https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

Natural Hazards and Earth System Sciences
Discussions

A Performance-Based Approach to Quantify Atmospheric River Flood Risk

Corinne Bowers[1], Katherine A. Serafin[2], and Jack Baker[1]

[1]Department of Civil and Environmental Engineering, Stanford University
[2]Department of Geography, University of Florida

Correspondence: Corinne Bowers (cbowers@stanford.edu)

Abstract.

Atmospheric rivers (ARs) are a class of meteorologic phenomena that cause significant precipitation and flooding on the US West Coast. This work presents a new Performance-based Atmospheric River Risk Analysis (PARRA) framework that adapts existing concepts from probabilistic risk analysis and performance-based engineering for application in the context of AR-driven fluvial flooding. The PARRA framework is a chain of physically based models that link the atmospheric forcings, hydrologic impacts, and economic consequences of AR-driven fluvial flood risk together at consistent “pinch point” variables. Organizing around these pinch points makes the framework modular, in that models between pinch points can be updated without affecting the rest of the model chain, and it produces a probabilistic result that quantifies the uncertainty in the underlying system states. The PARRA framework can produce results beyond analyses of individual scenario events and can look towards prospective assessment of events or system changes that have not been seen in the historic record.

The utility of the PARRA framework is demonstrated through a series of analyses in Sonoma County, California. Evaluation of a February 2019 case study AR event shows that the individual component models produce simulated distributions that capture the observed precipitation, streamflow, inundation, and damage. The component models are then run in sequence to generate a first-of-its-kind AR flood loss exceedance curve for Sonoma County. The prospective capabilities of the PARRA framework are presented through the evaluation of a hypothetical mitigation action. It was found elevating 150 homes, selected based on their proximity to the Russian River, was able to reduce the average annual loss by half. The loss results from the mitigated building portfolio were then compared against the original case. While expected benefits were minimal for the smallest events, the larger, more damaging ARs were expected to see loss reductions of approximately $50 million per event. These results indicate the potential of the PARRA framework for examining other changes to flood risk at the community level, including future changes to the hazard, through climate change; exposure, through development; and/or vulnerability, through flood mitigation investments.

1 Introduction

Atmospheric rivers (ARs) are long ($>$2000 km) and narrow (500-1000 km) corridors of strong horizontal water vapor transport, with water concentrated mostly in the lowest 3 km of the atmosphere (Ralph et al., 2018). ARs are the primary vector for

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

moving moisture from the tropics to the midlatitudes, responsible for up to 90% of longitudinal water transport while covering only 10% of the earth's surface (Zhu and Newell, 1998). California experiences ARs coming from a pathway called the Pineapple Express that transports moisture from Hawaii to the western coast of North America. ARs along this trajectory are crucial to the stability of California's water resources. In just a hundred hours of rain per year, ARs deposit up to half of the state's annual water supply (Lamjiri et al., 2018). However, this gift comes at a price: ARs cause well over three-quarters of all extreme precipitation events in California and over 90% of the state's record floods (Lamjiri et al., 2018), leading to almost $300 million in average annual losses (Corringham et al., 2019).

One particularly devastating AR event was the Great Flood of 1862. Central California received over ten feet of precipitation in just 43 days between December 1861 and January 1862, and cities from San Francisco to San Diego set precipitation records that still stand today. Based on this catastrophe, the US Geological Survey (USGS) created a hypothetical AR scenario named ARkStorm with a return period of approximately 500–1000 years (Porter et al., 2011). The study found that California's flood control infrastructure, mostly built to withstand 100– or 200–year return period events, was woefully underprepared for this disaster scenario, and that much needed to be done in terms of both mitigation investment and emergency planning to prevent an economic catastrophe. The ARkStorm report concluded with a call for an "end-to-end stochastic model of severe weather, physical impacts, and socioeconomic consequences" (Porter et al., 2011).

This paper presents a novel performance-based risk analysis framework that integrates meteorological, hydrological, and engineering models to assess AR-induced fluvial flood risk. The Performance-based AR Risk Analysis (PARRA) framework probabilistically models ARs from inception all the way to economic consequences, as first outlined a decade ago in the ARkStorm report. The core goal of the PARRA framework is to provide a conceptual outline to link research related to the atmospheric forcings, hydrologic impacts, and economic consequences of AR-induced flood risk. It is (a) physically based, meaning it supports precise modeling of the sequence of processes from inception to impacts; (b) modular, meaning individual component models can be modified without affecting the rest of the sequence; (c) probabilistic, meaning the full uncertainty quantification at each step is carried through the model sequence to assess confidence in the final results; and (d) prospective, meaning the framework can be used to assess "what-if" questions about events that have not yet occurred. These characteristics set it apart from previous scenario-based and statistical analyses of AR-driven losses and allow for a more comprehensive evaluation of AR-induced fluvial flood risk.

1.1 Disciplinary Context

Prediction of flood damage and loss due to ARs requires disciplinary expertise spanning meteorology, hydrology, engineering risk analysis, and more. Most previous research modeling ARs has focused on two pathways: the first linking atmospheric forcings to hydrologic impacts, and the second linking hydrologic impacts to economic consequences.

The first physical process to consider is the transformation of atmospheric phenomena into precipitation and runoff. Considerable effort has been invested into understanding the climatology of ARs, through collection of meteorological field data (Lavers et al., 2020), improvements to existing numerical weather prediction models (Nardi et al., 2018; Martin et al., 2018), and intercomparison between AR detection algorithms (Shields et al., 2018). Researchers have described the particular climatology of ARs affecting California and the US West Coast (Waliser and Guan, 2017; Guirguis et al., 2018). Characterizing these features and their spatial and temporal distributions allows us to better connect AR events with their hydrological impacts, namely extreme precipitation (Chen et al., 2018; Huang et al., 2020) and runoff (Konrad and Dettinger, 2017; Albano et al., 2020).

Translating from hydrologic impacts to economic consequences is generally a multi-step process, and in the literature there are both specific models for portions of the process and multi-model sequences. Hydrologic routing software such as HEC-HMS (USACE, 2020) determine the shape of the event hydrograph at a given point along the river based on a precipitation event, and hydrodynamic solvers such as HEC-RAS (Brunner, 2020) or LISFLOOD-FP (Bates and De Roo, 2000) generate 2D maps of the resulting inundation. These tools are generally designed for practitioners to assess the consequences of the "100 year event," which is the hydrological event with a 1% annual probability of occurrence. The 100 year event is a term with a long history in planning and engineering design and is generally set as the standard that flood control infrastructure must be built to withstand. In the US, the Federal Emergency Management Agency (FEMA) has been defining flood risk in these terms for decades, through the Federal Insurance Rate Maps (FIRMs) that delineate the 100 year floodplain and through their open-source loss estimation software Hazus-MH (FEMA, 2020).

1.2 Methodological Frameworks

Some studies have gone beyond domain-specific solutions to capture the entire chain of processes from atmospheric events to economic loss (Dominguez et al., 2018; Felder et al., 2018; Porter et al., 2011). However, these previous works have focused on the consequences of one specific disaster scenario rather than the full spectrum of AR flood risk in a given location. Corringham et al. (2019) subsequently connected ARs directly to their economic impacts and introduced the first estimate of total flood loss attributable to ARs in the Western US. While this was an important step in understanding the consequences of these storms, it included no physical modeling of the processes connecting ARs to losses. Capturing these physical connections with disciplinary models of atmospheric forcings, hydrologic impacts, and economic consequences allows for a greater depth of understanding of the underlying phenomenological drivers.

The framework proposed here provides a structure to overcome these challenges. The PARRA framework is an organized sequence of component models that connect AR occurrence to the damage and losses that result from AR-driven flooding. The idea of "model chains" or "model sequences" is born from probabilistic risk assessment, a field that originated in the regulation of nuclear reactors (NRC, 1975). Probabilistic risk assessment frames risk as a function of likelihood times consequence. The overall risk at a site is the product of the likelihood and consequence of a given outcome, integrated over all possible outcomes defined in the event space (Baker et al., 2021).

The report that first introduced probabilistic risk assessment relied on logic trees to organize these possible outcomes (NRC, 1975). Since then, more complex models and combinations of models have been adopted to improve the representation of the event space. The use of multi-model sequences is now common in natural disaster risk assessment and has grown in popularity as a way to address flood risk (Apel et al., 2004; Felder et al., 2018; Uhe et al., 2020). The insurance industry in particular has embraced model chains to create catastrophe risk models for an array of natural hazards (Pinelli and Barbato, 2019).

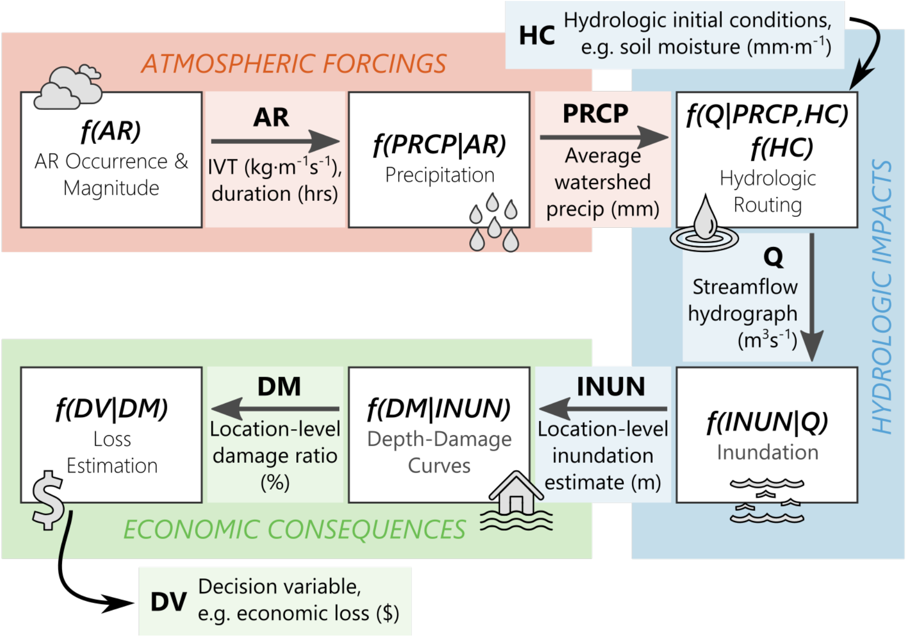

Figure 1. PARRA framework flowchart. White outlined boxes indicate component models, shaded arrows indicate pinch point variables, and colors broadly represent existing research domains.

While the insurance models prove the viability and utility of the proposed concept, they are often proprietary and so are not useful for generating public data or performing research. Therefore the PARRA framework is defined using the language of performance-based engineering, which focuses on the variables organizing the model sequence rather than the component models themselves.

Performance-based engineering (PBE) refers to designing a system to meet a target performance objective with a specified reliability rather than satisfying prescriptive requirements. Rather than defining the models that must be linked together, PBE frameworks are designed around "pinch points," or points in the modeling chain where only a handful of variables need to be passed from one step to the next (Garrick, 1984). The first PBE framework was introduced to estimate losses to buildings in future earthquakes (Krawinkler, 1999). The success in earthquake engineering led to development of PBE frameworks for other hazards, including wind (Ciampoli et al., 2011), hurricanes (Barbato et al., 2013), fire (Guo et al., 2013), and corrosion (Flint et al., 2014). Condensing the models down to a reduced set of pinch points at designated steps within the model chain dramatically improves the flexibility of the framework and helps organize research efforts into interchangeable and modular components.

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

2 Framework Description

The PARRA framework is composed of six component models connected by pinch points. We first develop the overall structure of the framework, then introduce the specific pinch points and component models.

