Title
The Racial Frontier: Biracials, Machine Learning, and the Future of Racial Group Boundaries

Author
Leslie, Gregory John

Publication Date
2022

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

The Racial Frontier:

Machine Learning, Biracials, and the Future of Racial Group Boundaries.

A thesis submitted in partial satisfaction of the

requirements for the degree Master of Science in Statistics

by

Gregory John Leslie

2022

ABSTRACT OF THE THESIS

The Racial Frontier:

Machine Learning, Biracials, and the Future of Racial Group Boundaries.

by

Gregory John Leslie

Master of Science in Statistics

University of California, Los Angeles, 2022

Professor Chad J. Hazlett, Chair

A growing thread of research uses Biracials—those who exist at the intersection of our major social cleavages (racial groups)—to reveal the current nature and future trajectory of our racial hierarchy. Specifically, researchers explore whether Minority-White Biracials (those with one White parent and one Minority parent) tend to be more similar to either Whites or to their Minority counterparts. The former circumstance would suggest a trajectory of assimilation for racial minority groups and waning intergroup prejudice, while the latter augurs enduring racial group boundaries and continued minority subjugation. Existing studies provide tremendous contributions to this genre, but are constrained in their data and methodology. In this study, I offer new data which measures Biracials by parentage (an important circumvention of endogeneity) and a machine learning approach which can use hundreds of variables at a time in order to measure how Biracials compare to their single-race counterparts. In terms of political attitudes, Black-White Biracials are more simliar to Blacks, while Asian-Whites exhibit political thinking

approximating that of single-race Whites. Latino-Whites remain ”in-between” their counterpart groups.

The thesis of Gregory John Leslie is approved.

Mark S. Handcock

Natalie Remi Masuoka

Chad J. Hazlett, Committee Chair

University of California, Los Angeles

2022

Table of Contents

List of Figures

List of Tables

1 Introduction

In recent decades, social scientists have greatly advanced our understanding of the nature of race, racial identity, and the roles they play in shaping American society. In particular, a rising scholarship creatively leverages the Biracial population to examine the effects of continuing trends in demographic change on racial group boundaries[1] (Alba, 2020; Davenport, 2018; Lee & Bean, 2012; Leslie & Sears, 2022; Masuoka, 2017). Building on theoretical assertions made nearly a century ago[2], scholars posit that comparing Biracials to their single-race counterparts on important benchmarks will reveal both the modern status and future trajectory of racial group boundaries. Specifically, scholars look to find out whether Black-White, Latino-White, and Asian-White Biracials are more similar to single-race Whites or to their single-race minority counterparts. The theoretical inference generally made is that if Minority-White Biracials are more similar to Whites it could indicate a trajectory of assimilation for racial minority groups and waning intergroup prejudice. However, if Minority-White Biracials are more similar to their minority counterparts this would augur enduring racial group boundaries and continuing minority subjugation (See also Leslie & Sears (2022) for more information on this framework and for a summary of related literature). The goal of this current study is to join previous scholars in comparing Biracials to their single-race counterparts, with a specific focus on political attitudes and machine learning.

In the past, studies have compared Biracials to single-race groups in terms of education and income (Alba, 2020; Hochschild et al., 2012; Masuoka, 2017), social networks and segregation (Alba et al., 2018; Bennett, 2011), poverty and inequality (Bratter, 2018), health outcomes (Bratter & Mason, 2016; Does et al., 2021; B. Miller et al., 2019), and political attitudes (Davenport, 2016; Davenport et al., 2022; Leslie & Sears, 2022; Masuoka, 2008). While these studies represent significant advances to the field, a goal of this study is to note and address their potential empirical limitations. First, studies tend to identify their Biracial samples via self-identification rather than by parentage. This method constrains the generalizability of findings given the plethora of research which shows that Biracial self-identification is highly unstable [3]. Moreover, Biracial self-identification is endogenous to the very

[1] Also known popularly as the racial hierarchy, or racial color lines.

[2] Robert E. Park (1928) in his musings on the "marginal man" argued that societies must undergo significant cultural and racial intermixing to defend against stagnation and achieve social and political progress. As a consequence, he—along with Stonequist (1935)—argued that in-between groups such as Biracials are the ideal group for studying the future trajectory of societies.

[3] 41.3% of Black-White and 45% of Asian-White self-identified individuals changed their identity choices after five years (Mihoko Doyle & Kao, 2007); 40.5% of Black-White and 54.1% of Asian-White self-identified individuals changed their identity choices between settings (home versus school; Harris & Sims, 2002); 40% of Black-White and 44% of Asian-White self-identified individuals changed their identity choices between 2000 and 2010 Censuses (Liebler et al., 2017)

outcomes we seek to explore[4]. To remedy this, I use both the 2020 Collaborative Multiracial Election Survey and the 2020 Nationscape datasets which allow for measurement of Biracials via parentage.

