

Response to Reviewer 2, M. Oskin, “The impact of earthquake cycle variability on neotectonic and paleoseismic slip rate estimates”

Richard Styron

Reviewer comments denoted as [Ox] where x is the comment number. Comments referenced as [Ax] or [Bx] refer to those from Reviewer 1 and 3, respectively.

Review by Michael Oskin

[O1] This paper presents a useful thought experiment on the impact of earthquake cycle variability on measured slip rates, and concludes that the convergence on the expected value is a function of the coefficient of variation. Overall this is a sensible conclusion. Underpinning this analysis are four assumed variants of earthquake recurrence, and a function to express the variability of slip per event.

[O2] I would like to see the effect of COV isolated from the slip per event distribution (use 1m slip for every event).

I performed this experiment; the figure with the slip rate results is shown below for comparison with Figure 5 in the paper, and are included in a new document in the supplemental material. The most relevant figure is also shown here (Figure 1)

The differences between the results of this experiment (fixed per-event displacement of 1 m) and of the numerical experiment performed in the manuscript are basically that these results are less smooth, but the total variance at any point in time (x -axis) is less. This is simply due to removing the stochasticity from one of the two variables in the system. The relative spread in the data and convergence rates are unchanged. This is to be expected as the variability in the per-event displacement is the same for all recurrence distributions, so even though it is a random variable in the simulations, it is not an experimental variable.

What I find the most interesting about this experiment is that the fluctuations in the




Figure 1: Slip rates with no displacement variability

estimated slip rates show very clearly the mean earthquake cycles once the noise from the per-event displacements has been removed. It is clear that these are kind of damped or averaged out after ~ 7 earthquake cycles. Nonetheless, though this is a cool pattern to see, I don't think it adds enough insight to be worth including in the manuscript. I have added the figure and table that show the slip rate variation through time to a new supplemental document, accompanied by a brief discussion.

Changes: New experiment added to new supplemental materials, and brief discussion in the main manuscript.

[O3] I would also like to see a more quantitative comparison of COV and a convergence on the mean to back up the assertion that COV of the distribution is more important than the distribution itself.

This comment is addressed in the response to [O18] below.

[O4] The paper would be improved by a more quantitative, empirical basis and discussion of physical processes that may drive such recurrence behavior.

I don't understand how the work could be more quantitative—it's a purely numerical study.

As for making it more empirical, there are several ways to do this:

1. *Using empirical distributions for earthquake recurrence and slip.* There isn't a lot of consensus on empirical (non-parametric) earthquake recurrence distributions; most of the community prefers parametric distributions such as the lognormal, Weibull or Brownian Passage Time distributions. As discussed in the paper by Matthews et al. (2002) that is the most prominent introduction to the Brownian Passage Time distribution, given the very small number of samples for earthquake recurrence that we will have for a given section of faulting, it is impossible to discriminate between these distributions, so going with the lognormal (as I have done in this study) is justifiable on empirical grounds as well as practical ones (it is familiar, implemented in many programming environments, and easy to manipulate). With regards to empirical slip rate distributions, I have added an experiment that uses one, which is explained in more detail in response to comment [O6a] below.
2. *Going through the literature and evaluating studies that claim that slip rates have changed (or have not changed) in light of the work presented here.* I considered this, and in fact a major motivation for me to begin the analysis was my skepticism over some recent literature claiming major slip rate changes over relatively short timespans. However, I opted not to do this in the paper, mainly because I didn't want to pick fights. There is a bit of a paradox here: If I claim that the conclusions in a paper claiming secular slip rate changes (or fluctuations) are actually due to aleatoric variability in earthquake recurrence, I will probably anger those authors and decrease the likelihood that they will consider these results in subsequent work. I'd rather write a more toothless

paper that doesn't single out any given researchers, and is therefore a bit easier to swallow by all.

3. *Incorporating measurement uncertainty.* Measurement uncertainty is a very large factor that affects the results of any slip rate measurements, and I fully agree with comment [O6b] that it is in most cases underreported, both in offset measurements (as that comment references) and in geochronologic dating of any sort. I chose to leave it out of this paper because I really wanted to focus on the aleatoric variability, which is generally *underappreciated* as opposed to *underreported*.