At a fundamental level every PBE framework is an implementation of the total probability theorem, which states that $P(A) = \sum_{i=1}^{n} P(A\,|\,B_i)\,P(B_i)$. In this equation, $P(A\,|\,B_i)$ represents the conditional probability of event A given that event B_i has occurred, and $P(B_i)$ represents the probability of event B_i out of some set of mutually exclusive, collectively exhaustive events. Summing these conditional probabilities over all possible instances of B_i gives us the total probability of event A.

Equation 1 modifies the statement of the total probability theorem to better fit the context of natural hazard assessment. We first transform the summation into an integral to capture continuous state variables rather than discrete events. Therefore, instead of finding the probability P of one event, we now move to calculating λ, the rate of occurrence of some specified condition. $P(DV > x \mid AR)$ represents the probability of the decision variable DV exceeding some loss value x, conditioned on the inducing event condition AR. $\lambda(AR)$ represents the rate density of that inducing event condition, which is then integrated over all possible events in the sample space. We refer to the resulting value, $\lambda(DV > x)$, as the loss exceedance curve, which is developed further in Sect. 4.2. Overall this equation forms the theoretical basis for the construction of a performance-based probabilistic framework for ARs.

$$\lambda(DV > x) = \int_{AR} P(DV > x\,|\,AR)\;\lambda(AR)\;dAR \tag{1}$$

While Eq. 1 is conceptually straightforward, calculating $P(DV > x\,|\,AR)$ is difficult. Therefore we expand this probability statement by introducing intermediary pinch point variables that decompose the calculations into a series of component models, as shown in Eq. 2.

$$\begin{aligned}\lambda(DV > x) = \iiint\iiint\; & P(DV > x\,|\,DM) * \\ & f(DM|INUN) * \\ & f(INUN|Q) * \\ & f(Q|PRCP, HC) * f(HC) * \\ & f(PRCP|AR) * \lambda(AR) \\ & dDM\; dINUN\; dQ\; dHC\; dPRCP\; dAR\end{aligned} \tag{2}$$

Each component model represents one discrete physical process in the series of processes connecting AR events to flood losses. The component models of the form $f(Y|X)$ are conditional probability density functions that describe the distribution of results from numerical analyses. The component model $P(DV > x\,|\,DM)$ is the complement of the cumulative distribution function, starting at 100% probability of occurrence and moving to a zero as x increases.

Conceptually the framework is executed by moving from right to left through Eq. 2, as each component model is conditioned on the one(s) to the right. This equation is additionally represented visually in Fig. 1.

2.1 Pinch Point Variables

The component models in Eq. 2 and Fig. 1 are linked together by pinch points that transfer information from one stage to the next. The pinch point variables are: AR, atmospheric river (some measure of intensity for a specific AR event); $PRCP$, precipitation (accumulated rainfall at the location or watershed of interest); HC, antecedent hydrologic conditions (some measure of the preexisting water balance within the watershed); Q, streamflow (the inflow hydrograph for the study area of interest); $INUN$, inundation (the height of water at buildings or locations of interest within the study area); DM, damage measure (damage ratios at buildings or locations of interest within the study area); and DV, decision variable (some metric of impact or consequence for the study area). Pinch point variables are represented by the arrows in Fig. 1.

AR magnitude can be categorized using the intensity scale from Ralph et al. (2019), which ranks ARs on a scale from 1 to 5 to summarize their severity (from weak to exceptional) and potential impacts (from beneficial to hazardous). Category 1 ARs are classified as primarily beneficial storms, replenishing the water supply without causing adverse consequences. Category 5 ARs are classified as primarily hazardous with a high likelihood of flooding and damage. The AR category is determined based on two quantities: the maximum recorded integrated water vapor transport (IVT) $(\mathrm{kg}\cdot\mathrm{m}^{-1}\mathrm{s}^{-1})$ and the duration (h) of sustained IVT exceeding 250 $\mathrm{kg}\cdot\mathrm{m}^{-1}\mathrm{s}^{-1}$. Therefore we summarize each AR event in terms of its maximum IVT and duration and record the pinch point AR as a vector with these two elements.

Precipitation ($PRCP$) measures the amount of water released by the storm as measured at the ground surface, generally recorded as a height or depth. Precipitation is variable in both space and time and can potentially be summarized in several ways. In this work we define the precipitation pinch point variable as a scalar value representing the storm-total accumulated rainfall averaged across the spatial extent of the watershed.

Hydrologic conditions (HC) refer to the saturation state before the storm, i.e. how much water is already present in the watershed system. Previous research has found antecedent soil moisture to be an important factor in predicting which precipitation events will lead to flood events (Cao et al., 2019). Therefore the pinch point variable representing antecedent hydrologic conditions is a scalar value measuring the watershed-average soil moisture equivalent height.

The streamflow hydrograph (Q) is the timeseries of flow measurements vs. time recorded at the study area inlet for the duration of the AR event. Instead of storing this entire vector of flow vs. time values, we parameterize the hydrograph with three variables: Q_p, peak streamflow $(\mathrm{m}^3\mathrm{s}^{-1})$; t_p, time to peak streamflow (h); and m, a unitless shape parameter. This parameterization process is discussed in more detail in Sect. 3.4.

Inundation ($INUN$) is the surface water depth at locations of interest within the study area. These values should be zero or positive-valued only, as they are represent a height of water above the ground surface. The vector of heights at N locations of interest is stored for the next component model.

The damage measure (DM) is defined in this work as a damage ratio, or the expected cost to repair a building divided by the total value of that building. The damage ratio is assumed to be purely a function of water depth with respect to the first floor elevation. The result of the depth-damage calculation is a length N vector of the same size as $INUN$ where 0 signifies no damage and 1 signifies a cost of repair equal to the value of the building.

Finally, the decision variable (DV) is some actionable measure of AR impacts, usually a number representing expected monetary losses. The expected loss for each structure is the value of exposed assets, namely building and contents valuations, multiplied by the damage ratio. The result is a length N vector that represents the expected loss for each location of interest.

2.2 Component Models

The pinch point variables described above are used to decompose AR-induced fluvial flooding into six discrete, self-contained physical processes, which are referred to as component models. The six component models are: AR occurrence/magnitude, precipitation, hydrologic routing, inundation, depth-damage curves, and loss estimation. These are represented by the white boxes in Fig. 1.

In practice, Eq. 2 and other performance-based frameworks typically cannot be reduced to a closed form equation and are therefore solved through Monte Carlo simulation. Each one of the component models is implemented as a numerical analysis that produces a best-fit projection for the next pinch point variable in the model sequence plus some representation of the uncertainty. Rather than explicitly parameterizing the uncertainty at every step, we define empirical relationships based on the overall distribution of the historic record, then generate Monte Carlo samples to produce multiple stochastic realizations of each pinch point variable. This stochastic realizations of potential system states are propagated through the model chain to build an empirical distribution of expected loss. A more in-depth discussion of framework implementation and Monte Carlo simulation can be found in Sect. 5.2.

3 Case Study: Sonoma County

We describe a case study application of the PARRA framework for the lower Russian River in Sonoma County, California, as shown in Fig. 2. The case study is used to discuss implementation choices for each one of the component models introduced in the previous section and to illustrate the potential value and insights that can be provided by this framework.

Flood losses in Sonoma County have totaled over five billion dollars in the last forty years, with over 99% of that due to ARs (Corringham et al., 2019). Sonoma County also has the highest proportion of state disaster assistance payouts in California (34%), more than six times the second highest county (Sonoma County, 2017). The spatially repetitive, locally severe flooding seen in Sonoma County is a signature characteristic of ARs.

Because of its location and its extensive history of AR-induced flooding (Ralph et al., 2006), Sonoma County is an excellent place to perform a case study implementation of the PARRA framework. The vast majority of AR-induced flooding is the result of fluvial flooding from the lower Russian River, which has overtopped its banks 36 times in the last 80 years (Rogers, 2019). The Russian River flows 110 miles through Mendocino and Sonoma counties, draining a watershed that covers 1,485 square miles (3,846 km^2). The main tributaries within Sonoma County that flow into the Russian River are Dry Creek, which joins just south of Healdsburg; Mark West Creek, which joins around Forestville; and Austin Creek, which joins between Monte Rio

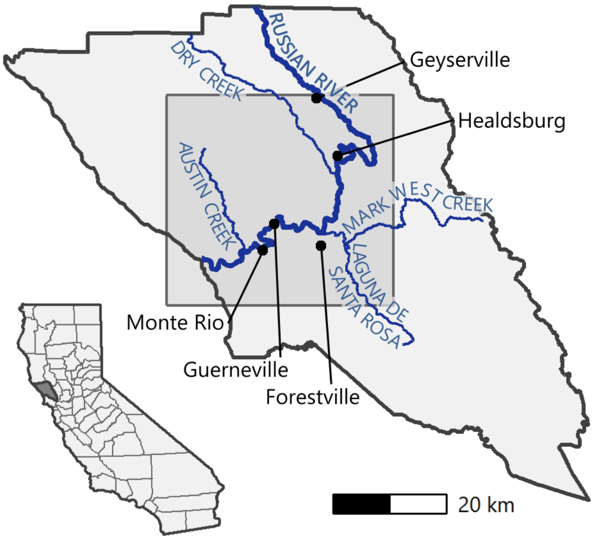

Figure 2. Map of Sonoma County. The shaded box indicates the study area, and the labels highlight notable creeks and riverside communities. A map of California with Sonoma County shaded in dark grey is included for geographic context.

and the Pacific Ocean. The Laguna de Santa Rosa is a protected wetlands complex that serves as an important overflow area for flood control.

Within Sonoma County we use the PARRA framework to examine the impacts of an AR storm from February of 2019 (the "2019 event"). This severe storm caused approximately \$155 million dollars in damage with \$91.6 million attributed to residential buildings (Chavez, 2019). Several towns along the lower Russian River were flooded and required emergency evacuations for residents. The combination of this storm's recency and its severe impact on riverside communities mean that datasets are available to validate many of the individual component models.

We treat the 2019 event as a case study, comparing the simulated Monte Carlo realizations of each pinch point variable to observed data for every component model within the framework sequence. Each subsection discusses how the component model was implemented and calibrated for Sonoma County, then presents a comparison between the simulated realizations and the observed values from the 2019 event.

3.1 Atmospheric River (AR)

3.1.1 Component Model

The first step of the case study was to create a catalog of historic ARs affecting the study area. We started with IVT records from MERRA-2 (Gelaro et al., 2017), which records IVT at six hour intervals on a grid of $50\text{km} \times 50\text{km}$ cells. We then used the detection algorithm from Rutz et al. (2014) to identify landfalling ARs in each grid cell and recorded the maximum IVT

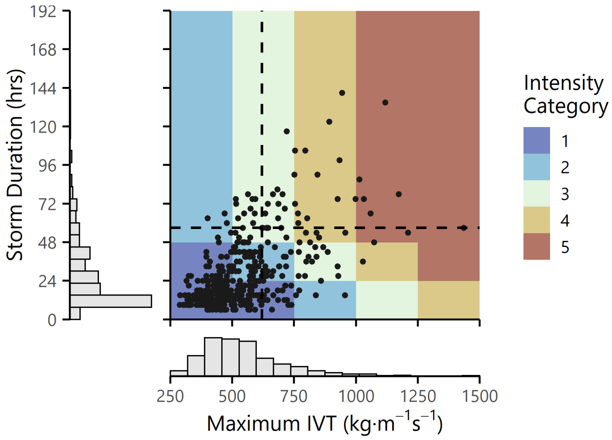

Figure 3. Summary of Sonoma County historic catalog. Each point represents one AR event recorded in Sonoma County between 1988 and 2019 (n = 382). The maximum IVT and duration for the 2019 event are represented by the two dashed lines, and the Ralph et al. (2019) intensity categories are represented by the background colors.

