherty, E. S. Cochran, and P. M. Shearer (2017). Source spectral properties of small to moderate earthquakes in southern Kansas, *J. Geophys. Res.* **122,** no. 10, 8021–8034, doi: 10.1002/2017JB014649.

van Stiphout, T., J. Zhuang, and D. Marsan (2012). Seismicity declustering, *Community Online Resource for Statistical Seismicity Analysis*, doi: 10.5078/corssa-52382934.

Viegas, G., R. E. Abercrombie, and W.-Y. Kim (2010). The 2002 M 5 Au Sable Forks, NY, earthquake sequence: Source scaling relationships and energy budget, *J. Geophys. Res.* **115,** no. B7, doi: 10.1029/2009JB006799.

Waldhauser, F., and W. L. Ellsworth (2002). Fault structure and mechanics of the Hayward fault, California, from double-difference earthquake locations, *J. Geophys. Res.* **107,** no. B3, doi: 10.1029/2000JB000084.

Waldhauser, F., and D. P. Schaff (2008). Large-scale relocation of two decades of northern California seismicity using cross-correlation and double-difference methods, *J. Geophys. Res.* **113,** no. B8, doi: 10.1029/2007JB005479.

Wei, S., S. Barbot, R. Graves, J. J. Lienkaemper, T. Wang, K. Hudnut, Y. Fu, and D. Helmberger (2015). The 2014 M_w 6.1 South Napa earthquake: A unilateral rupture with shallow asperity and rapid afterslip, *Seismol. Res. Lett.* **86,** no. 2A, 344–354, doi: 10.1785/0220140249.

Wooddell, K. E., and N. A. Abrahamson (2014). Classification of main shocks and aftershocks in the NGA-West2 database, *Earthq. Spectra* **30,** no. 3, 1257–1267, doi: 10.1193/071913EQS208M.

Yenier, E., and G. M. Atkinson (2014). Equivalent point-source modeling of moderate-to-large magnitude earthquakes and associated ground-motion saturation effects, *Bull. Seismol. Soc. Am.* **104,** no. 3, 1458–1478, doi: 10.1785/0120130147.

Yenier, E., and G. M. Atkinson (2015). Regionally adjustable generic ground motion prediction equation based on equivalent point source simulations: Application to central and eastern North America, *Bull. Seismol. Soc. Am.* **105,** no. 4, 1989–2009, doi: 10.1785/0120140332.

Yenier, E., G. M. Atkinson, and D. F. Sumy (2017). Ground motions for induced earthquakes in Oklahoma, *Bull. Seismol. Soc. Am.* **107,** no. 1, 198–215, doi: 10.1785/0120160114.

Zaliapin, I., and Y. Ben-Zion (2013). Earthquake clusters in southern California. I: Identification and stability, *J. Geophys. Res.* **118,** no. 6, 2847–2864, doi: 10.1002/jgrb.50179.

Institute of Geophysics and Planetary Physics
Scripps Institution of Oceanography—UC San Diego
8800 Biological Grade, Room 307
La Jolla, California 92037
dtrugman@lanl.gov

Manuscript received 30 August 2017;
Published Online 6 March 2018