Earth's Future

RESEARCH ARTICLE
10.1029/2023EF003571

Key Points:

- Flood insurance purchase in the US is dominated by reactive behavior after severe floods
- The Community Rating System (CRS) fosters proactive insurance adoption irrespective of socio-economic background
- The CRS should further balance existing inequalities by targeting specific population segments

Supporting Information:

Supporting Information may be found in the online version of this article.

Correspondence to:

N. Veigel,
nadja.veigel@tu-berlin.de

Citation:

Veigel, N., Kreibich, H., & Cominola, A. (2023). Interpretable machine learning reveals potential to overcome reactive flood adaptation in the continental US. *Earth's Future*, *11*, e2023EF003571. https://doi.org/10.1029/2023EF003571

Received 5 FEB 2023
Accepted 6 SEP 2023

Author Contributions:

Conceptualization: Nadja Veigel, Heidi Kreibich, Andrea Cominola
Data curation: Nadja Veigel
Funding acquisition: Heidi Kreibich, Andrea Cominola
Methodology: Nadja Veigel, Andrea Cominola
Resources: Heidi Kreibich, Andrea Cominola
Software: Nadja Veigel, Andrea Cominola
Supervision: Heidi Kreibich, Andrea Cominola
Validation: Nadja Veigel
Visualization: Nadja Veigel
Writing – original draft: Nadja Veigel

Interpretable Machine Learning Reveals Potential to Overcome Reactive Flood Adaptation in the Continental US

Nadja Veigel[1,2,3], Heidi Kreibich[2], and Andrea Cominola[1,3]

[1]Chair of Smart Water Networks, Technische Universität Berlin, Berlin, Germany, [2]Section 4.4 Hydrology, GFZ German Research Centre for Geosciences, Potsdam, Germany, [3]Einstein Center Digital Future, Berlin, Germany

Abstract Floods cause average annual losses of more than US$30 billion in the US and are estimated to significantly increase due to global change. Flood resilience, which currently differs strongly between socio-economic groups, needs to be substantially improved by proactive adaptive measures, such as timely purchase of flood insurance. Yet, knowledge about the state and uptake of private adaptation and its drivers is so far scarce and fragmented. Based on interpretable machine learning and large insurance and socio-economic open data sets covering the whole continental US we reveal that flood insurance purchase is characterized by reactive behavior after severe flood events. However, we observe that the Community Rating System helps overcome this behavior by effectively fostering proactive insurance purchase, irrespective of socio-economic backgrounds in the communities. Thus, we recommend developing additional targeted measures to help overcome existing inequalities, for example, by providing special incentives to the most vulnerable and exposed communities.

Plain Language Summary Flood resilience of individuals and communities can be improved by bottom-up strategies, such as insurance purchase, or top-down measures like the US National Flood Insurance Program's Community Rating System (CRS). Our interpretable machine learning approach shows that flood insurances are mostly purchased reactively, after the occurrence of a flood event. Yet, reactive behaviors are ill-suited as more extreme events are expected under future climate, also in areas that were not previously flooded. The CRS counteracts this behavior by fostering proactive adaptation across a widespread range of socio-economic backgrounds. Future risk management including the CRS should support and motivate individuals' proactive adaptation with a particular focus on highly vulnerable social groups to overcome existing inequalities in flood risk.

1. Introduction

Average flood losses in the USA amount to US$32.1 billion per year and flood risk is estimated to grow by 26.4% in 2050 (Wing et al., 2022). Impact patterns are inequitable across different socio-economic groups (Wing et al., 2022). Thus, there is an urgent need for improved and more equitable adaptation to flood risk in the US.

Proactive adaptation strategies, that is, informed actions based on risk assessments that are adopted before a flood hits, can be driven by top-down measures (e.g., community-level policies or infrastructural interventions for flood protection) or bottom-up actions (e.g., purchase of flood insurance by households). Household adaptation, for instance through insurance purchase (Cremades et al., 2018), has the potential to lower future flood risk substantially by reducing the financial vulnerability to residual risk (Aerts et al., 2018; Jongman et al., 2015; Kundzewicz et al., 2018).

In the US, top-down and bottom-up flood adaptation measures are interlinked within the mechanism that relates individuals' flood insurance access and communities' flood risk management. Flood insurance adoption as a bottom-up measure is regulated within the National Flood Insurance Program (NFIP). The NFIP was founded in 1968 as the only source of flood insurance in the US. Communities maintain minimum floodplain management standards to make insurance accessible for their inhabitants (Horn & Brown, 2018). Within the NFIP, the Community Rating System (CRS) (FEMA, 2021c) is a top-down strategy to encourage insurance adoption. If a community agrees to undertake flood risk mitigation and floodplain management measures, it is ranked with a higher CRS class and the insurance premiums in the community will be lowered (Sadiq et al., 2019).

Writing – review & editing: Nadja Veigel, Heidi Kreibich, Andrea Cominola

Previous studies have shown that individuals' adaptation actions are potentially impacted by a large number of demographic, socio-economic, and psychological factors (Koerth et al., 2017; McPhillips et al., 2018), including risk perception (Bubeck et al., 2012; Sanders et al., 2020), preparedness, and social networks (Hu, 2020). For instance, clear increases of flood insurance take-up rates were observed after the occurrence of catastrophic floods by several studies (Gallagher, 2014; Kuang & Liao, 2020). Moreover, resilience differs quite strongly between different socio-economic groups and there are problems of inequitable access to adaptation measures and unbalanced consideration of different socio-demographic groups in flood risk management policies (Knighton et al., 2021; Wing et al., 2022). A US wide study with focus on selected metropolitan areas showed a disproportional exposure of metropolitan areas with a higher percentage of Black, Indigenous, and People of Color (BIPoC) residents (Knighton et al., 2021). Another study in Georgia showed that African-Americans as well as educated and older people were more likely to purchase flood insurance (Atreya et al., 2015). Conversely, Cannon et al. (2020) found that the dominant ethnicity in a neighborhood was not significantly associated with insurance coverage, whereas Knighton et al. (2021) describe unequal exposure patterns.

Overall, the literature on the socio-economic drivers of bottom-up adaptation measures in the US is heterogeneous and in some cases contradictory. This could be a result of spatially heterogeneous individual reactions or analysis run on small samples. This reveals a research gap: to generalize beyond fragmented results, studies are needed that provide large scale evidence on the drivers of bottom-up flood adaptation strategies in combination with community-level policies.