Table 1. Statistics of the Sonoma County historic catalog by AR intensity category.

Category	**Events**	**Mean Values**				
		Maximum IVT $(\mathrm{kg \cdot m^{-1}s^{-1}})$	Duration (h)	Precipitation (mm)	Peak Flow $(\mathrm{m^3s^{-1}})$	Time to Peak Flow (h)
1	238	464	16	19	85	19
2	67	597	35	51	199	22
3	49	665	56	81	214	21
4	20	849	75	156	513	23
5	8	1,139	75	179	471	24

and duration for each. Because AR activity occurs almost exclusively during the first half of the water year (October 1st –April 1st), we kept only ARs that occurred within those annual ranges.

A rectangular study area was defined as shown in Fig. 2 to encompass the lower Russian River, starting at USGS gage 11463500 near Geyserville and ending at the outlet to the Pacific Ocean. We followed the process outlined by Albano et al. (2020) to take ARs identified on the 50km × 50km MERRA-2 grid and downscale them to the area of interest to create a catalog of historic AR events. The generated catalog contains 382 AR events recorded in the Russian River watershed over a 32 year period spanning water years 1988–2019. The maximum IVT and duration of these events are summarized in Fig. 3. The historic catalog is assumed to accurately represent the climatology of the region, although future work could expand upon this catalog to include the characteristics of storms not yet seen or recorded in Sonoma County.

We also collected additional information about each AR, including precipitation, streamflow, and more. Precipitation values are the storm-total cumulative precipitation areally averaged over the inlet watershed, and peak flow and time to peak flow were both calculated based on data from USGS gage 11463500 (study area inlet). Table 1 summarizes the statistics of some of these additional parameters as a function of the Ralph et al. (2019) AR intensity categories.

3.1.2 2019 Case Study Event

To better understand the case study storm, we did not simulate any new values, but instead used the real maximum IVT (620 $\mathrm{kg \cdot m^{-1} s^{-1}}$) and duration (57 h) recorded for the 2019 event as inputs to the next component model.

3.2 Precipitation (PRCP)

3.2.1 Component Model

We used the historic catalog generated in the previous subsection to estimate a statistical relationship between the predictors (maximum IVT and duration) and the outcome (storm-total precipitation averaged over the watershed of interest). The goal in quantifying this relationship was to be able to generate simulated precipitation realizations that are consistent with the historic climatology of the region.

We implemented a weighted least squares (WLS) linear regression to predict precipitation as a function of maximum IVT, duration, and an interaction term between the two. We chose WLS over ordinary least squares methods because of the heteroskedasticity in the residuals, which can be identified by visual inspection of Fig. 4(a–b). Even after applying weights to correct for the heteroskedasticity, though, there were still some extreme precipitation events in the historic catalog that were not well represented by a Gaussian error model. We therefore characterized the standard errors with a mixture model: 90% of residuals were calculated using the WLS standard errors, and 10% were calculated with a distribution fit the to the largest 10% of AR events. Regression coefficients are reported in Eq. 3, and Fig. 4(a–b) shows the regression line plotted over the historic catalog data at selected values of maximum IVT and duration.

$$\mathrm{E}[PRCP_i] = -3.516 + 0.0212\,(IVT_i) - 0.5044\,(DUR_i) + 2.737 \times 10^{-3}\,(IVT_i{*}DUR_i) \tag{3}$$

where $\mathrm{E}[PRCP_i]$ is the expected total precipitation for event i, and IVT_i (maximum IVT) and DUR_i (duration) are the two elements of the pinch point variable AR_i.

Because our goal was to match the overall distribution of historical precipitation, we calculated goodness of fit metrics that focus on success in replicating distribution shape rather than individual events. Fig. 4(c) shows a quantile–quantile (Q–Q) plot comparing a prediction from the fitted regression including errors against the observed distribution of precipitation from the historic catalog. Visually the simulated results fall very close to the parity line on this plot, which indicates good distributional fit. A more quantitative metric of distributional fit is the two sample Kolmogorov–Smirnov (K–S) test, which is a nonparametric test designed to measure the goodness of fit between two empirical distributions. We calculated the test statistic d_{PRCP} such that a value of $d_{PRCP} > 1$ would reject with 95% confidence the null hypothesis that the two datasets are drawn from

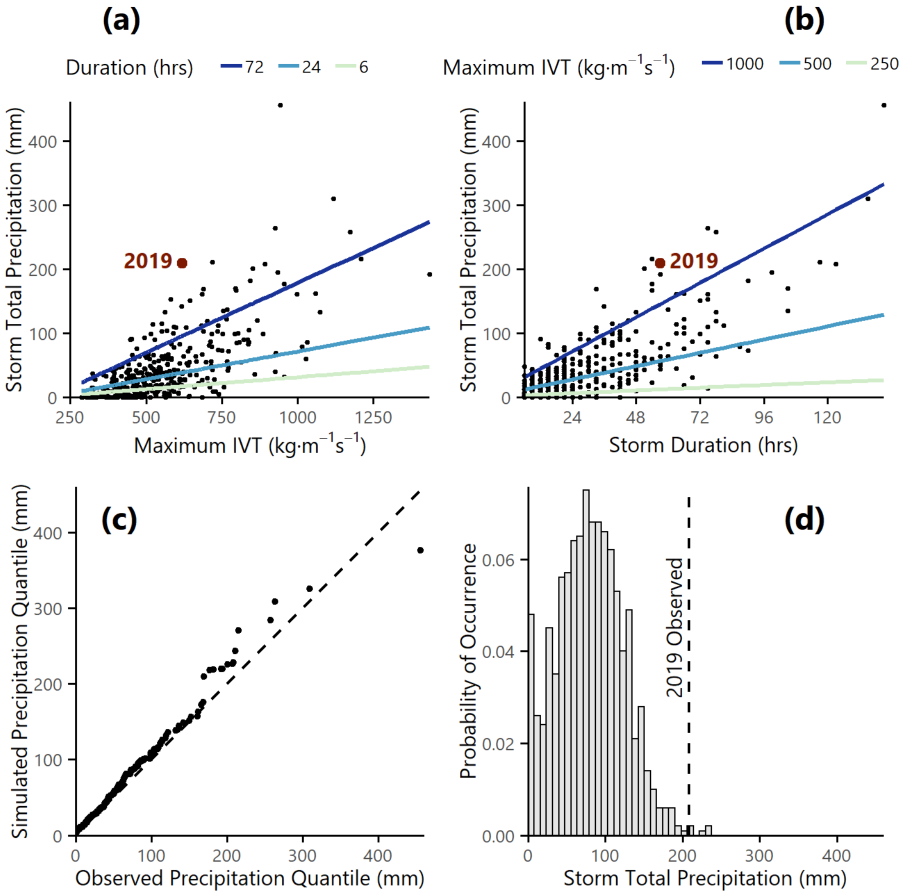

Figure 4. Precipitation component model. Scatterplots of **(a)** maximum IVT vs. precipitation and **(b)** duration vs. precipitation, with the fitted WLS regression line shown at the values specified in the legends at the top. **(c)** Q–Q plot of observed vs. simulated precipitation including uncertainty for all AR events in the historic catalog. **(d)** Distribution of simulated precipitation realizations including uncertainty for the 2019 event, with the 2019 observed precipitation marked with a dashed vertical line.

the same underlying probability distribution. Comparing the observed vs. simulated distribution yielded a test statistic of $\mathrm{E}[d_{PRCP}] = 0.478$, which is well under the rejection threshold for the null hypothesis.

3.2.2 2019 Case Study Event

255 We next considered whether the fitted precipitation component model adequately captured the observed precipitation from the 2019 event. We generated 1,000 Monte Carlo realizations from the precipitation component model using the 2019 observed IVT and duration values and plotted the results as a histogram in Fig. 4(d). The histogram represents realizations of potential precipitation if another AR occurred in Sonoma County with the same characteristics. The 2019 observed storm-total precipitation (208 mm) is marked in Fig. 4(d) by the vertical dashed line.

The 2019 observed precipitation falls in the upper tail of the simulated precipitation distribution. Previous research has suggested that AR orientation and the resultant orographic enhancement from mountains along California's North Coast can play a significant role in overall precipitation accumulation (Hecht and Cordeira, 2017; Ralph et al., 2013), which is not explicitly accounted for in the current implementation. Another explanation may be selection bias—the 2019 event was chosen as a case study scenario because of its severe consequences, so it is reasonable to see precipitation on the higher end of the expected range relative to the size and strength of the driving AR. Because the observed precipitation was captured within the simulated distribution, though, we accept this implementation as a suitable representation of reality.

It is also interesting to note that about 5% of realizations did not cause any precipitation. This finding is consistent with the rest of the historic catalog, which has zero observed precipitation for 4.5% of AR events over the last 32 years. These represent instances where atmospheric conditions in the study area met AR threshold criteria but no precipitation was recorded. One possible explanation may be that the AR passed overhead without releasing moisture, or that precipitation occurred outside of the spatial/temporal bounds of the catalog. We note that Sonoma County is not guaranteed to see any impacts, beneficial or hazardous, from the combination of AR characteristics recorded during the 2019 event.

3.3 Hydrologic Conditions (HC)

3.3.1 Component Model

To characterize antecedent hydrologic conditions we used the CPC Soil Moisture dataset as provided by the NOAA/ OAR/ESRL Physical Sciences Laboratory (van den Dool et al., 2003). This simulated dataset reports the monthly average soil moisture as an equivalent height of water (mm) found in the top meter of the subsurface. While this does not necessarily represent the true soil moisture in Sonoma County, the dataset has global coverage at a 0.5° latitude × 0.5° longitude resolution and reports monthly values back to 1948, which covers the full spatial and temporal extent of the historic catalog.

Soil moisture values were matched with AR events in the historic catalog, and based on these records the antecedent hydrologic conditions component model was characterized as a lognormal distribution with $\mu_{\mathrm{ln}} = 4.18$ and $\sigma_{\mathrm{ln}} = 0.43$. The values in the historic catalog and the fitted lognormal distribution are both shown in Fig. 6(a). When simulating future events we sample soil moisture from this distribution, under the assumption that soil moisture records from past ARs implicitly capture the effects of seasonality on both AR occurrence and soil moisture. It is worth restating that the historic catalog generated for Sonoma County only considers ARs during the six month wet season, so by construction only soil moisture values from October to April were included in this fitted distribution. The distribution of soil moisture values would shift downward if values from the dry season were included.