Per a discussion of the physics and mechanics behind recurrence behavior: I have added a short discussion, but I don't want to really dig into the topic, for two reasons:

1. The intended audience for the paper is not only crustal deformation researchers, but others in the seismic hazard community as well—this includes engineers, geotechnical workers, analysts in the insurance industry, etc. In my experience as a member of this community, many others are only interested in these sorts of phenomena to the degree that they are consequential and actionable; their intellectual interests are often oriented towards their fields of expertise (structural engineering, ground motions, human and economic exposure, etc.). I want this paper to be a straightforward reference for how to evaluate slip rate data in light of aleatory variability that is not tied down in jargon or linked to specific geological or geophysical models or ideas that may not stand the test of time. Because of how variable and poorly-understood earthquake recurrence and fault interaction phenomena are, an in-depth discussion without resolution may well be off-putting to much of the audience that I would like to read this paper.
2. I don't think that we have a great understanding of the real mechanisms, yet. There are a variety of mechanisms under consideration (e.g., co- and post-seismic elastic and viscoelastic Coulomb stress changes, stress transients, dynamic triggering, pore fluid pressure fluctuations, fluctuations in the frictional failure threshold on a fault) in addition to actual secular changes in tectonic loading rates. The time-dependent mechanisms (particularly post-seismic processes) often show different behavior with regard to whether they are 'spun-up' and at a dynamic equilibrium, or not. And all of these mechanisms are necessarily linked to uncertainty as to how (and where) faults are loaded to begin with—whether the loading is in the elastic crust, in the viscoelastic/viscous mid- or lower crust (in a continuum style), or on a discrete creeping dislocation down-dip of the brittle fault. There is a big range of scientific opinion on all of these questions. As a community we are begging for a big review paper to at least concatenate and organize these ideas and potentially test them or at least sort them into compatible vs. mutually exclusive sets for future testing. But we don't have that right now, so the topic is kind of a big mudhole. I will dip my toe in but I really want to avoid falling in for the purposes of this paper.

Changes: New experiment with empirical slip distribution, and discussion of physical mechanisms behind aleatory recurrence variability.

[O5] There is a literature of ideas to draw upon, such as post-seismic fault reloading (Kenner and Simons, 2005), earthquake super cycles (Sieh et al., 2008; Weldon et al., 2004), isolated versus fault-network behavior (Berryman et al., 2012). Some of these ideas are discussed briefly but need more explanation.

I have added a short discussion (two paragraphs) on the topic, but as noted in my response to [O4] I don't think that a more full discussion is warranted. I want this paper to be a simple, easy-to-digest paper and I think that a long and necessarily unsatisfying discussion (as we don't have answers yet) on the mechanisms behind recurrence variability will be an obstacle, and many readers will just put the paper down.

Changes: discussion added.

[O6a] Likewise one could examine actual earthquake slip distributions (not landform offsets of historic events, which convolve landscape processes with tectonic slip) to develop an empirical basis for the slip function.

Such a distribution is given by Biasi and Weldon (2009), following work done by Hemphill-Haley and Weldon (1999). It is a bit different than the lognormal distribution used in that the probability of relatively low values (zero or near-zero) is higher than in a lognormal distribution. The sample COV of this distribution is 0.67, slightly lower than the lognormal slip distribution used in the paper, with a COV of 0.75.

As an experiment, I have re-done the simulation sampling from this distribution; the data are given as 1313 discrete points from earthquakes worldwide, normalized to the mean slip per event. I have sampled randomly from this finite set, with replacement, instead of interpolating the set into a continuous distribution and sampling from that. The results are in the new supplemental materials (Figures S1 and S2, Table S1, and some discussion).

The results are nearly indistinguishable.

Changes: new experiment added.

[O6b] Some of the scatter in slip distributions is likely due to under-reported measurement uncertainty (Gold et al., 2013), and thus the cancellation of this error over multiple earthquakes should let cumulative slip converge more quickly than may be predicted from the author's model.

Either I don't understand this comment (which is quite possible) or it is a bit misapplied. The only reading of this comment under which one expects faster convergence than I have modeled is if the reviewer believes that measurement error is some component of the total variability represented. But that is not the case in the modeling;

these distributions are taken to represent only aleatory variability and the study is performed with assumptions of zero measurement error.

Changes: None.

[O7] Page 1, line 4. The open interval problem is well known and attempts to quantify it do exist on case-by-case basis.