Here, we investigate the relationship between socio-economic backgrounds, household flood insurance purchase (bottom-up), and the community-scale CRS (top-down) in the whole continental US. We aim to reveal the main drivers of insurance uptake and explore the potential of the CRS to foster proactive adaptation and overcome inequitable flood risk management across household and community scales, without any a priori assumption on individuals' motivations, or utility maximization. We formulate the following research questions:

1. What are the main socio-economic and behavioral drivers of flood insurance purchase in the US? Is the CRS effective in fostering individuals' flood insurance purchase?
2. Which patterns emerge in the socio-economic composition of communities that are active in the CRS? Are there inequalities in the representation of different social groups?

2. Materials and Methods

To address the above questions, we develop a data-driven method based on interpretable machine learning (ML). We investigate the role of socio-economic characteristics, flood history, and participation in the CRS with respect to flood insurance purchase as for the whole continental US. Similarly, we model communities' participation in the CRS as a function of flood history and socio-economic characteristics.

2.1. Data Sources and Data Preparation

We analyze open-access flood insurance data by the NFIP and the American Community Survey (ACS) by the US Census Bureau with a census tract resolution to model the policy records since 2009 that were in-force at the time of data retrieval (FEMA, 2020). From the ACS we initially selected 400 different potential behavioral and socio-economic predictors to model individuals' and communities' flood resilience and compare it across different census tracts. Additionally we model CRS participation probability (FEMA, 2021a).

The locations of the NFIP policies-in-force (FEMA, 2020) are redacted to 0.1° precision upon retrieval. To counteract the resulting high concentration of points (in a few census tracts), we apply a random spread to the points. Finally, we spatially aggregate this data to match the census tract scale of the ACS data. Thus, we total the data points within the census tracts and divide by the number of housing units. We then use the resulting processed data as output variable to train a ML model to predict household flood insurance purchase as a function of the potential behavioral and socio-economic predictors. In building the set of candidate model input features, or predictors, we manually select 400 variables from more than 25,000 census variables contained in the ACS data set. Feature values were retrieved as total counts and divided by the number of inhabitants in a census tract to estimate the percentage of people showing a respective feature. We consider the claims filed within the NFIP as a proxy to estimate past average annual number of floods (feature called flood history, *regular*) and the maximum amount of people affected in a flood per census tract (FEMA, 2019) to estimate flood severity (flood history, *severe*). This strategy may result in an under-estimation of under-insured census tracts for the flood severity variable. The area within the flood zone in a census tract is calculated based on the official Special Flood Hazard Area (SFHA) by FEMA (FEMA, 2021b) and distinguished between coastal (flood zone V) and fluvial flood zones (flood zone A). Two features describing whether the insurance policy covered a building or its contents are calculated based on the metadata of the policies in force indicating the percentage of policies covering contents and those covering buildings. We find no significant pair-wise variable correlations in our input variables, which we assume to be a result of spatial heterogeneity and locally differing effects. The resulting data set contains 400 features for 72,366 samples (census tracts) across all states in the continental US.

2.2. Regression of Flood Insurance Purchase per Household and Classification of CRS Participation

We first perform a regression task to estimate insurance coverage in a census tract based on the 400 features described in the previous section using a Gradient Boosting Decision Trees (GBDT) framework. In GBDT an ensemble model is assembled by step-wise addition of decision trees to minimize the residual error of a tree ensemble. We use an efficient form of state-of-the-art GBDT called Light Gradient Boosting Machine (LightGBM) (Ke et al., 2017). LightGBM builds on Gradient-based-One Side-Sampling and Exclusive Feature Bundling to accelerate model training and achieve higher prediction accuracy (Ke et al., 2017). We train the LightGBM using data from 80% of the census tracts (n = 54,274) and test on the remaining 20% (n = 18,092) (see Text S1.3 in Supporting Information S1). We train the LightGBM model both at the aggregate level for the whole US, as well as a series of individual models, one for each state in the continental US. For the regression problem of insurance purchase modeling, we tune the hyperparameters of the LightGBM via grid search and k-fold cross-validation ($k = 10$) on the training set. For all simulations we select a Tweedie objective for extremely unbalanced zero-inflated distributions as in Zhou et al. (2020). Its gradient is calculated as the log likelihood of a Tweedie distribution (dispersion parameter equivalent to 1.7). For details on the implementation in this specific framework please refer to Veigel et al. (2022). We estimate the Tweedie distribution dispersion parameter by fitting a Tweedie distribution to the target variable in the training data set before training the LightGBM. Finally, we assess the model performance using multiple performance metrics, including the Root-Mean-Square Error (RMSE), the Mean Absolute Error (MAE), and the coefficient of determination (R^2). LightGBM robustness to possible correlation and redundancy in the input data is achieved by selection of one feature within a correlated group. We apply a hierarchical grouping strategy (Text S1.2 in Supporting Information S1) to counteract this model limitation.

In the second part of our research, we use LightGBM also in classification mode to tackle the binary classification problem of predicting whether a community takes part in the CRS. LightGBM hyperparameters are trained similarly to the above regression task, but the area under the receiver operating characteristic curve was used as an objective function (Bradley, 1997). Estimated participation probabilities above 0.5 are classified as participating communities. Class imbalance in the training data set is equalized prior to training by random under-sampling in the class with more instances to contain the same amount of observations as in the underrepresented class (Guo et al., 2008).