3.3.2 2019 Case Study Event

We used the "observed" soil moisture for the 2019 event ($140\ \mathrm{mm}$) as input to this case study. This was a 96^{th} percentile soil moisture among all ARs in the historic catalog, indicating that the subsurface was already quite saturated when this AR

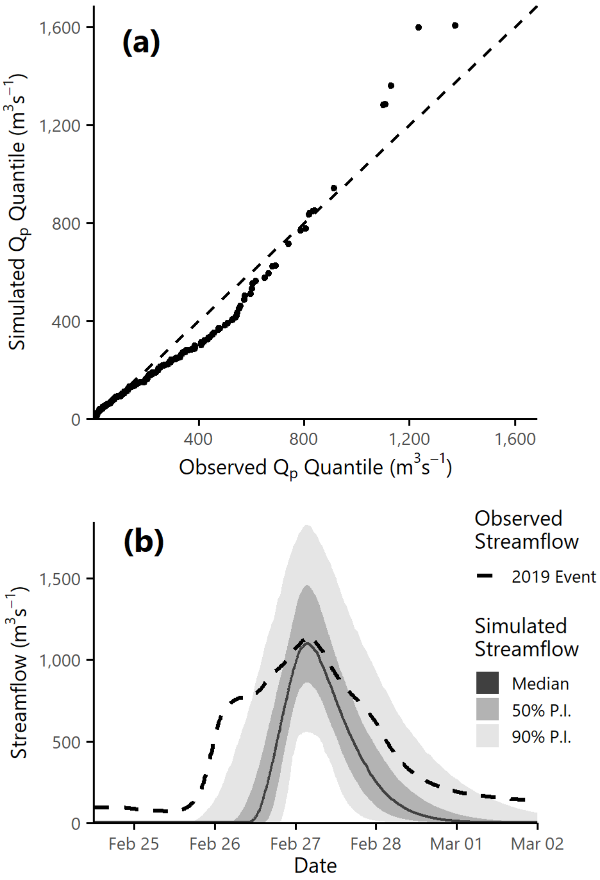

Figure 5. Streamflow component model. **(a)** Q–Q plot of observed vs. simulated peak streamflow (Q_p) values for all events in the historic catalog. **(b)** Distribution of simulated streamflow hydrograph realizations for the 2019 event. The observed hydrograph timeseries is shown as a black dashed line. The solid line represents the median of the simulated realizations, and the dark and light grey shaded areas represent the 50^{th} and 90^{th} percentile prediction intervals, respectively.

made landfall. This "observed" soil moisture value was used along with the 2019 observed precipitation as input to the next component model.

3.4 Flow (Q)

3.4.1 Component Model

295 As mentioned in Sect. 2.1, the streamflow hydrograph Q is characterized in terms of three parameters: the peak streamflow (Q_p), the time to peak streamflow (t_p), and the hydrograph shape parameter (m). Equation 4 from the National Engineering

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

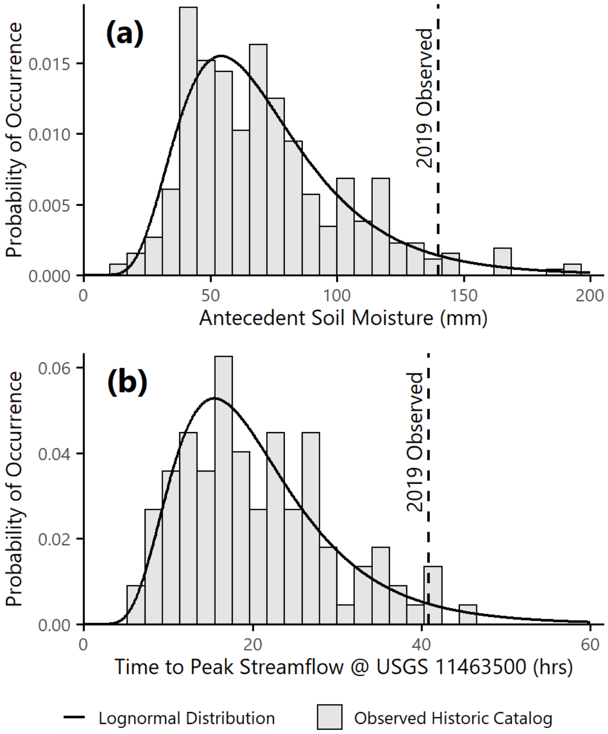

Figure 6. Fitted lognormal distributions. (a) HC, antecedent soil moisture; **(b)** t_p, time to peak streamflow at USGS gage 11463500 (study area inlet). In each panel, the histogram represents the values in the historic catalog, the solid line represents the lognormal fit, and the dashed vertical line represents the 2019 observed value.

Handbook (NRCS, 2004, Chapter 16) converts these three parameters into the full streamflow hydrograph for further analysis.

$$\frac{Q}{Q_p} = \left(\frac{t}{t_p}\right)^m \exp\left[m\left(1-\frac{t}{t_p}\right)\right] \quad (4)$$

where Q is the instantaneous streamflow timeseries recorded at USGS gage 11463500 (the study area inlet) at a vector of time values t, and Q_p, t_p, and m are constant parameters defined for each AR event.

We step through the fit and calibration of each of these parameters, starting with Q_p. Rather than using a complex hydrologic routing model, we implemented a simplified method to estimate Q_p by calculating runoff as an intermediary variable. Runoff is the portion of precipitation that does not go towards evapotranspiration or infiltration but instead flows over the ground surface. To calculate runoff as a function of precipitation we used the curve number (CN) method (NRCS, 2004, Chapter 10), which is an empirical method that computes runoff as a function of precipitation and a unitless regional "curve number" parameter that captures catchment characteristics. This is one of the most widely used solutions for rainfall-runoff modeling. The CN method also includes the effects of antecedent hydrologic conditions by defining an upper (wet) CN bound and a lower (dry)

CN bound. The curve number was found by following the method described in the National Engineering Handbook (NRCS, 2004, Chapter 10), and uncertainty bounds on the CN estimate were calculated with a bootstrapping technique.

We fit an ordinary least squares linear regression to predict Q_p as a function of precipitation, runoff, and an interaction term between the two. Similar to Sect. 3.2, we once again used a mixture model to capture the long tails of the residual distribution, with 90% of errors calculated based on the bulk of the historic catalog values and 10% calculated based on the extremes. The regression form is shown in Eq. 5, and the Q–Q plot for this regression fit is shown in Fig. 5(a). We again validated regression fit by comparing the shape of the observed vs. simulated distribution rather than comparing individual records. The K–S statistic for the OLS regression fit was calculated to be $\mathrm{E}[d_Q] = 0.889 > 1$, so we conclude that the regression produced an acceptable fit to the data.

$$\mathrm{E}\left[Q_{p,i}\right] = 8.986 + 0.3626\left(PRCP_i\right) + 11.56\left(R_i\right) - 0.0162\left(PRCP_i * R_i\right) \tag{5}$$

where $\mathrm{E}[Q_{p,i}]$ is the expected peak streamflow at USGS gauge 11463500 during event i, R_i is the watershed-average runoff, and $PRCP_i$ is the watershed-average total precipitation.

The time to peak streamflow t_p was calculated based on the distribution of observed values in the historic catalog, which was found to be well represented by a lognormal distribution with $\mu_{\ln} = 2.95$ and $\sigma_{\ln} = 0.44$. This distribution is shown in Fig. 6(b).

For the hydrograph shape parameter we chose $m = 4.0$, which was recommended by the National Engineering Handbook (NRCS, 2004, Chapter 16) and was found to be a reasonable approximation for this section of the Russian River through comparison to observed streamflow records.

3.4.2 2019 Case Study Event

Given the 2019 observed precipitation and antecedent soil moisture, we compared the predicted streamflow hydrograph from the calibrated component model to the observed hydrograph from the 2019 event. The hydrograph at USGS gage 11463500 (study area inlet) recorded a peak streamflow of $Q_p = 1{,}130\ \mathrm{m^3 s^{-1}}$ and a time to peak streamflow of $t_p = 41\ \mathrm{hrs}$. We generated 1,000 Monte Carlo realizations from the streamflow component model to estimate simulated distributions of Q_p and t_p given the observed precipitation and soil moisture values as input. Figure 5(b) indicates that the fitted component model captures the observed values within the range of uncertainty.

3.5 Inundation (INUN)

3.5.1 Component Model

The inundation component model accepts a hydrograph at the inlet to the study domain, routes water through the river channel, and distributes it through the floodplain based on site-specific information defined by the user. We used the hydrodynamic model LISFLOOD-FP (Bates and De Roo, 2000), referred to as LISFLOOD, to perform these calculations as it is lightweight and computationally efficient yet capable of capturing hydrologic processes over complex terrain. The model grid was defined

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

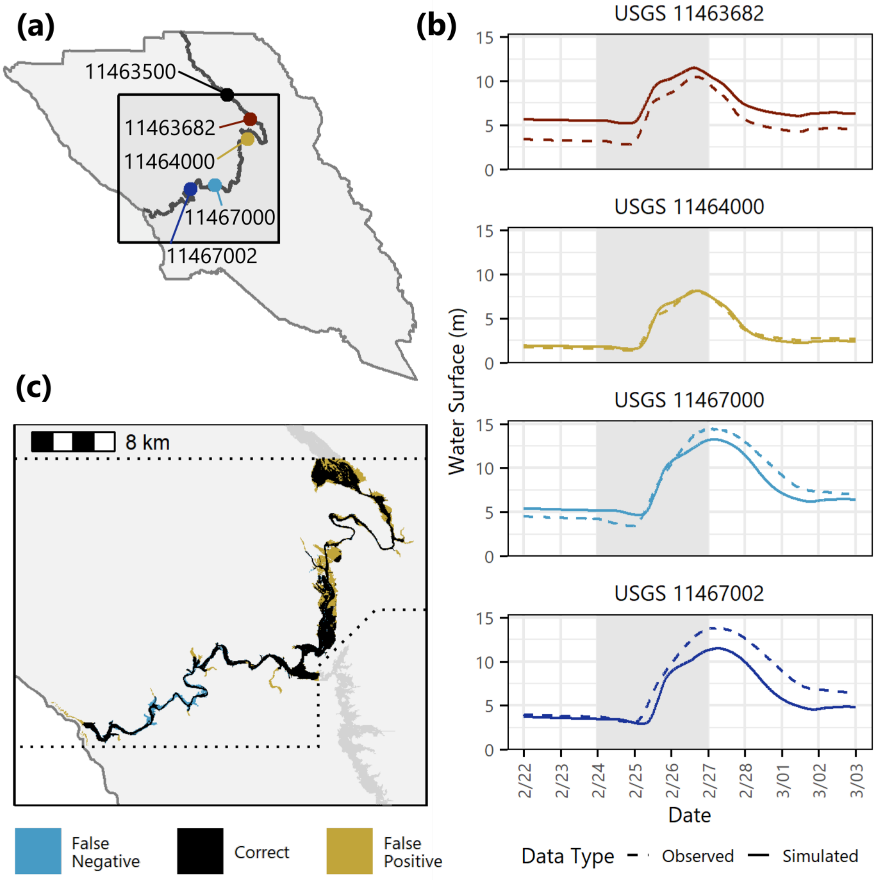

Figure 7. Inundation component model. **(a)** Locations of USGS gages with available stage data for the 2019 event. **(b)** Comparison of observed vs. simulated streamflow at selected USGS gage locations for the 2019 event. The Y axis represents the water surface height above the datum of the gage, as recorded in the USGS National Water Information System. The grey shaded areas indicate the days with sustained AR conditions. **(c)** Comparison of "observed" (Sonoma GIS) vs. simulated inundation extent for the 2019 event. The dotted lines denote the extent of the Sonoma GIS map, which is slightly smaller than the PARRA study area and does not include the Laguna de Santa Rosa.

with 925,000 cells at $40\mathrm{m} \times 40\mathrm{m}$ resolution. LISFLOOD has several parameters, including floodplain roughness, channel roughness, channel shape, and hydrologic boundary conditions, that must be specified by the user to best represent the study area. We chose twenty parameters of interest within the LISFLOOD environment, generated 500 Latin hypercube sample sets of these twenty parameters, and calculated inundation maps for each sample set. We then conducted sensitivity testing to determine best-fit parameters for the subset of "sensitive" parameters that were determined to significantly impact inundation results. Best-fit values were chosen such that when the 100 year peak flow value from USGS StreamStats was given as input, the LISFLOOD inundation extent matched the 100 year floodplain from the FEMA National Flood Hazard Layer (NFHL) along the Russian River. We adopted the accuracy metrics from Wing et al. (2017) used for validation of their nationwide

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

flood hazard map and referred to their values as benchmarks of acceptable performance. Overall the fitted LISFLOOD model was able to reach a critical success index of 69%, which means that when either the FEMA NFHL or the LISFLOOD model predicted inundation, the prediction of the LISFLOOD model was correct over two thirds of the time. More information about this fitting process can be found in the supplemental code release.