Yes, and some of these cases are cited in the introduction. However, the open-interval problem simply deals with the uncertainty in a single recurrence interval (the present one), and not the variability that is present throughout all of the closed earthquake intervals that have contributed to the measured offset; this larger issue is the topic of the manuscript.

Changes: None.

[O8] Page 1, line 13. It seems odd to characterize uncertainty due to a random distribution as epistemic. Isn't this unreported aleatory uncertainty?

Aleatory and epistemic uncertainty are not mutually exclusive categories. Much epistemic uncertainty results from aleatoric variability, particularly when the underlying distributions that characterize the aleatoric variability are not known.

This is one of those instances: The framing of the situation is that one has made a single slip-rate 'measurement' (net slip / time) without knowledge of where the fault is in its earthquake cycle, what the past earthquake history is, and what the distributions of slip and recurrence are for that fault to begin with. Thus the condition is one of ignorance, i.e. epistemic uncertainty, and this section of the study shows how to approximately quantify this uncertainty under different assumptions of the slip and recurrence distributions.

Changes: None.

[O9] Page 1, line 20: Why is marker in quotes?

I wanted to declare that it was a technical term and not a word that I arbitrarily applied to the situation. But this isn't necessary.

Changes: Quotes removed.

[O10] Page 2, line 5. afterslip and creep also contribute.

Truth.

Changes: Afterslip and creep added to sentence.

[O11] Page 2, line 10-11. Awkward sentence. Break into two.

Ok.

Changes: Sentence broken.

[O12] Page 2, line 13 and other citations: persistent use of ‘e.g.’ after citing only one or two articles is poor form and makes this reader think that the author has not adequately explored the literature.

The use of ‘e.g.’ denotes that the given citations are not authoritative or canonical in the sense that the cited works are where the concepts given are first introduced or best developed, as this isn’t true. The cited works are generally just modern, high quality studies that exemplify the topic at hand.

I don’t really care what readers may think of the depth of my scholarship.

For what it’s worth, ‘e.g.’ should be before the references but was placed after by a LaTeX bug that I hadn’t diagnosed.

Changes: None.

[O13a] Page 3, line 11. This is not the correct definition of an exponential / poisson distribution. There is no prescribed number of events, only a prescribed time-independent probability.

The time-independent probability is the mean rate of events. The mean rate of events is the mean number of events that occur within some time interval.

Obviously with finite sample sets (of time, or of events) there will be some variation—otherwise I wouldn’t have written the paper.

Nevertheless, the statement actually made is that the spacing between uniform random samples in some interval is characterized by an exponential distribution, which is true. It is not stated that this is the definition of the distribution.

Changes: None.

[O13b] It is also worth noting that this is physically unrealistic at short time intervals because it violates elastic rebound.

Elastic rebound is a hypothesis, not a law, and is phenomenological instead of physical in nature. It is unfortunately a step removed from the modern understanding of the mechanics of earthquakes, which are based around stress, not strain. These map to each other nicely in the case of elastic and Newtonian viscoelastic rheologies, but not as nicely with rate-dependent rheologies, which are often considered the best characterization of the lower crust and upper mantle (e.g., Hetland and Hager, 2006). It’s also hard to put strain in a framework with friction, for example. ‘Physically realistic’ modeling has to make a lot of assumptions and use heavy duty equipment (finite elements, for example) to incorporate strain.

There is also a separate issue with elastic rebound: It’s not very easy to tell whether all the accumulated shear strain was released in an earthquake or not. What kind of measurements would tell us this?

I strongly suspect that we are fundamentally underestimating the frequency of very short recurrence intervals on faults. They’re close to invisible to paleoseismology, which is our main source of data for recurrence interval statistics, because very

closely-spaced events may not each produce differentiable colluvial wedges or other signs of surface deformation. This could plausibly result in a strong sample bias in the statistics. Nonetheless, we have clear observations of short recurrence intervals in the past few years. For example, some parts of the Monte Vettore fault slipped about 20 cm in the Amatrice earthquake and then ~2 m in the Norcia earthquake a few months later (Gruppo di Lavoro INGV sul Terremoto di Amatrice-Visso. (2016, October 29). PRIMO RAPPORTO DI SINTESI SUL TERREMOTO DI VISSO ML 5.9 DEL 26 OTTOBRE 2016 (ITALIA CENTRALE). Zenodo. <http://doi.org/10.5281/zenodo.163818>).