2.3. Interpretable Machine Learning

We use LightGBM in combination with SHapley Additive exPlanations (SHAP) (S. Lundberg & Lee, 2017) with the two-fold goal of, first, predicting insurance coverage and CRS participation and, second, understanding which predictors (features) affect the model outcome and with which magnitude. SHAP values are an Interpretable ML method that quantifies the shift of the predicted value caused by each individual input feature. They represent the marginal change we observe in a conditional expectation function upon feature introduction at each step in a unique and additive manner. Since the calculation of marginal impacts depends on the features previously evaluated, the contribution should be averaged over all previously implemented features. Several methods to achieve this estimation were previously proposed, tough they are computationally demanding (S. Lundberg & Lee, 2017). Conversely, SHAP values can be estimated efficiently by exploiting the structure of tree based models such as LightGBM. They are calculated by evaluating all possible combinations of branches that end in the respective node with the Tree SHAP algorithm (S. M. Lundberg et al., 2020). In this work, we adopt SHAP values over other measures of feature importance to quantify the effect of features on model output. SHAP enhances interpretation by quantifying whether a specific predictor has a positive or negative influence on model output, along with whether specific ranges of interest emerge, when a predictor has a stronger/weaker influence on the model output. We calculate SHAP values both for the regression problem of predicting household flood insurance purchase in a census tract, where the SHAP value represents the change in insurance policies per housing unit in a census tract that can be attributed to a specific input variable, and for the classification problem of predicting communities' participation in the CRS, where the SHAP values represent the changes in participation probability attributed to a specific input variable. To mitigate the potential presence of weakly correlated features in our data set containing 400 variables, we apply a grouping strategy, in which we combine the SHAP values for thematically related features. Feature importance within a group is quantified as the weighted sum of absolute SHAP values. The combined feature importance within a group is represented even if the importance gets assigned to one variable or split across multiple variables depending on the model structure. The results and discussion are based on those variables that show the highest SHAP, either individually, or in groups.

2.4. Kruskal Wallis and Dunn's Test

As an output of the binary classification problem of CRS participation probability estimation we obtain the SHAP values representing the changes observed for the communities' participation probability in the CRS. If a community was less likely to participate because of a specific feature, such feature would be attributed with negative SHAP values more frequently. Positive SHAP values represent a shift toward a higher participation probability, and SHAP values close to zero (± 0.01) indicate no change in participation probability. To understand if the effects of a feature on communities' participation in the CRS as quantified by SHAP values are statistically significant, that is, the distributions of that feature across different SHAP groups (positive, negative, null) are statistically different, we performed a Kruskal Wallis (KW) test (Kruskal & Wallis, 1952) and post-hoc Dunn's test (Dunn, 1964). The KW test is a non-parametric test to evaluate the similarity of variances and estimate if the samples of a feature originate from the same distribution when taken from different SHAP groups. The effect size describes how strongly the groups differ from each other by comparing the feature value variances. Prior to running the KW test we apply a floodzone correction, dividing the percentage of inhabitants that are characterized by a certain feature by the floodplain area in a census tract. We further apply a post-hoc Dunn's test (Dunn, 1964) on the variables with significant results in the KW test to identify which groups exactly show significant differences with each other (Text S1.3 in Supporting Information S1).

3. Results

3.1. Flood Insurance Coverage and CRS Participation Model Performance

We here present results both for an aggregate model trained over the entire data set at the continental US scale, and models fitted individually for each state. The aggregate US LightGBM model we implement to predict flood insurance coverage at the census tract level and identify its significant drivers attains a very good training performance ($R^2 = 0.987$, RMSE = 0.004, MAE = 0.001) and satisfactory test performance ($R^2 = 0.677$, RMSE = 0.019, MAE = 0.003). For further model performance analysis broken down by individual US states please refer to Figure S1 and Table S1 in Supporting Information S1, which show that our model generalizes well for the majority of the US states, even if overfitting occurs in a small number of cases (e.g., Pennsylvania). This result demonstrates the ability of our ML approach to model overall flood insurance purchase patterns at the continental scale. Conversely, flood insurance coverage models fitted to individual states achieve varied accuracy levels, suggesting that some local effects might not be captured by the set of predictors we included in our analysis. Notably, R^2 higher than 0.6 (obtained for the US states included in Figure 1) are achieved for high-risk US states such as Georgia (test $R^2 = 0.81$) and Louisiana (test $R^2 = 0.78$). Model performance results confirm that despite the existence of several climate zones and heterogeneous socio-demographic characteristics among different states in the continental US, a relation between flood insurance coverage and local socio-economic and behavioral drivers exists.

3.2. Flood Insurance Policies Are Purchased as a Reactive Flood Adaptation Behavior

Individual state-level and aggregate US model alike indicate that flood history is the most important predictor of flood insurance coverage (Figure 1). The top-ranked *flood history* feature summarizes both flood frequency

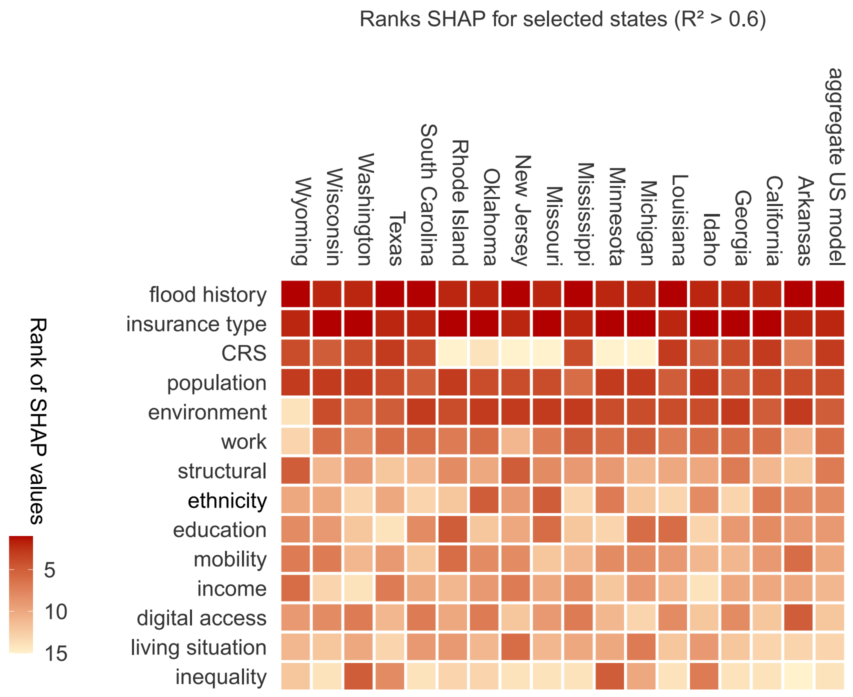

Figure 1. Feature importance for individual state models and an aggregated US model. Feature importance is quantified as the weighted sum of absolute SHapley Additive exPlanations (SHAP) values, ranked from top to bottom in descending order based on the values obtained for the aggregate US model. The graph shows feature groups including age, gender, and population density (*population*), the area within floodzones (*environment*), work-related characteristics such as the average daily working hours (*work*), and information about the building (*structural*). Only individual-state models that achieved a training performance with $R^2 > 0.6$ are selected for representation. Maximum number of affected people or annual average number of floods (combined in the feature group of *flood history*) is the most relevant determinant of flood insurance purchase for all states.