Despite its efficiency, the runtime of LISFLOOD was found to be prohibitive for Monte Carlo analysis, which involves repeated iterations of the same component model calculations. Therefore we calibrated and used a low dimensional surrogate model to quickly and accurately reproduce the results of the hydrodynamic simulation. Surrogate models are popular for testing multiple scenarios, exploring uncertainty, and making near-real-time predictions without the time and computational expense of high-fidelity model runs (Bass and Bedient, 2018; Razavi et al., 2012). 1,000 LISFLOOD runs were conducted over a range of Q_p and t_p values to populate the parameter space. We used the inverse distance weighting spatial interpolation method to generate new inundation maps based on this existing database of LISFLOOD runs. Based on a specified Q_p and t_p, the spatial interpolator identifies the "closest" points in parameter space and weights them based on distance to produce a best-fit estimate of the LISFLOOD inundation map. The hyperparameters of the surrogate model, which control the size of the search neighborhood, the distance weighting power function, and the relative importance of Q_p vs. t_p, were fit by tenfold cross validation. The error metric was the root mean squared error (RMSE) of all LISFLOOD grid cells. Replacing the LISFLOOD hydrodynamic simulation with this new surrogate model significantly reduced the computational demand of the PARRA framework while maintaining high levels of accuracy. The final fitted surrogate model reduced the runtime of a single inundation calculation from hours to seconds and had a median RMSE of 3.5 cm, which is a tolerable tradeoff when compared to the 3 cm median relative vertical accuracy reported by the digital elevation model.

The gridded inundation maps from the surrogate model were overlaid with building information to estimate inundation heights at locations of interest, namely residential buildings in Sonoma County. We used building footprints from SonomaVegMap LIDAR data and building parcel information from the Sonoma County Clerk Recorder Assessor to identify residential locations. Reporting inundation heights at only these locations rather than everywhere in the study area reduced the amount of information that needed to be transferred from one component model to the next.

3.5.2 2019 Case Study Event

Given the 2019 observed hydrograph at USGS gage 11463500 (study area inlet), we used LISFLOOD to generate a simulated inundation map for the 2019 event. We examine three different strategies to validate the accuracy of the component model result. The first focuses on the height of the river within the channel. There were four USGS gages downstream of the study area inlet that recorded stage heights at 15 minute intervals during the 2019 event, as shown in Fig. 7(a). Each of the timeseries plots in Fig. 7(b) compares the observed vs. simulated stage height in the channel at one of these four locations.

Outside of the channel there exists little recorded information about floodplain inundation resulting from the 2019 event. However, the Geographical Information Systems (GIS) branch of the Permit and Resource Management Department in Sonoma County has made available some of the results from their internal inundation modeling of the lower Russian River. These maps, rather than representing specific scenario flood events, are indexed to stage heights at the Guerneville bridge (USGS gage

11467002, Fig. 7(a)). Together they constitute a suite of design events that model stepwise increases in inundation severity across the Russian River watershed.

We chose the inundation map from the design event that most closely matched the peak stage height observed at USGS gage 11467002 during the 2019 event. We refer to this as the Sonoma GIS map. While we cannot expect a perfect match, as there are specifics of hydrograph shape and site conditions from the 2019 event that are not captured in an idealized design event, we assume that the Sonoma GIS map serves as a reasonable representation of the true inundation observed during this AR, and if our LISFLOOD model can approximate this map then it is considered a successful implementation. Our LISFLOOD model was able to reproduce the Sonoma GIS map with a critical success index of 67.77%, which is a very good match to the "observed" case. Figure 7(c) shows the spatial extent of flooding from both sources.

As a third and final validation strategy we calculated the number of inundated buildings due to both the Sonoma GIS map and the LISFLOOD model as a proxy for flood impacts. News reports following the 2019 event estimated that about 1,900 homes were damaged (Chavez, 2019). When overlaid with residual building locations, the Sonoma GIS map estimated 1,678 buildings with nonzero flood height, and the LISFLOOD model estimated 1,380.

3.6 Damage Measure (DM)

3.6.1 Component Model

The damage measure (DM) is the ratio of expected building repair costs to building value and is computed from inundation using a two step process. First, the first floor water level at each building is calculated by taking the inundation height relative to the ground surface ($INUN$ from the previous component model) and subtracting the height of the foundation. Foundation height information was collected at an aggregate level from Hazus-MH (FEMA, 2020), then assigned to individual buildings based on the distribution of foundation types and heights within that building's census tract. The foundation options were: basement (6 ft below grade), slab (1 ft above grade), crawl space (3–4 ft above grade), piers (5–6 ft above grade), or piles (7–8 ft above grade).

Second, the first floor water level was converted to a damage ratio using a depth-damage curve. We chose two sets of relationships: Hazus-MH depth-damage curves, which are widely used in engineering applications, and the beta distributions from Wing et al. (2020), which capture the uncertain and bimodal nature of residential flood damage at the household scale. A comparison of the damage ratios predicted by these relationships at various water depths is provided in Fig. 8(a).

3.6.2 2019 Case Study Event

Because there was little available in the way of site-specific damage information, we used building safety as a proxy variable for validation of the damage component model. Sonoma County performed over 2,000 building inspections as part of the Rapid Evaluation Safety Assessment (RESA) that immediately followed the 2019 event. Buildings were assigned colors based on these rapid inspections: green tags indicated that the structure was safe to enter, yellow indicated some risk, and red indicated an imminent safety threat. The RESA tags are categorical measures of safety and are thus an imperfect analog to the continuous

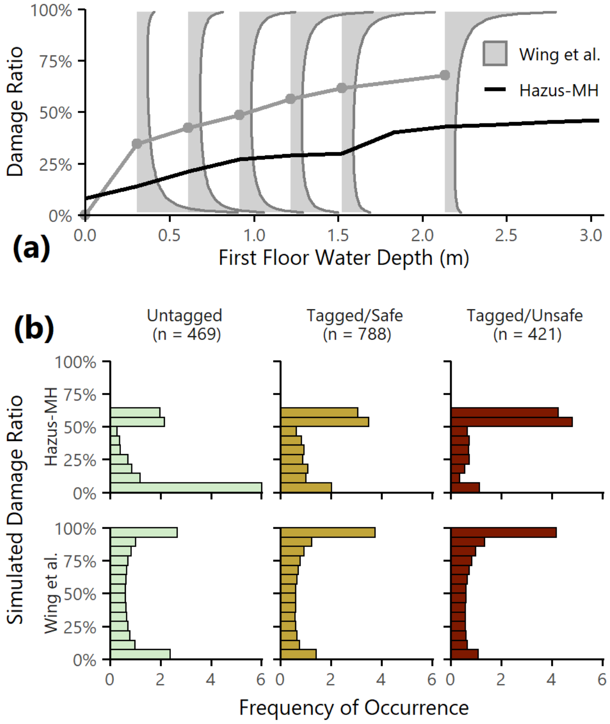

Figure 8. Damage component model. (a) Depth-damage relationships from Wing et al. (2020) and from Hazus-MH. The light grey areas represent the shapes of the beta distributions for the five defined curves (1, 2, 3, 4, 5, & 7 ft), and the grey line follows the distribution means. The black line is the Hazus-MH depth-damage function for single story residential buildings without a basement. **(b)** Damage ratio distributions for inundated buildings based on 1,000 Monte Carlo realization, using the Hazus-MH (top) and Wing et al. (2020) (bottom) depth-damage relationships and grouped by safety category.

damage ratios estimated by the damage component model. A building may have been tagged as unsafe for reasons beyond just inundation (i.e. roof damage or downed trees), and conversely a building that experienced nonzero inundation may still have been deemed safe to enter. Despite these limitations, comparing tagged buildings to our prediction of damaged buildings provides some intuition that damage is being predicted where actual damage was likely to occur.

We matched 1,253 of the RESA tags to residential buildings along the Russian River and aggregated them into three safety categories: untagged, tagged/safe (green), and tagged/unsafe (yellow+red). Using the Sonoma GIS map from the previous subsection we estimated how many buildings within the study area saw nonzero inundation at the ground surface (not accounting for foundation height). We see in Table 2 that 96–97% of all tagged building locations were predicted to have some level of inundation, compared to less than 5% for untagged building locations. This is consistent with our expectations because by definition any building with a RESA tag is one that was identified as a top priority for safety inspection and was therefore much more likely to have experienced inundation compared to an untagged building.

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

Table 2. RESA tags and inundated buildings in the study area following the 2019 event.

	Untagged	Tagged/ Safe	Tagged/ Unsafe
Number of Buildings	10,513	820	433
Inundated Buildings	469	788	421
Percent Inundated	4.46%	96.10%	97.23%

Figure 8(b) shows the expected distributions of damage ratios by safety category among the inundated buildings in the study area, as a function of the Hazus-MH (top) and Wing et al. (2020) (bottom) damage relationships. Although the shapes of the damage ratio distributions are notably different, we see overall increasing levels of damage with increasing tag severity. Of the small proportion of untagged buildings that saw nonzero inundation, most were predicted to receive little to no damage, and the tagged/unsafe category saw more severe damage outcomes than the tagged/safe category. These plots show that we are, in a broad sense, capturing more intense damage where we expected to do so.

3.7 Decision Variable (DV)

3.7.1 Component Model

The final component model in the PARRA framework converts the building-level damage ratios to a decision variable. The decision variable (DV) is some measurement of the impact or consequence of a hazard event. The consequences of flooding are traditionally described using the two axes of direct vs. indirect and tangible vs. intangible (Merz et al., 2004). This creates four categories: direct tangible (e.g. structural damage), indirect tangible (e.g. business interruption costs), direct intangible (e.g. loss of life), and indirect intangible (e.g. post-traumatic stress). The decision variable of interest for this case study is direct tangible loss, specifically the estimated total repair cost of residential buildings.

Losses were estimated for each building by multiplying the damage ratio times the expected value of the building as determined from Sonoma tax assessor roll data. Tax assessor data has some inherent limitations in California due to Proposition 13, which in many cases prevents assessments from directly tracking property values, so we applied correction factors at a census tract level such that the median from the tax assessor roll matched median value of owner-occupied housing units by census tract from the American Community Survey (ACS). When used in conjunction with the ACS correction factors, this high resolution tax assessor dataset allows us to match valuations to inundated buildings at the household level and more precisely estimate resulting losses.