From a fault mechanics perspective, some researchers (for example Mark Zoback and his students, primarily Townend, as well as myself) believe that the shear stress on a fault required to initiate failure is much greater than the stress drop during the event, i.e. shear stress does not go to zero. Failure is decently described by Mohr-Coulomb models, and at, say, 10 km depth, the confining pressure is almost 300 MPa. With a reasonable rock density (2700 kg/m^3), coefficient of static friction (say 0.5), and pore fluid pressure (say 0.3 times lithostatic pressure), the shear stress at failure is 94.5 MPa. Stress drops are generally on the order of 0.1-10 MPa (see Peter Shearer's work on Brune-type stress drop estimates, for example). So if less than 10% of shear stress is relieved in an earthquake, what are the implications for the elastic rebound hypothesis?

My take on this is that elastic rebound is a great way to describe the phenomenology of earthquakes to a non-geologist. Scientifically, it was an idea of absolute genius in 1910, but it isn't a thorough or mechanically sound framework for earthquake science a century later. Stress-based frameworks are much better suited to both conceptual and quantitative treatments.

Changes: None.

[O14] Page 3, line 24. It would be useful to briefly discuss how shape and scale affect distributions generally. Shape governs the how tailed and is dimensionless; scale determines the spread of the distribution and is dimensioned (in years for this case).

Good idea. I added brief definitions of the scale, shape and location parameters to an earlier paragraph where the terms are first written.

Changes: definitions added.

[O15] Page 4, line 6. Pareto distribution is another, simpler distribution needing only shape and scale to describe $\text{COV} > 1$

The Pareto distribution isn't used to describe earthquake recurrence, as far as I know. I believe that its only use in seismology is the tapered Pareto distribution for magnitude-frequency distributions by Yan Kagan (and perhaps others); this is not a similar-enough use to include here.

Changes: None.

[O16] Page 4, line 15. Akciz et al (2012) revised Grant and Sieh (1994) and found much more periodic behavior.

The referenced sentence simply states “However, one can find examples of studies indicating the opposite conclusions”, to reinforce the paragraph’s opening statement that “No consensus exists among earthquake scientists as to the most appropriate recurrence interval distribution.” That a study revised a previous study and found different results further reinforces this point, but I don’t think there is additional benefit to citing it.

Changes: None.

[O17] Page 4, line 28. The author should consider non-dimensionalizing the results of this study to facilitate more general use of its results. Instead of mean slip of 1m, one would refer to non-dimensional slip of 1 and multiply by average slip per event to scale the results. This is effectively what the author describes already, though without formal non-dimensionalization.

I considered formal nondimensionalization but decided against it.

In the end, I decided to present the results dimensionally because most geologists think dimensionally (myself included).

Dimensional thinking allows for different heuristics to be used when analyzing methods and results, than non-dimensional thinking. Basically, non-dimensionalizing parameters forces parameters to only be defined in terms of their relationship with each other, and to only exist within the context of this specific problem.

If I say a fault has a mean slip of 1, and a mean recurrence interval of 1, and a mean slip rate of 1, it’s hard to picture such a fault. Are those reasonable values for these parameters, individually? Can’t say. Furthermore, the slip rate isn’t even 1, it’s 1/1, because it’s a rate and the non-dimensional units are still not arithmetically compatible units—multiplication and division are possible (kind of, but not reducible) but addition and subtraction are not at all defined. (Please forgive me for not knowing enough measure theory to rigorously state this...).

Non-dimensionalization in many geological problems reduces the clarity of the analysis or solution because it strips it of context. At the same time it can facilitate the manipulation of the equations within the study, or during programming, etc. Physicists like it because it makes their work easier, I would say, and they’re all pretty used to it.

Compare this to the dimensionalized problem in the paper. Can you imagine a fault with a mean per-event slip of 1 m, a recurrence interval of 1000 years, and a slip rate of 1 mm/yr? I can, and I can place it—it’s a small but pretty active intraplate fault, or a splay in a plate boundary. It’s the kind of fault one might actually go and trench.

However, all the numbers are 1 (or 1000) so that it is easy to scale to other faults by multiplying by non-dimensional scaling factors (i.e., for a Mongolian fault with a slip rate of 3 mm/yr and a per-event slip of 10 m, you scale the slip by 10 and

the slip rate by 3). It's not accidental that I have chosen these numbers, and I have explicitly described how to do the scaling.