(*regular*) and severity (*severe*) (see Section 2). Both the feature ranking in Figure 1 and spatial patterns of flood severity, flood frequency, and their SHAP values in Figures 2a–2e reveal that flood insurance purchase is a reactive behavior, triggered by the occurrence of severe or frequent flooding events.

Flood insurance policies are either covering damage to buildings or to their contents. This is the second most important determinant of flood insurance coverage (*insurance type* in Figure 1). If the majority of insurance policies in a census tract covers building damage, the SHAP value for this variable increases, that is, there is a strong correlation with overall flood insurance purchase.

Spatial patterns of population density also correlate with flood insurance purchase, with lower insurance coverage associated with high population density, that is, in urban areas (*population* in Figure 1). Typically metropolitan areas experience high turnover of residents leading to lesser knowledge of the flood history in those areas. Therefore urban areas should be in the focus of resilience increasing strategies.

Conversely, other socio-economic backgrounds only marginally correlate with insurance coverage. Features related to income, structural characteristics of buildings (*structural* in Figure 1), education, and overall living situation are ranked with lower importance, as shown by the lowest weighted sum of absolute SHAP values.

3.3. Community-Level Policies Foster Proactive Adaptation

The results in Figure 1 show that the CRS stands out for positively influencing insurance coverage in a census tract and helps overcome reactive behaviors, being ranked right after flood history and flood insurance type. People are more likely to purchase a flood insurance to protect their property and belongings if the community they live in participates in the CRS. Therefore we implement a second model to classify which communities participate or do not participate in the CRS. The CRS follows a categorization system with 10 categories with

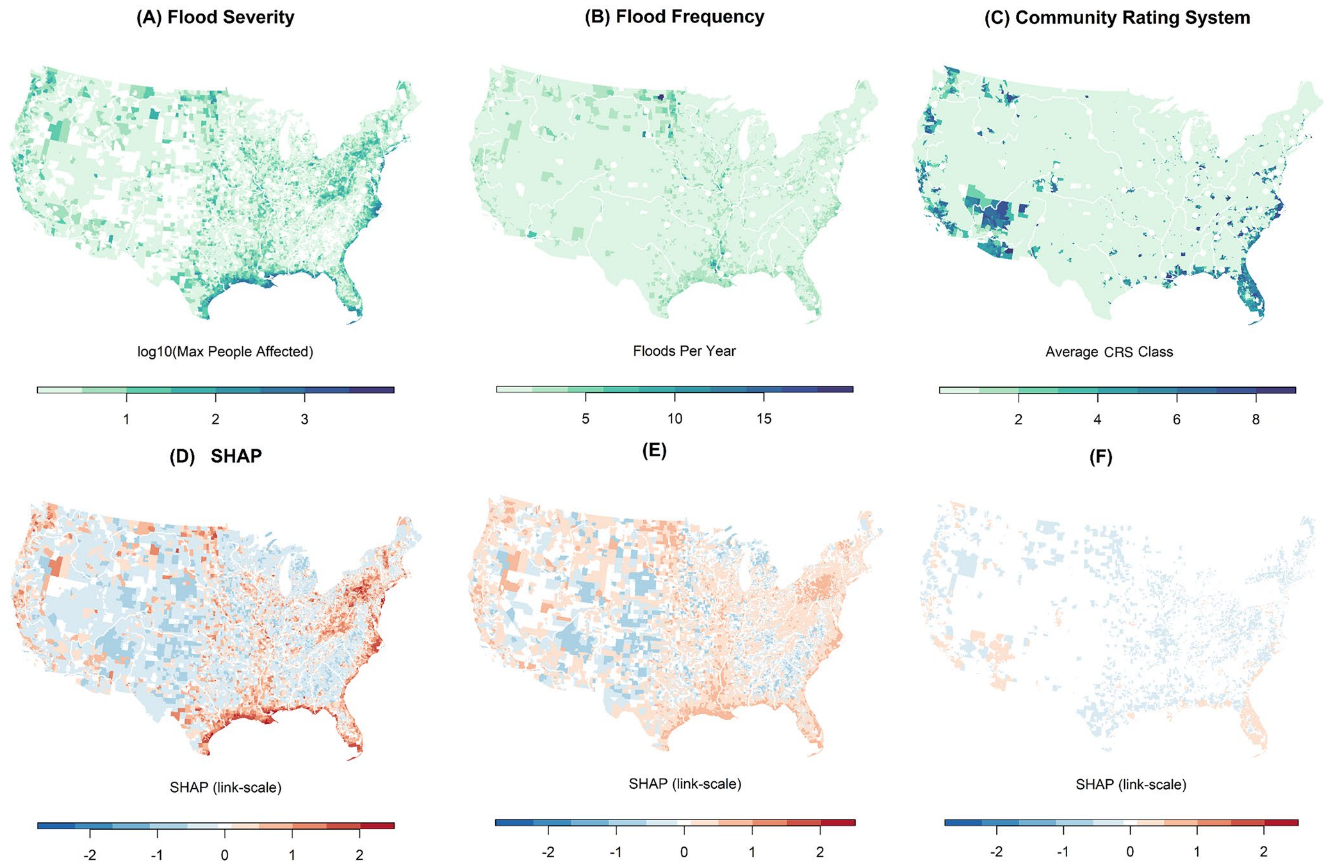

Figure 2. Feature values and related SHapley Additive exPlanations (SHAP) values for the number of people affected per flood (a–d), the average annual number of floods (b–e), and Community Rating System (CRS) average class (c–f). SHAP values are computed from modeling insurance coverage in each census tract. The maximum number of people affected by a flood (a) is skewed toward the hurricane-affected coastal communities. This historical flood severity directly translates to higher insurance coverage values (d). Flood Frequency (b) shows similar patterns with a reduced magnitude and high SHAP values centered around fluvial flood zones, such as the Lower Mississippi and the Minnesota Red River (e). The features related to flood history (flood severity and flood frequency) are the most important predictors of insurance coverage, followed by the CRS class (c) that leads to increased insurance coverage in participating communities (f).

incremental reduction in insurance premiums (Figure 2c). In the results from the insurance coverage simulation (Figure 2e) we find that the magnitude of increase in insurance purchase is similar for all positive CRS classes. Thus, the increase in insurance adoption does not depend on the specific CRS class achieved by a community, but it rather correlates with whether the communities participate at all. However, participation in the CRS is not selected as an important predictor of flood insurance in some states such as Rhode Island, Michigan, and Oklahoma (Figure 1). These states had low overall participation rates (FEMA, 2021c), which results in low absolute changes in insurance coverage attributed to the CRS. CRS premium reductions are higher within the SFHA. When we look at the changes of predicted insurance coverage for the proportion of SFHA in a census tract by CRS class we see an equal increase of flood insurance coverage for census tracts rated in all CRS classes (Figure S1 in Supporting Information S1). This indicates that the amount of subsidization within the CRS does not necessarily affect the insurance coverage.