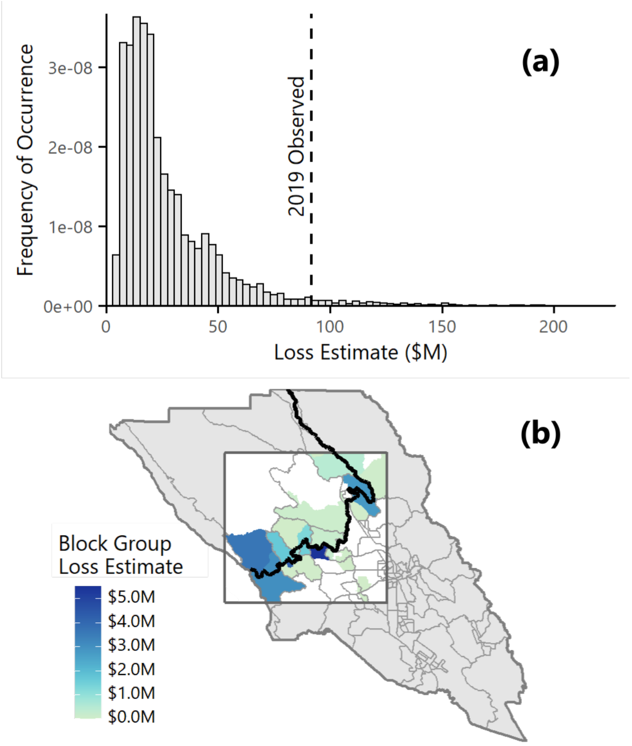

Figure 9. PARRA case study loss estimates. **(a)** Distribution of simulated loss realizations from the full PARRA framework for the 2019 event. The 2019 observed loss to residential buildings is marked by the dashed vertical line. **(b)** Spatial distribution of mean losses due to the 2019 event, aggregated to the census block group level.

3.7.2 2019 Case Study Event

It was not possible to validate the loss component model individually as we have done for all preceding component models. With neither damage ratios nor loss information at the household level, we had no accurate input data and no validation data to compare against the simulated output. Instead we considered a fully probabilistic estimate from the entire PARRA framework. Given the 2019 observed maximum IVT, duration, and soil moisture values, all subsequent component models were run probabilistically to produce a distribution of possible loss values. We generated 10,000 Monte Carlo loss realizations to estimate the full spectrum of potential loss outcomes, i.e. the flood losses that could have occurred for the 2019 event if any of the other pinch point variables had been different. We compared the observed vs. simulated losses and examined how the losses were spatially distributed within the study area.

The histogram of simulated loss realizations is shown in Fig. 9(a). The vertical dashed line marks $91.6 million, the estimate of true losses experienced by residential buildings in Sonoma County. The PARRA framework estimates this historical event to have been a 98th percentile loss event based on the given AR characteristics. The 2019 event was a Category 3 AR, which is broadly expected to produce a mixture of beneficial and hazardous conditions. Part of the reason this AR was chosen as a

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

case study, though, was because of the significant and severe consequences, so it is reasonable for the observed loss total to fall in the upper tail of the simulated distribution. It is also worth noting that the upper tail extends out to well over double the observed estimate, which indicates that while the loss experienced during the 2019 event was higher than expected for this particular combination of AR characteristics and antecedent conditions, it does not necessarily represent the worst case scenario of what we could see.

Figure 9(b) shows the spatial distribution of building losses, averaged across all simulated realizations and aggregated to the census block group level. Losses are concentrated along the banks of the Russian River, with particular hotspots near Healdsburg, Guerneville, and the mouth of the river near the Pacific Ocean. These areas are consistent with the locations of warnings and evacuation orders issued due to this storm event.

4 Results

In this section we utilize the framework to move beyond individual scenarios to capture the full distribution of AR-driven fluvial flood losses in Sonoma County. We define and report metrics of interest for community loss assessment, then consider how the framework can be used to evaluate potential mitigation policies.

4.1 Average Annual Loss

The first metric of interest is the average annual loss (AAL), or the long-run average of expected loss per year. The AAL is widely used in the insurance sector to price policies and is a convenient and well known metric to summarize risk.

To calculate the AAL we generated 38,200 simulated loss realizations: 382 events in the 32 year historic catalog times 100 realizations per event. This produced a synthetic stochastic record representing 3,200 (32 × 100) years of potential AR losses. We rank ordered the stochastic record and assigned each event a $\frac{1}{3200}$ annual rate of occurrence, then summed the product of the event rate times expected consequence, as described in Eq. 6. This calculation produces an empirical estimate of the AAL due to AR-induced flooding.

$$\mathrm{E}[DV] = \sum_{i=1}^{38,200} dv_i * \lambda(event_i) \tag{6}$$

where $event_i$ is one event in the stochastic record, dv_i is a sample realization of the decision variable DV due to event i, and $\lambda(event_i) = \frac{1}{3200}$ is the annual rate of occurrence.

The estimated AAL from the stochastic record for AR-induced flood losses to residential structures is \$163 million. We compared this to another loss estimate based on claims from the National Flood Insurance Program. We collected flood insurance claims within the counties that overlap with the study area for the years 1979-2018, then estimated the flood insurance penetration rate in these counties by dividing the average annual number of policies by the total number of households as determined from the American Community Survey. The average annual loss for the study area is then the average annual insured loss divided by the penetration rate, or \$156 million. It is important to acknowledge the significant uncertainty around both of these estimates that stems from the difficulties in characterizing the types of low probability, high consequence events that cause damaging impacts. However, we can use this comparison to confirm that the PARRA framework is able to broadly capture the full spectrum of AR-induced losses along the lower Russian River.

4.2 Loss Exceedance Curve

The second metric is the loss exceedance curve, which measures the expected annual rate of occurrence for a range of possible losses. The loss exceedance curve was introduced in Sect. 2 as the final outcome of the PARRA framework. This curve is the desired end product of a probabilistic risk assessment, summarizing information about expected likelihood and consequence across a range of potential risk events. It also provides insight about the character of the risk, such as whether frequent small events or a few rare events dominate the AAL. To our knowledge there has never been a loss exceedance curve created for AR-induced flood risk in Sonoma County.

We use the synthetic stochastic record generated above as a modeled representation of the event space. The loss exceedance curve is estimated by evaluating Eq. 7 at a range of x values. The result is shown as the black line in Fig. 10.

$$\lambda(DV > x) = \sum_{i=1}^{38,200} \mathrm{I}(dv_i > x) \, * \, \lambda(event_i) \tag{7}$$

where $\mathrm{I}(dv_i > x)$ is the binary indicator function measuring whether the the loss for event i exceeds the specified threshold x. Equation 7 is the empirical approximation of Eq. 1, using sample realizations of dv_i in place of an analytical $P(DV > x \,|\, AR)$.

With the loss exceedance curve we can examine risk thresholds such as the “100 year event” (1% annual rate of occurrence). This event, marked by the horizontal dashed line in Fig. 10, has a loss of \$197 million. Figure 10 also indicates that the loss associated with the 2019 event, \$91.6 million, has an annual occurrence rate of $\lambda = 0.152$ (approximately a 1 in 7 year event).

4.3 Mitigation Analysis

We show the performance-based capabilities of the PARRA framework by defining a loss reduction target and identifying mitigation actions that meet the specified threshold. We then quantify the impacts of the system changes on the shape of the loss exceedance curve to highlight the framework’s capability to prospectively assess events without historical precedent. This process of designing a system to meet a target performance objective is a key unique capability of performance-based frameworks.

Sonoma County is well aware of its flood risk and since 1995 has spent almost \$45 million dollars through the FEMA Hazard Mitigation Assistance Program to acquire or elevate 542 private properties. We explored the benefit of a hypothetical extension of the home elevation program and estimated how many homes would need to be elevated in order to reduce the AAL by half, from \$163 million to \$81 million. We ranked homes in order of increasing distance from the Russian River, elevated them one by one, and iteratively calculated the new AAL until it fell below the target threshold. We found that this strategy would require elevating approximately 150 additional homes in order to achieve the performance target. This is not necessarily the

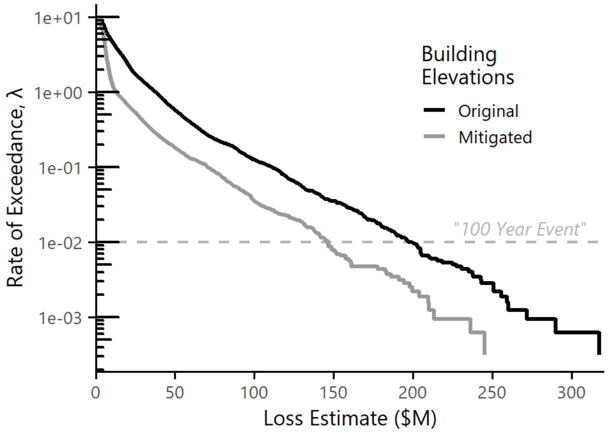

Figure 10. AR flood loss exceedance curve for Sonoma County. Loss exceedance is shown before and after the hypothetical mitigation action of elevating 136 households to reduce the AAL by half.

most efficient or effective possible strategy, but it illustrates in a straightforward manner how prospective evaluations can be performed.

The resultant change between the original and mitigated loss exceedance curves is shown in Fig. 10. The horizontal distance between the original and mitigated curves represents the benefit of mitigation at a specified rate of occurrence λ. For the the most frequent events ($\lambda > 1$) this distance is relatively small, but the benefit increases for larger events. The 100 year loss estimate drops from $197 million to $146 million, a difference of over $50 million dollars. The expected benefit for events with rates of occurrence rarer than the 100 year event ($\lambda \leq 0.01$) stays relatively constant around this value.

In addition to understanding the difference in expected loss at a constant return period, we can also pick a constant loss value and examine how the return period of that event has changed. For example, the 2019 case study event had a flood loss of $91.6 million, which was originally a 1 in 7 year return period loss. The new estimated occurrence rate for an event similar to the 2019 case study is of the 2019 event is $\lambda = 0.0491$, or a return period of approximately 20 years. This analysis indicates that there is significant potential benefit to continuing investment in home elevations as a flood loss mitigation strategy in Sonoma County. These types of insights can help community members and stakeholders make more informed decisions about their level of flood risk and the effect of potential mitigation actions.

5 Discussion

Implementing a large, complex, multi-disciplinary model sequence such as the PARRA framework involves challenges and compromises, but offers significant potential for novel insights. In this section we highlight the more subtle benefits of the probabilistic, performance-based approach proposed here and discuss some of the practicalities involved in implementation and validation.

5.1 Scenario Disaggregation

The PARRA framework provides insights regarding which of the pinch points made a particular event "extreme" in terms of its associated losses, and what factors might have led to a more or less desirable outcome. For example, if the peak streamflow was higher than expected for a given storm, the framework can indicate whether the precipitation or the antecedent hydrologic conditions were contributing more to that deviation. The abnormally high loss experienced during the 2019 event was likely attributable to the high antecedent soil moisture and the relatively high precipitation given the maximum IVT and duration (Fig. 4). We found that the estimated streamflow and inundation, conditioned on these higher-than-average forcing variables, were not significantly different from the observed values. The new information gained from each component model individually paints a fuller picture of these phenomena overall.

5.2 Framework Implementation

Another strength of the PARRA framework is its ability to track and quantify uncertainty across multiple component models. Each component model produces a best-fit estimate and a range of uncertainty for the output pinch point variable of interest through Monte Carlo simulation. The benefit of Monte Carlo simulation is that the component models do not need to be resolved into analytical PDFs; instead pinch point variables are calculated directly using the component models, then propagated through the full equation to produce an empirical distribution of total loss. However, a downside is the computational effort involved in generating these realizations, especially for time- or resource-intensive component models.

We made several implementation decisions to minimize the computational expense of the PARRA framework without compromising accuracy of the results. Two of those decisions are highlighted here. First, we implemented a functional programming paradigm in R to run the framework and leveraged parallelization and high performance computing resources to speed calculations. Second, we limited computational costs of individual component models by relying on statistical relationships rather than high-fidelity physical models (e.g. $f(Q|PRCP, HC)$, Sect. 3.4) and implementing surrogate models to replace expensive calculation steps (e.g. $f(INUN|Q)$, Sect. 3.5). The repetitive nature of Monte Carlo simulation means that savings in any one component model lead to multiplicative effects when running the entire framework in sequence.