Changes: None.

[O18] Page 5, line 15. The statement 'appears to be related to COV' is disappointing. Given that this paper is entirely simulation, the author should be able to make a quantitative comparison of slip-rate variance to COV.

Sorry to disappoint.

A comparison between COV and the slip-rate variance at any time is actually done in the study (it is in fact the core of the study)—I vary the COV of a single distribution (the lognormal distribution) while changing nothing else, then I keep the lognormal distribution with a COV of 1 and compare that to a very different distribution (the exponential distribution) with a COV of 1. The results, which are described unambiguously in the study, show that the slip rate variance changes with COV but isn't much affected by the shape of the distribution (i.e. lognormal vs. exponential).

There is no reason to pursue this farther here. The work that I have done here has covered the (small) space of geologically reasonable distributions; there aren't huge gaps left unexplored.

There are two main families of distributions that are in consideration for earthquake recurrence: Exponential distributions (possibly with modifications such as a stretched or hyperexponential distribution) and lognormal-like distributions (lognormal, Weibull, BPT, etc., which are not distinguishable in real paleoseismic datasets). I have compared both types of distributions for a single COV and the differences are very minor. There are no alternate distribution families in any consideration within seismology that are *more* different than these two families.

This work is not meant to offer a mathematical proof, which is why I gave the soft 'appears to be related to...' statement instead of a more firm 'is demonstrated to be caused by...' or 'is proven to be a function of...'. Demonstrating some correlation or relation is enough to further the general point, which is that more earthquake recurrence variability will lead to more short-term slip rate variability (which seems self-evident in any case).

Going farther would either mean doing many more numerical experiments (which do not offer real proof and would just clog the paper) or invoking more complicated mathematical tools such as stochastic calculus, which I don't really know and am unwilling to teach myself for this paper.

Changes: None.

[O19] Page 5, line 29. This problem has been studied (Weldon et al., 2004; Sieh et al., 2008). The Sumatran subduction zone work is particularly relevant and completely overlooked here.

The work by Weldon et al. (2004) is referenced on Page 6, Line 3. Neither of these papers deal with the topic of autocorrelated recurrence intervals in any quantitative or otherwise explicit manner, and I don't read anything that I can interpret as a qualitative discussion either. One can read both papers and not get a sense of whether a short recurrence interval implies the next recurrence interval will be short or long, much less any quantification. Both papers deal with the concept of 'earthquake supercycles' or groups of earthquakes that are relatively tightly-spaced and separated from other groups by long recurrence intervals. This may share a conceptual link, but from a technical perspective this likely has more to do with periodicity than autocorrelation, and these are mathematically separate such that a periodic sequence may have positive or negative autocorrelation. Goh and Barabasi (2008, Europhysics Letters) is a useful discussion on the topic of autocorrelation vs. periodicity in regards to quantifying clustering behavior (though one may safely ignore the references to seismicity in that paper).

[Note: I've analyzed Weldon's data (from K. Scharer's refined OxCal earthquake dates) and found that earthquake recurrence interval duration at the Wrightwood and Pallett Creek sites has a negative autocorrelation, i.e. a short recurrence interval is likely to be followed by a long recurrence interval, and vice versa. This is quite unlike the autocorrelation in the Puget Lowlands of Washington State, which is a network of generally similar faults; here autocorrelation is positive, so a long recurrence interval will more likely be followed by another long recurrence interval. I don't know what this means, or whether it's all noise, but it's intriguing to me and I put this bit in the paper in hopes of catching the attention of others who may be looking for a project.]

Changes: None.

[O20] Page 6, line 1. Zero friction at rupture arrest is very unrealistic, and not a prerequisite for characteristic behavior.

It is not a necessary condition but it is a sufficient one, when coupled with fairly regular reloading and failure conditions. I personally think it's unlikely as well (see Styron and Hetland 2015 for example), but zero (or very near zero) friction is part and parcel with complete stress drop in major earthquakes, which is supported by many studies (e.g., Hardebeck, 2012; Hasegawa et al., 2011) though I am pretty suspicious of the results. Furthermore, there are a host of laboratory experiments which suggest that friction during slip decreases to very low values (e.g., Di Toro et al., 2004).

Changes: None.