3.4. Characteristics of CRS Communities Are Heterogeneous

Using a second model to classify the communities in the CRS, we analyze statistically significant correlations between communities' probability of participation in the CRS and their socio-economic characteristics via KW test (see Section 2). The features spatially represented in Figure 3 with their corresponding average SHAP values are a subset of grouped socio-economic variables that significantly correlate with different levels of probability of participation in the CRS. The amount of meaningful associations represented in Figure 3 shows that the CRS altogether covers a wide range of socio-economic backgrounds and that a variety of socio-economic variables

10.1029/2023EF003571

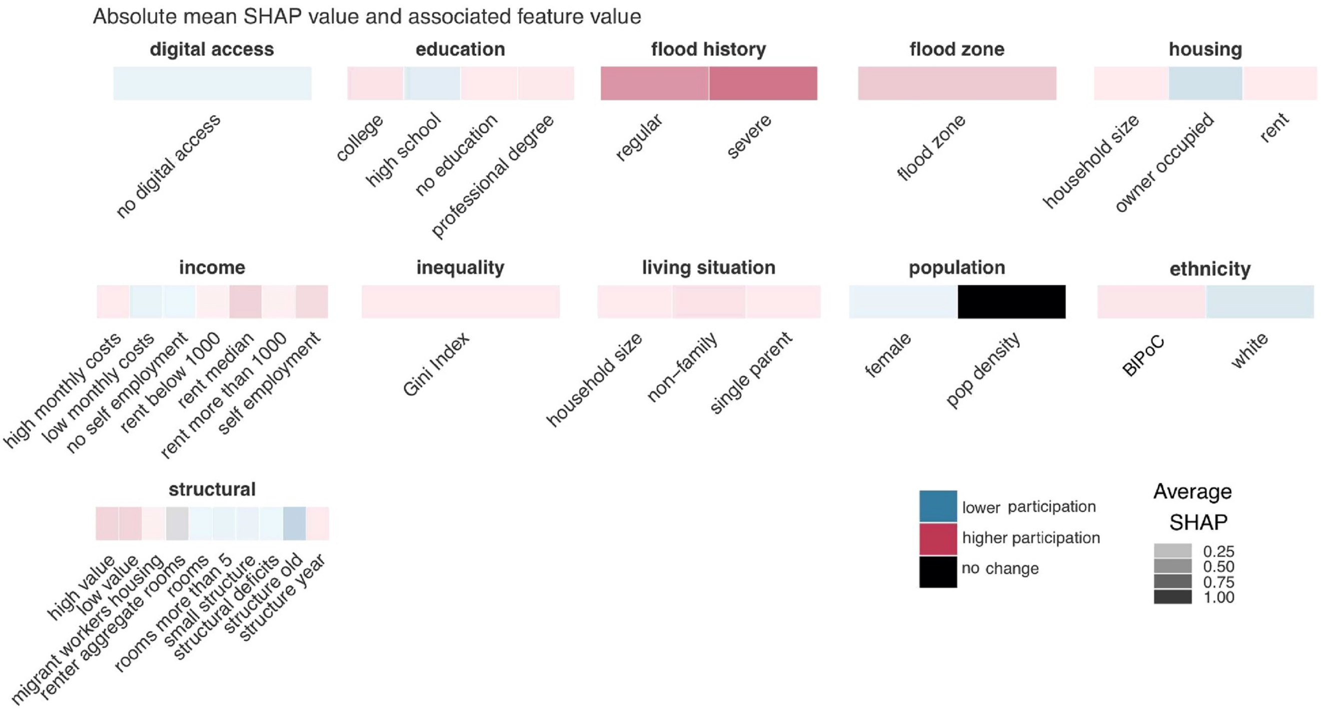

Figure 3. Socio-economic and external features grouped by type and their impact on communities' probability of participation in the Community Rating System (CRS). We train a Gradient Boosting Decision Trees (GBDT) classifier that successfully simulates and predicts CRS participation (training and testing accuracy 94%, see Figure S2 in Supporting Information S1). We analyze the output of the GBDT classifier with the most detailed aggregation level separating the features thematically. The impact of a feature on communities' CRS participation probability is categorized based on its SHAP value in lower participation (blue), higher participation (red), and no change (black). The strength of this impact is represented by the absolute average SHAP value for that feature (color opacity). *Severe* is a continuous variable that quantifies the number of people affected per flood, while *regular* refers to the annual number of floods. Both of the features are combined in the flood history variable. High feature values of each variable indicate either more severe flood characteristics or more regular flooding. If the highest mean feature values are associated with negative SHAP values (before transformation to absolute SHAP value), the feature group is appearing in blue and inversely correlates with participation in CRS (lower participation). If the highest mean feature values are associated with positive SHAP values, the feature is marked in red and directly correlates with higher CRS participation. Average absolute SHAP values close to zero indicate no change in CRS participation. Only variables that were found to significantly correlate with different levels of probability of participation in the CRS are represented (see Section 2). Feature values are divided by the area inside the floodplain in the respective census tract.