5.3 Validation Data

In the Sect. 3 case study, significant effort went into finding validation data, both for the overall component model fit/calibration and for the specific 2019 event under consideration. The datasets collected as part of this project varied widely in spatial and temporal resolution, as well as in the ease of access. Atmospheric information, precipitation, and streamflow timeseries data were all readily available from academic or governmental institutions. Most of these institutions were involved not only in data collection, but also in the curation and maintenance of web architectures that allowed easy navigation of these datasets. Gridded maps of inundation extents, while not available for many historical events, are becoming more prevalent as remote sensing moves this particular task from one that must be completed on-the-ground to one that can be automated and released in near real time. Damage and loss data, though, are still primarily recorded at the local level through surveys, home inspections, and

https://doi.org/10.5194/nhess-2021-337
Preprint. Discussion started: 12 November 2021

other resource intensive methods. Difficulties surrounding data collection and individual privacy concerns mean that there are far less publicly available data for analyses of these steps of the model sequence. As a consequence, moving from inundation to damage and from damage to loss are by far the most uncertain aspects of flood risk assessment due to the documented difficulties in validation against observed data (Apel et al., 2009; Gerl et al., 2016).

Two potential pathways could help to close this data gap. First, insurance companies often have extensive information about flood damage and loss at the household level, but offer few opportunities for access. New models for academic partnerships that permit data sharing while allowing insurance companies to retain their competitive market advantage could help further research in this area. Second, Sonoma County, like a growing number of cities and counties nationwide, maintains an open data portal that proved to be invaluable to this project. Expanding the quantity of these data offerings, and adding damage and loss estimation data at the local level, would improve the outputs of analytical modeling efforts and in turn produce more relevant findings for community level decision making.

5.4 Component Model Alternatives

The case study developed in Sect. 3 shows only one possible implementation choice for each component model. As long as connections between the pinch point variables are maintained, any of the component models could be replaced with an alternative implementation. There are multiple possible methods to characterize some of the pinch point variables that would improve fidelity to the underlying physical processes, at the cost of increased computational demand. For example, the linear regression used for the precipitation component model could be replaced with a numerical weather prediction scheme that explicitly captures the effects of additional factors such as AR orientation. The precipitation pinch point variable itself could be modified to represent the full spatial distribution of precipitation rather than an areal average. The depth-damage relationships used in the the damage component model could be replaced with an assembly-based vulnerability model, such that the damage to a building is measured as the sum of all damages to its component parts. Many other substitutions and modifications to the component models are possible and could be implemented without changing the underlying underlying structure or the output metrics produced by the PARRA framework.

6 Conclusions

This paper introduced the Performance-based Atmospheric River Risk Analysis (PARRA) framework to quantify AR-induced flood risk. The framework captures the physical processes connecting atmospheric forcings, hydrologic impacts, and economic consequences of AR-driven fluvial flooding. Using a performance-based engineering paradigm, this approach offers several benefits. It quantifies the uncertainty surrounding the physical processes by following a deliberate, ordered simulation procedure. It connects multiple physical processes in sequence by constructing a chain of discrete component models that link together at defined pinch points. Pinch points serve to facilitate intercompatibility of models from different disciplines to better understand this multifaceted hazard.

We demonstrated the uncertainty quantification capabilities and the modularity of the component models through a case study of a historic AR event affecting the lower Russian River in Sonoma County. Section 3 discussed the fit and calibration of five individual component models: precipitation estimation, hydrologic routing, inundation modeling, depth-damage relationships, and loss estimation. We performed a step-by-step comparison between each of these component models and ground-truth data from a damaging 2019 AR event in Sonoma County to show that the PARRA framework could reproduce observed results. We then ran a fully probabilistic simulation using only the observed AR characteristics and antecedent soil moisture as input and examined both the probabilistic range and the spatial distribution of the predicted losses. The total losses to residential homes associated with the 2019 case study event fell within the expected probabilistic range estimated by the PARRA framework, and the regions with the highest predicted total losses matched communities with known historic vulnerability to flooding.

In Sect. 4, the PARRA framework was used to generate a first-of-its-kind loss exceedance curve for the lower Russian River to understand the return periods of a full spectrum of potential loss events rather than a single scenario event or a long term annual average. We compared the current loss exceedance curve to one with elevated homes to quantify the reduction in flood risk from a hypothetical mitigation decision. When 150 homes were elevated above the 100 year flood elevation, the AAL was reduced by half, and the average benefit for events with return periods of 100 years and longer was found to be $50 million per AR. The PARRA framework is ideally suited to prospectively evaluate potential impacts for events outside of the historic record, or events that have not yet occurred but could in an evolving climate. It can similarly be used to estimate changes in future flood risk due to land use shifts, population change, and more. The framework presented here has been shown to work in a real world implementation and has the potential to greatly expand our understanding of the risks associated with AR-induced flooding.

Code and data availability.

Readers are encouraged to explore the supplemental code release associated with this paper, hosted at https://github.com/corinnebowers/PARRA.

All data used for the historic catalog and for the case study are free and publicly available. Selected sources are included below.

Watershed boundaries and hydrography information were retrieved from the USGS National Map at https://apps.nationalmap.gov/downloader/.

Census boundaries and designations were retrieved from the US Census Bureau (https://www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html) using the R package *tigris* (https://CRAN.R-project.org/package=tigris).

Precipitation data came from the CPC Hourly US Precipitation dataset, provided by the NOAA/OAR/ESRL PSL in Boulder, CO and retrieved using the R package *rnoaa* (https://CRAN.R-project.org/package=rnoaa).

Streamflow data was retrieved from the USGS National Water Information System (http://dx.doi.org/10.5066/F7P55KJN) using the R package *dataRetrieval* (https://code.usgs.gov/water/dataRetrieval).

The bare earth digital elevation model for the study area was retrieved from Sonoma Veg Map at http://sonomavegmap.org/.

The 100-year return period streamflow at USGS gage 11463500 (study area inlet) was calculated using the USGS StreamStats application at http://streamstats.usgs.gov/.

The 100-year calculated inundation extent for Sonoma County was retrieved from the FEMA National Flood Hazard Layer at https://www.fema.gov/flood-maps/national-flood-hazard-layer.

Predicted inundation maps for the lower Russian River are available from the County of Sonoma at https://sonomacounty.maps.arcgis.com/home/item.html?id=9d8d63558c6b4124b000e6476a0a020d.

Building footprints are available from the County of Sonoma at https://sonomacounty.maps.arcgis.com/home/item.html?id=0f5982c3582d4de0b811e68d7f0bff8f.

Building occupancy classes and valuations are available from the County of Sonoma, Clerk Recorder Assessor at https://sonomacounty.maps.arcgis.com/home/item.html?id=2202c1cd6708441f987ca5552f2d9659.

Rapid Evaluation Safety Assessment (RESA) tag information following the 2019 flood event in Sonoma County was retrieved from Permit Sonoma at https://sonomacounty.ca.gov/PRMD/Eng-and-Constr/Building/RESA-2019-Flooding/.

Median values of owner-occupied housing units by census tract were retrieved from the American Community Survey using the R package *censusapi* (https://cran.r-project.org/web/packages/censusapi/vignettes/getting-started.html).

Information about National Flood Insurance Program claims and policies and about the FEMA Hazard Mitigation Assistance program was retrieved from the OpenFEMA data platform at https://www.fema.gov/about/openfema/data-sets.

All figures with color were generated using the *roma* scientific color map from Crameri et al. (2020).

Author contributions. CCB formulated the research concept and conceptualized the PARRA framework with input and assistance from JWB and KAS. CCB performed numerical analyses, generated figures and results, and prepared the original draft of the manuscript. All authors were involved in editing and review of the submitted manuscript.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This material is based upon work supported by both the Stanford Gabilan Graduate Fellowship and the National Science Foundation Graduate Research Fellowship under Grant No. 1000265549. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Some of the computing for this project was performed on the Sherlock cluster. We would like to thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that contributed to these research results.

References

Albano, C. M., Dettinger, M. D., and Harpold, A. A.: Patterns and Drivers of Atmospheric River Precipitation and Hydrologic Impacts across the Western United States, Journal of Hydrometeorology, 21, 143–159, https://doi.org/10.1175/jhm-d-19-0119.1, 2020.

Apel, H., Thieken, A. H., Merz, B., and Blöschl, G.: Flood risk assessment and associated uncertainty, Natural Hazards and Earth System Sciences, 4, 295–308, https://www.nat-hazards-earth-syst-sci.net/4/295/2004/nhess-4-295-2004.pdf, 2004.

Apel, H., Aronica, G. T., Kreibich, H., and Thieken, A. H.: Flood risk analyses - How detailed do we need to be?, Natural Hazards, 49, 79–98, https://doi.org/10.1007/s11069-008-9277-8, 2009.

Baker, J. W., Bradley, B. A., and Stafford, P. J.: Seismic Hazard and Risk Analysis, Cambridge University Press, Cambridge, England, 2021.

Barbato, M., Petrini, F., Unnikrishnan, V. U., and Ciampoli, M.: Performance-Based Hurricane Engineering (PBHE) framework, Structural Safety, 45, 24–35, https://doi.org/10.1016/J.STRUSAFE.2013.07.002, 2013.

Bass, B. and Bedient, P.: Surrogate modeling of joint flood risk across coastal watersheds, Journal of Hydrology, 558, 159–173, https://doi.org/10.1016/j.jhydrol.2018.01.014, 2018.

Bates, P. D. and De Roo, A.: A simple raster-based model for flood inundation simulation, Journal of Hydrology, 236, 54–77, https://doi.org/10.1016/S0022-1694(00)00278-X, 2000.

Brunner, G. W.: HEC-RAS 6.0 Hydraulic Reference Manual, Tech. rep., US Army Corps of Engineers (USACE), Davis, CA, 2020.

Cao, Q., Mehran, A., Ralph, F. M., and Lettenmaier, D. P.: The Role of Hydrological Initial Conditions on Atmospheric River Floods in the Russian River Basin, Journal of Hydrometeorology, 20, 1667–1686, https://doi.org/10.1175/JHM-D-19-0030.1, 2019.

Chavez, N.: Flooding in Sonoma County causes estimated $155 million in damage, https://www.pressdemocrat.com/article/news/flooding-in-sonoma-county-causes-estimated-155-million-in-damage/, 2019.

Chen, X., Leung, L. R., Gao, Y., Liu, Y., Wigmosta, M., and Richmond, M.: Predictability of Extreme Precipitation in Western U.S. Watersheds Based on Atmospheric River Occurrence, Intensity, and Duration, Geophysical Research Letters, 45, 693–11, https://doi.org/10.1029/2018GL079831, 2018.

Ciampoli, M., Petrini, F., and Augusti, G.: Performance-Based Wind Engineering: Towards a general procedure, Structural Safety, 33, 367–378, https://doi.org/10.1016/J.STRUSAFE.2011.07.001, 2011.

Corringham, T. W., Ralph, F. M., Gershunov, A., Cayan, D. R., and Talbot, C. A.: Atmospheric rivers drive flood damages in the western United States, Science Advances, 5, http://advances.sciencemag.org/, 2019.

Crameri, F., Shephard, G. E., and Heron, P. J.: The misuse of colour in science communication, Nature Communications, 11, 1–10, https://doi.org/10.1038/s41467-020-19160-7, 2020.

Dominguez, F., Dall'erba, S., Huang, S., Avelino, A., Mehran, A., Hu, H., Schmidt, A., Schick, L. J., and Lettenmaier, D. P.: Tracking an atmospheric river in a warmer climate: from water vapor to economic impacts, Earth System Dynamics, 9, 249–266, https://doi.org/10.5194/esd-9-249-2018, 2018.