correlates with communities' participation in the CRS. However, some patterns stand out. The overall highest SHAP values (darker color opacity in Figure 3) originate from flood history as well as population density, consistently with the feature importance ranking reported in Figure 1. Flood history contributes to an increase in CRS participation probability between 50% and 100%. Conversely, the observed SHAP values for population density see not to suggest any influence on CRS participation (opaque black color in Figure 3). However, their high average SHAP (low opacity) originates from the combined effect of decreased participation rates in rural areas and participation increase in semi-urban areas (Figures S3, S6, and Table S2 in Supporting Information S1). According to the CRS classification model, census tracts with a history of severe flooding are more likely to be participating in the CRS. The insurance coverage regression model shows that they also have a higher insurance coverage. Evaluation of the CRS participation as a predictor in the insurance coverage model shows the multi-dimensional relationship between CRS class, past events, and insurance adoption (see Figure S4 in Supporting Information S1). In areas were past floods were frequent, those communities that participate in the CRS show a higher insurance coverage. Conversely, with increasing number of people affected by past floods, communities that do not participate in the CRS show a slightly higher insurance coverage. Within our model for CRS participation we can identify the socio-economic characteristics of communities that participate in the CRS. Those communities show higher insurance coverage values, despite having experienced less flooding in the past. Multiple socio-economic predictors (i.e., digital access, education, housing, income, inequality, living situation, ethnicity, and structural characteristics of buildings) are also statistically correlated with CRS participation probability. However, their impact along this correlation is marginal compared to the shift in CRS participation probability caused by the above external factors (flood history and population density).

3.5. Effects of Socio-Economic Structure on CRS Participation

Insights about the correlation of socio-economic variables with CRS participation (from the CRS classification model) emerge from the statistically significant differences in feature distributions displayed in Figure 3 and grouped by SHAP value (see also Figure S7, Tables S3, and S4 in Supporting Information S1). To further assess the magnitude, or strength, of their effect on CRS participation beyond averaged SHAP values, we further analyze their effect size (see Section 2). The features we find to have very high differences between the groups (KW effect size $f > 0.998$, corresponding to the 0.8 quantile of all significant effect sizes) are external factors. Thus, consistently with our previous analysis, the areas with past flood experience of severe or regular events are more likely to be participating in the CRS. However, also features indicating vulnerable communities (structural deficits, population density, and no digital access), describing the amount and type of rented places (high rental prices, rented properties), and employment type displayed significant effect sizes, suggesting that CRS can help overcome inequalities. Finally, both high and low property values are associated with higher participation probability, reflecting the accumulation of wealth in floodplains, the overvaluation of housing prices inside floodplains (Hino & Burke, 2021), and the increased exposure of low-income population (Wing et al., 2022).

4. Discussion

The findings from our data-driven analysis advance knowledge on human behavior in relation to flood risk with two key insights. First, previous studies have shown that vulnerable groups are systematically more exposed to floods (Knighton et al., 2020; Wing et al., 2022). Thus, there is a necessity for more vulnerable inhabitants of areas that are more exposed to floods in the US to increase their resilience to a higher degree (e.g., BIPoC residents (Knighton et al., 2021; Wing et al., 2022)).

We observe that community-level policies such as the CRS can overcome such inequality bias and foster proactive adaptation across heterogeneous socio-economic backgrounds, especially in areas with a record of experienced severe floods. We show evidence that the CRS constitutes a useful framework to encourage increased resilience for different types of communities. The observed patterns of unequal exposure imply a necessity for targeted risk communication. We find that the CRS is not targeting specific groups, but instead it covers a wide range of socio-economic characteristics. However, the CRS should improve its focus toward low income and BIPoC population segments to further balance existing inequalities. Future research could help understand more specific varying exposure to evaluate where to prioritize efforts to increase resilience via top-down and bottom-up disaster risk reduction (DRR) measures.

Second, previous studies demonstrated that individuals' ability of making informed and rational decisions to maximize their utility is hampered, particularly in high risk scenarios that occur with a low probability, for example, flooding (Conlisk, 1996). Furthermore, subsidization has been proven ineffective in improving flood insurance uptake (Kunreuther & Pauly, 2004), leaving risk communication as an alternative option to encourage informed insurance decisions (Robinson & Botzen, 2019). Our results confirm and generalize this behavioral pattern on a continental scale: insurance purchase is triggered by previous flood experience and by participating in the CRS but does not rise consistently within the classes, indicating that further subsidization does not encourage risk-averse behavior alone. We do not see a spread in insurance adoption between census tracts that expect higher reductions of insurance premiums (larger areas within the SFHA), our results suggest that the reduction of premiums is not the defining factor for higher insurance coverage in CRS communities. Rather, improved risk communication, which is part of the CRS point collection system, could make a difference in rising insurance coverage.

The unexplained portion of behaviors we model, that is, the residual in our models, might be emotionally driven by non-tangible anticipated or immediate emotions as suggested by Loewenstein and Lerner (2003). However, future research is needed to investigate these emotional components, when existing, as they cannot be described through the socio-economic variables and at the scale and level of aggregation considered in this study. The complexity of behavioral drivers is a limiting factor for model performance. The large-scale nature of this study also introduces some uncertainty with regards to collinearity in the data. In the case of localized collinearity unaccounted for by our method, the SHAP values could be proportionately reduced. To mitigate this effect we apply the predictor grouping strategy to distribute the importance between potentially correlated variables and therefore we can draw conclusions for the whole US. Future research could focus on evaluating correlations that might occur within smaller administrative units and how they relate to our findings. Further, inside the SFHA households with governmentally backed mortgages are required to purchase an insurance with the NFIP. In some cases the recipients of post-disaster aid may be required to hold an insurance for a limited amount of time. In this study we still include this small fraction of policy holders in the analysis, although their decision to acquire insurance is not self motivated. Furthermore, we do not aim to assume causality within this framework. However, our findings that areas in which households have previously been exposed to flooding are more likely to purchase insurance and that subsidization is not a major factor driving insurance purchase are in line with existing research (Cannon et al., 2020; Kunreuther & Pauly, 2004).

With the SHAP values and LightGBM we develop an interpretable ML-based analysis that minimizes biases because of the data-driven model structure without conceptualizations/hypothesis. Our results regarding the overall marginal impact of socio-economics could indicate a potential confirmation bias in conceptual or socio-hydrologic models. Due to a selection of socio-economic variables performed before model building and implementation based on prior expectation of the human-environment system, many existing models may show significant relationship that are inherent in their structure which is based on conceptualization. However, the error introduced through location redaction of NFIP insurance policies data, along with the potentially heterogeneous composition of socio-economic variables in a census tract that might be unevenly associated with the local insurance coverage after spatial aggregation at the census tract level are two limitations in this study deriving from the data processing sequence we implemented.

Since our results show that the increase of insurance purchase is similar for all CRS classes, we can conclude that even simple methods of DRR (e.g., making flood maps available to the community) can make an impact to improve individuals' motivations to undertake resilience actions. Our results suggest that the strategies and structure of the CRS has been so far successful in promoting insurance purchase throughout the US and supported people of all socio-economic backgrounds. The observation that with increasing CRS class there is no strong increase in insurance purchase could also imply that there is a limit to which communities can be motivated to take up an insurance. The defining factor of improving insurance adoption seems to be the participation in the CRS and slightly lowered insurance rates. However, the extent to which the insurance rates are lowered is not decisive for insurance purchase. These findings are further supported by the observation that income is not a significant predictor for flood insurance coverage, indicating that financial constraints on a household level do not have a predominant role in driving the decision to purchase insurance policies.

Our research shows that currently the purchase of flood insurance is mainly due to reactive behavior after severe floods. Affected people learn from experiencing a flood and improve their resilience via insurance. However, the CRS effectively motivates proactive insurance adoption and can increase community resilience in areas where extreme events have not occurred recently. We find that the CRS covers a wide range of socio-economic backgrounds in a non-discriminatory way and supports both, vulnerable communities (e.g., with low education, high BIPoC population) and less vulnerable communities (e.g., with high housing value, high median rent). Since, for instance, the BIPoC population faces the largest increase in future flood exposure (Knighton et al., 2021; Wing et al., 2022), they would need most help. Thus, the CRS should address and balance such existing inequalities, in focusing on high risk areas with an increasingly exposed and vulnerable population.

Data Availability Statement

All data used in this study is available from the U.S. Census Bureau and the Federal Emergency Management Agency (U.S. Census Bureau, 2018; FEMA, 2019, 2021a, 2021b, 2021c). The R code used to generate the results reported in this article is available at Veigel et al. (2023).

Acknowledgments

The authors would like to thank the Helmholtz Einstein International Berlin Research School in Data Science (HEIBRiDS) for supporting this project and the research environment that made it possible. We are grateful for the computing resources provided by the HPI FutureSOC Lab. Open Access funding enabled and organized by Projekt DEAL.

References

Aerts, J. C., Botzen, W. J., Clarke, K. C., Cutter, S. L., Hall, J. W., Merz, B., et al. (2018). Integrating human behaviour dynamics into flood disaster risk assessment. *Nature Climate Change*, *8*(3), 193–199. https://doi.org/10.1038/s41558-018-0085-1

Atreya, A., Ferreira, S., & Michel-Kerjan, E. (2015). What drives households to buy flood insurance? New evidence from Georgia. *Ecological Economics*, *117*, 153–161. https://doi.org/10.1016/j.ecolecon.2015.06.024

Bradley, A. P. (1997). The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, *30*(7), 1145–1159. https://doi.org/10.1016/S0031-3203(96)00142-2

Bubeck, P., Botzen, W. J., & Aerts, J. C. (2012). A review of risk perceptions and other factors that influence flood mitigation behavior. *Risk Analysis*, *32*(9), 1481–1495. https://doi.org/10.1111/j.1539-6924.2011.01783.x

Cannon, C., Gotham, K. F., Lauve-Moon, K., & Powers, B. (2020). The climate change double whammy: Flood damage and the determinants of flood insurance coverage, the case of post-Katrina New Orleans. *Climate Risk Management*, *27*(September 2019), 100210. https://doi.org/10.1016/j.crm.2019.100210

Conlisk, J. (1996). Why bounded rationality? *Journal of Economic Literature*, *34*(2), 669–700.

Cremades, R., Surminski, S., Máñez Costa, M., Hudson, P., Shrivastava, P., & Gascoigne, J. (2018). Using the adaptive cycle in climate-risk insurance to design resilient futures. *Nature Climate Change*, *8*(1), 4–7. https://doi.org/10.1038/s41558-017-0044-2

Dunn, O. J. (1964). Multiple comparisons using rank sums. *Technometrics*, *6*(3), 241–252. https://doi.org/10.1080/00401706.1964.10490181

FEMA. (2019). FIMA NFIP redacted claims—FEMA.gov [Dataset]. FEMA. Retrieved from https://www.fema.gov/about/openfema/data-sets#nfip

FEMA. (2020). FIMA NFIP redacted policies—V1—FEMA.gov [Dataset]. FEMA. Retrieved from https://www.fema.gov/about/openfema/data-sets#nfip

FEMA. (2021a). Community rating system overview and participation. Retrieved from https://www.fema.gov/fact-sheet/community-rating-system-overview-and-participation

FEMA. (2021b). FEMA flood map service center [Dataset]. FEMA. Retrieved from https://msc.fema.gov/portal/home

FEMA. (2021c). Information about the community rating system [Dataset]. FEMA. Retrieved from https://www.fema.gov/case-study/information-about-community-rating-system

Gallagher, J. (2014). Learning about an infrequent event: Evidence from flood insurance take-up in the United States. *American Economic Journal: Applied Economics*, *6*(3), 206–233. https://doi.org/10.1257/app.6.3.206

Guo, X., Yin, Y., Dong, C., Yang, G., & Zhou, G. (2008). On the class imbalance problem. In *2008 Fourth international conference on natural computation* (Vol. 4, pp. 192–201).

Hino, M., & Burke, M. (2021). The effect of information about climate risk on property values. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(17). https://doi.org/10.1073/pnas.2003374118

Horn, D. P., & Brown, J. T. (2018). *Introduction to the National Flood Insurance Program (NFIP)*. Congressional Research Service, 3–4. Retrieved from https://fas.org/sgp/crs/homesec/R44593.pdf

Hu, Z. (2020). Salience and households' flood insurance decisions. *SSRN Electronic Journal*. https://doi.org/10.2139/SSRN.3759016

Jongman, B., Winsemius, H. C., Aerts, J. C., Coughlan De Perez, E., Van Aalst, M. K., Kron, W., & Ward, P. J. (2015). Declining vulnerability to river floods and the global benefits of adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(18), E2271–E2280. https://doi.org/10.1073/PNAS.1414439112

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., et al. (2017). LightGBM: A highly efficient gradient boosting decision tree. *Advances in Neural Information Processing Systems*, *30*, 3146–3154.