Felder, G., Gómez-Navarro, J. J., Zischg, A. P., Raible, C. C., Röthlisberger, V., Bozhinova, D., Martius, O., and Weingartner, R.: From global circulation to local flood loss: Coupling models across the scales, Science of the Total Environment, 635, 1225–1239, https://doi.org/10.1016/j.scitotenv.2018.04.170, 2018.

FEMA: Hazus-MH Technical Manual Flood Model, Washington D.C: Department of Homeland Security, 2020.

Flint, M. M., Baker, J. W., and Billington, S. L.: A modular framework for performance-based durability engineering: From exposure to impacts, Structural Safety, 50, 78–93, https://doi.org/10.1016/j.strusafe.2014.03.003, 2014.

Garrick, B. J.: Recent Case Studies and Advancements in Probabilistic Risk Assessment, Risk Analysis, 4, 1984.

Gelaro, R., McCarty, W., Suárez, M. J., Todling, R., Molod, A., Takacs, L., Randles, C. A., Darmenov, A., Bosilovich, M. G., Reichle, R., Wargan, K., Coy, L., Cullather, R., Draper, C., Akella, S., Buchard, V., Conaty, A., da Silva, A. M., Gu, W., Kim, G. K., Koster, R., Lucchesi, R., Merkova, D., Nielsen, J. E., Partyka, G., Pawson, S., Putman, W., Rienecker, M., Schubert, S. D., Sienkiewicz, M., and Zhao, B.: The modern-era retrospective analysis for research and applications, version 2 (MERRA-2), Journal of Climate, 30, 5419–5454, https://doi.org/10.1175/JCLI-D-16-0758.1, 2017.

Gerl, T., Kreibich, H., Franco, G., Marechal, D., and Schröter, K.: A Review of Flood Loss Models as Basis for Harmonization and Benchmarking, PLoS ONE, 11, https://doi.org/10.1371/journal.pone.0159791, 2016.

Guirguis, K., Gershunov, A., Clemesha, R. E. S., Shulgina, T., Subramanian, A. C., and Ralph, F. M.: Circulation Drivers of Atmospheric Rivers at the North American West Coast, Geophysical Research Letters, 45, https://doi.org/10.1029/2018GL079249, 2018.

Guo, Q., Shi, K., Jia, Z., and Jeffers, A. E.: Probabilistic Evaluation of Structural Fire Resistance, Fire Technology, 49, 793–811, https://doi.org/10.1007/s10694-012-0293-6, 2013.

Hecht, C. W. and Cordeira, J. M.: Characterizing the influence of atmospheric river orientation and intensity on precipitation distributions over North Coastal California, Geophysical Research Letters, 44, 9048–9058, https://doi.org/10.1002/2017GL074179, 2017.

Huang, X., Swain, D. L., Walton, D. B., Stevenson, S., and Hall, A. D.: Simulating and Evaluating Atmospheric River-Induced Precipitation Extremes Along the U.S. Pacific Coast: Case Studies From 1980–2017, Journal of Geophysical Research: Atmospheres, 125, 1–18, https://doi.org/10.1029/2019JD031554, 2020.

Konrad, C. P. and Dettinger, M. D.: Flood Runoff in Relation to Water Vapor Transport by Atmospheric Rivers Over the Western United States, 1949–2015, Geophysical Research Letters, 44, 456–11, https://doi.org/10.1002/2017GL075399, 2017.

Krawinkler, H.: Challenges and progress in performance-based earthquake engineering, International Seminar on Seismic Engineering for Tomorrow, pp. 1–10, 1999.

Lamjiri, M. A., Dettinger, M. D., Ralph, F. M., Oakley, N., and Rutz, J. J.: Hourly Analyses of the Large Storms and Atmospheric Rivers that Provide Most of California's Precipitation in Only 10 to 100 Hours per Year, San Francisco Estuary and Watershed Science, 16, https://doi.org/10.15447/sfews.2018v16iss4art1, 2018.

Lavers, D. A., Bruce Ingleby, N., Subramanian, A. C., Richardson, D. S., Martin Ralph, F., Doyle, J. D., Reynolds, C. A., Torn, R. D., Rodwell, M. J., Tallapragada, V., and Pappenberger, F.: Forecast errors and uncertainties in atmospheric rivers, Weather and Forecasting, 35, 1447–1458, https://doi.org/10.1175/WAF-D-20-0049.1, 2020.

Martin, A. C., Ralph, F. M., Demirdjian, R., DeHaan, L., Weihs, R., Helly, J., Reynolds, D., and Iacobellis, S.: Evaluation of atmospheric river predictions by the WRF model using aircraft and regional mesonet observations of orographic precipitation and its forcing, Journal of Hydrometeorology, 19, 1097–1113, https://doi.org/10.1175/JHM-D-17-0098.1, 2018.

Merz, B., Kreibich, H., Thieken, A. H., and Schmidtke, R.: Estimation uncertainty of direct monetary flood damage to buildings, Natural Hazards and Earth System Sciences, 4, 153–163, https://hal.archives-ouvertes.fr/hal-00299090, 2004.

Nardi, K. M., Barnes, E. A., and Ralph, F. M.: Assessment of Numerical Weather Prediction Model Reforecasts of the Occurrence, Intensity, and Location of Atmospheric Rivers along the West Coast of North America, Monthly Weather Review, 146, https://doi.org/10.1175/MWR-D-18, 2018.

NRC: Reactor Safety Study. An Assessment of Accident Risks in US Commercial Nuclear Power Plants. Executive Summary, Tech. rep., WASH-1400. United States Nuclear Regulatory Commission, https://doi.org/10.2172/7134131, 1975.

NRCS: Part 630 Hydrology, in: National Engineering Handbook, US Department of Agriculture, Washington, DC, 2004.

Pinelli, J.-P. and Barbato, M.: A comparison study of Catastrophe Modeling vs. Performance-based design, in: 13th International Conference on Applications of Statistics and Probability in Civil Engineering, ICASP 2019, 2019.

Porter, K., Wein, A., Alpers, C., Baez, A., Barnard, P., Carter, J., Corsi, A., Costner, J., Cox, D., Das, T., Dettinger, M. D., Done, J., Eadie, C., Eymann, M., Ferris, J., Gunturi, P., Hughes, M., Jarrett, R., Johnson, L., Dam Le-Griffen, H., Mitchell, D., Morman, S., Neiman, P. J., Olsen, A., Perry, S., Plumlee, G., Ralph, F. M., Reynolds, D., Rose, A., Schaefer, K., Serakos, J., Siembieda, W., Stock, J., Strong, D., Sue Wing, I., Tang, A., Thomas, P., Topping, K., and Wills, C.: Overview of the ARkStorm scenario, Tech. rep., US Geological Survey, Reston, VA, http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Overview+of+the+ARkStorm+Scenario#0, 2011.

Ralph, F. M., Neiman, P. J., Wick, G. A., Gutman, S. I., Dettinger, M. D., Cayan, D. R., and White, A. B.: Flooding on California's Russian River: Role of atmospheric rivers, Geophysical Research Letters, 33, https://doi.org/10.1029/2006GL026689, 2006.

Ralph, F. M., Coleman, T., Neiman, P. J., Zamora, R. J., and Dettinger, M. D.: Observed Impacts of Duration and Seasonality of Atmospheric-River Landfalls on Soil Moisture and Runoff in Coastal Northern California, Journal of Hydrometeorology, 14, https://doi.org/10.1175/JHM-D-12-076.1, 2013.

Ralph, F. M., Dettinger, M. D., Cairns, M. M., Galarneau, T. J., and Eylander, J.: Defining "Atmospheric river" : How the glossary of meteorology helped resolve a debate, Bulletin of the American Meteorological Society, 99, 837–839, https://doi.org/10.1175/BAMS-D-17-0157.1, 2018.

Ralph, F. M., Rutz, J. J., Cordeira, J. M., Dettinger, M. D., Anderson, M., Reynolds, D., Schick, L. J., and Smallcomb, C.: A scale to characterize the strength and impacts of atmospheric rivers, Bulletin of the American Meteorological Society, 100, 269–289, https://doi.org/10.1175/BAMS-D-18-0023.1, 2019.

Razavi, S., Tolson, B. A., and Burn, D. H.: Review of surrogate modeling in water resources, Water Resources Research, 48, https://doi.org/10.1029/2011WR011527, 2012.

Rogers, P.: Guerneville floods more than anywhere in the Bay Area. Why can't it be fixed?, https://www.mercurynews.com/2019/03/04/why-does-guerneville-flood-so-much-and-why-cant-it-be-fixed/, 2019.

Rutz, J. J., James Steenburgh, W., Ralph, F. M., and Steenburgh, W. J.: Climatological characteristics of atmospheric rivers and their inland penetration over the western united states, Monthly Weather Review, 142, 905–921, https://doi.org/10.1175/MWR-D-13-00168.1, 2014.

Shields, C. A., Rutz, J. J., Leung, L.-Y. Y., Ralph, F. M., Wehner, M., Kawzenuk, B., Lora, J. M., Mcclenny, E., Osborne, T., Payne, A. E., Ullrich, P., Gershunov, A., Goldenson, N., Qian, Y., Ramos, A. M., Sarangi, C., Sellars, S., Gorodetskaya, I. V., Kashinath, K., Kurlin, V., Mahoney, K. M., Muszynski, G., Pierce, R., Subramanian, A. C., Tome, R., Waliser, D. E., Walton, D., Wick, G. A., Wilson, A., Lavers, D. A., Collow, A., Krishnan, H., Magnusdottir, G., Nguyen, P., Guan, B., and Prabhat: Atmospheric River Tracking Method Intercomparison Project (ARTMIP): project goals and experimental design, Geoscientific Model Development, 11, 2455–2474, https://doi.org/10.5194/gmd-11-2455-2018, 2018.

Sonoma County: Flood Hazards, Tech. rep., 2017.

Uhe, P., Mitchell, D., Bates, P. D., Allen, M. R., Betts, R. A., and Sanderson, B. M.: Method-uncertainty is essential for reliable confidence statements of precipitation projections, Journal of Climate, pp. 1–45, https://doi.org/10.1175/JCLI-D-20-0289.1., 2020.

USACE: HEC-HMS Technical Reference Manual, Tech. rep., US Army Corps of Engineers, 2020.

van den Dool, H., Huang, J., and Fan, Y.: Performance and analysis of the constructed analogue method applied to U.S. soil moisture over 1981-2001, Journal of Geophysical Research: Atmospheres, 108, 1–16, https://doi.org/10.1029/2002jd003114, 2003.

Waliser, D. E. and Guan, B.: Extreme winds and precipitation during landfall of atmospheric rivers, Nature Geoscience, 10, 179–183, https://doi.org/10.1038/ngeo2894, 2017.

Wing, O. E. J., Bates, P. D., Sampson, C. C., Smith, A. M., Johnson, K. A., and Erickson, T. A.: Validation of a 30 m resolution flood hazard model of the conterminous United States, Water Resources Research, 53, 7968–7986, https://doi.org/10.1002/2017WR020917, 2017.

Wing, O. E. J., Pinter, N., Bates, P. D., and Kousky, C.: New insights into US flood vulnerability revealed from flood insurance big data, Nature Communications, 11, 1444, https://doi.org/10.1038/s41467-020-15264-2, 2020.

Zhu, Y. and Newell, R. E.: A proposed algorithm for moisture fluxes from atmospheric rivers, Monthly Weather Review, 126, 725–735, https://doi.org/10.1175/1520-0493(1998)126<0725:APAFMF>2.0.CO;2, 1998.