Knighton, J., Buchanan, B., Guzman, C., Elliott, R., White, E., & Rahm, B. (2020). Predicting flood insurance claims with hydrologic and socioeconomic demographics via machine learning: Exploring the roles of topography, minority populations, and political dissimilarity. *Journal of Environmental Management*, *272*, 111051. https://doi.org/10.1016/j.jenvman.2020.111051

Knighton, J., Hondula, K., Sharkus, C., Guzman, C., & Elliott, R. (2021). Flood risk behaviors of United States riverine metropolitan areas are driven by local hydrology and shaped by race. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(13), 2021. https://doi.org/10.1073/pnas.2016839118

Koerth, J., Vafeidis, A. T., & Hinkel, J. (2017). Household-level coastal adaptation and its drivers: A systematic case study review. *Risk Analysis*, *37*(4), 629–646. https://doi.org/10.1111/risa.12663

Kruskal, W. H., & Wallis, W. A. (1952). Use of ranks in one-criterion variance analysis. *Journal of the American Statistical Association*, *47*(260), 583–621. https://doi.org/10.1080/01621459.1952.10483441

Kuang, D., & Liao, K. H. (2020). Learning from floods: Linking flood experience and flood resilience. *Journal of Environmental Management*, *271*(June), 111025. https://doi.org/10.1016/j.jenvman.2020.111025

Kundzewicz, Z. W., Hegger, D. L. T., Matczak, P., & Driessen, P. P. J. (2018). Flood-risk reduction: Structural measures and diverse strategies. *Proceedings of the National Academy of Sciences*, *115*(49), 12321–12325. https://doi.org/10.1073/pnas.1818227115

Kunreuther, H., & Pauly, M. (2004). Neglecting disaster: Why don't people insure against large losses? *Journal of Risk and Uncertainty*, *28*(1), 5–21. https://doi.org/10.1023/b:risk.0000009433.25126.87

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. In *Handbook of affective sciences* (pp. 619–642).

Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., et al. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, *2*(1), 56–67. https://doi.org/10.1038/s42256-019-0138-9

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in neural information processing systems, 2017-December* (pp. 4766–4775). Retrieved from https://arxiv.org/abs/1705.07874v2

McPhillips, L. E., Chang, H., Chester, M. V., Depietri, Y., Friedman, E., Grimm, N. B., et al. (2018). Defining extreme events: A cross-disciplinary review. *Earth's Future*, *6*(3), 441–455. https://doi.org/10.1002/2017EF000686

Robinson, P. J., & Botzen, W. W. (2019). Determinants of probability neglect and risk attitudes for disaster risk: An online experimental study of flood insurance demand among homeowners. *Risk Analysis*, *39*(11), 2514–2527.

Sadiq, A.-A., Tyler, J., Noonan, D. S., Norton, R. K., Cunniff, S. E., & Czajkowski, J. (2019). Review of the federal emergency management agency's community rating system program. *Natural Hazards Review*, *21*(1), 03119001. https://doi.org/10.1061/(ASCE)NH.1527-6996.0000320

Sanders, B. F., Schubert, J. E., Goodrich, K. A., Houston, D., Feldman, D. L., Basolo, V., et al. (2020). Collaborative modeling with fine-resolution data enhances flood awareness, minimizes differences in flood perception, and produces actionable flood maps. *Earth's Future*, *8*(1), 1–23. https://doi.org/10.1029/2019EF001391

U.S. Census Bureau. (2018). 2018 American Community Survey 5-year public use microdata samples [Dataset]. Census Data API. Retrieved from https://api.census.gov/data/2018/acs/

Veigel, N., Kreibich, H., & Cominola, A. (2022). A gradient boosting approach to identify behavioral and policy determinants of flood resilience in the continental US. *IFAC-PapersOnLine*, *55*(33), 85–91. (2nd IFAC Workshop on Control Methods for Water Resource Systems CMWRS 2022). https://doi.org/10.1016/j.ifacol.2022.11.014

Veigel, N., Kreibich, H., & Cominola, A. (2023). Code for interpretable machine learning reveals potential to overcome reactive flood adaptation in the continental US [Code]. *Zenodo*. https://doi.org/10.5281/zenodo.8067448

Wing, O. E., Lehman, W., Bates, P. D., Sampson, C. C., Quinn, N., Smith, A. M., et al. (2022). Inequitable patterns of US flood risk in the Anthropocene. *Nature Climate Change*, *12*(2), 156–162. https://doi.org/10.1038/S41558-021-01265-6

Zhou, H., Qian, W., & Yang, Y. (2020). Tweedie gradient boosting for extremely unbalanced zero-inflated data. *Communications in Statistics Simulation and Computation*, *51*(9), 1–23. https://doi.org/10.1080/03610918.2020.1772302

References From the Supporting Information

Hijmans, R. J. (2021). raster: Geographic data analysis and modeling [Software]. (R package version 3.5-2). Retrieved from https://CRAN.R-project.org/package=raster

Ke, G., Soukhavong, D., Lamb, J., Meng, Q., Finley, T., Wang, T., et al. (2021). lightgbm: Light gradient boosting machine [Software]. (R package version 3.2.1). Retrieved from https://CRAN.R-project.org/package=lightgbm

Liu, Y., & Just, A. (2021). Shapforxgboost: Shap plots for 'xgboost'. [Software]. (R package version 0.1.1). Retrieved from https://CRAN.R-project.org/package=SHAPforxgboost

Pebesma, E. (2018). Simple features for R: Standardized support for spatial vector data [Software]. The R Journal, *10*(1), 439–446. https://doi.org/10.32614/RJ-2018-009

R Core Team. (2021). R: A language and environment for statistical computing [Software]. Vienna, Austria. Retrieved from https://www.R-project.org/

Tomczak, M., & Tomczak, E. (2014). The need to report effect size estimates revisited. an overview of some recommended measures of effect size. *Trends in Sport Sciences*, *1*(21), 19–25.

Wickham, H. (2016). ggplot2: Elegant graphics for data analysis [Software]. Springer-Verlag. Retrieved from https://ggplot2.tidyverse.org

Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D., François, R., et al. (2019). Welcome to the tidyverse. *Journal of Open Source Software*, *4*(43), 1686. https://doi.org/10.21105/joss.01686