Second, existing models tend to compare Biracials to their single-race counterparts using one (or maybe a few) outcomes at a time. For example, a study might investigate whether Black-White Biracials are more similar to single-race Whites or single-race Blacks in terms of their sociodemographic profiles by comparing their mean poverty rates. While this approach works well for single, objective measures, it may be unsuitable for comparing groups in terms of multiple, subjective outcomes such as political attitudes. The first concern is that scholars must make assumptions a priori regarding which outcomes best characterize the differences between Whites and each minority racial group. For example, comparing Black-White Biracials, single-race Blacks, and single-race Whites in terms of their partisanship may seem like an ideal empirical test of the political attitudinal differences between these racial groups. However, existing research tells us that partisanship is not as salient among newer immigrant groups like Latinos and Asians as it is for Blacks (Abrajano & Alvarez, 2010), so alone it may not well capture the differences in political thinking between all racial groups. My empirical approach elides these concerns by using machine learning to agnostically identify the political attitudinal features which best explain the differences between each racial group. Moreover, ensemble modeling weights each specific feature according to its relevance and explanatory power which allows me to leverage very large datasets to provide comprehensive estimates of the differences between Biracials and their single-race counterparts.

To summarise my empirical approach, I begin by using ensemble machine learning to train three binary categorization models to predict single-race[5] class membership among three separate subsamples including minority and White single-race pairings (Blacks and Whites; Latinos and Whites; and Asians and Whites). Next—the key methodological innovation of this study—I apply these trained categorization models onto three different Biracial[6] subgroup samples: Black-Whites[7], Latino-Whites, and Asian-Whites, respectively. Each model returns the mean probabilities with which each Biracial subgroup sample would be categorized either as White or with their racial minority group, and I comport this quantity as an empirical measurement of the theoretical distances between Biracials and their single-race White and minority counterparts. If Biracials are more often categorized as minorities than

[4]If we want to know whether individuals with dual-racial heritage are systematically more similar to one single-race group or the other, identifying a Biracial sample via self-identification is problematic because it allows individuals with dual-racial heritage to select out of the "Biracial" sample which confounds our inference.

[5]"Single-race individuals" hereafter denotes individuals who have indicated having two biological parents—a mother and father—from one, shared racial group.

[6]"Biracial individuals" hereafter denote those who have indicated that their biological mother and father belong to two different (and singular) racial groups.

[7]"Black-Whites" denote individuals who have one single-race Black parent and one-single-race White parent.

as Whites, I infer that the expectations of hypodescent or the one drop rule—a social and institutional norm which decrees that one drop of Black heritage makes an individual de facto Black (hypodescent has been widely extended to non-Black minority heritage as well)—continue to ring true in the modern era.

For the purposes of this study political attitudes represent an ideal dimension on which to compare Biracials to their single-race counterparts. First, political attitudes are in large part a reflection of individuals' lived experiences and are heavily impacted by what is known as "reflected appraisals" or the racialized nature in which one is treated by society (Sims, 2016). If Biracials are systematically characterized by society as either Whites or as a member of their minority counterpart group we should expect that their political attitudes will mirror those of the group with which they have been externally assigned. Second, according to social hierarchical theories such as social dominance orientation, group consciousness, and linked fate, we should also expect a tight linkage between the degree with which each individual experiences or perceives racial discrimination against their counterpart groups and their political dispositions (A. H. Miller et al., 1981; Sidanius & Pratto, 2001; Tate, 1994). Third, political attitudes represent a diffuse and complex amalgamation of psychological dispositions measured through innumerable survey questions. Machine learning can help generalize measurement given the multiple items involved. All told, the machine learning methodology proposed offers a useful tool for scholars who wish to empirically characterize any group which is theorized to be "in-between" two poles, and for which multiple variables are necessary or helpful for characterizing.

2 Data

2.1 Datasets

This study relies primarily on two datasets, the 2020 Collaborative Multiracial Post-Election Survey[8] (CMPS) and the 2020 Nationscape.[9] The CMPS is particularly useful for this study as it boasts a nearly exhaustive set of identificational, attitudinal, and behavioral features which are related to politics and race. As a collaboration of over 200 scholars from over 100 different universities, the CMPS sampled an incredibly large pool of minority participants: $N_{CMPS} = 14,819$ Black, White, Latino, or Asian participants. Moreover, the CMPS uses both the "mark one or more" format on the racial self-identification question as well as two seldom-used items which ask participants to indicate

[8]See https://cmpsurvey.org/2020-survey/ for more information

[9]See https://www.voterstudygroup.org/nationscape for more information

the race of their biological mother and father.[10] These latter items allow for the enumeration of Biracials by parentage. Lastly, the CMPS contains an incredibly vast set of political attitudinal features ($X_{CMPS} > 400$) that can inform the categorical predictions of each machine learning algorithm.

Next, the Nationscape data is also unique in that it includes items measuring the race(s) of participants and their parents.[11] The Nationscape data, however, is much more limited than the CMPS in terms of the number of features which were asked consistently across each wave.[12] However, the sample size is much larger ($N_{Nat} = 142,065$).

2.2 Participants

In each dataset, I identify seven subsets belonging to different racial group classes. These classes are comprised of three Biracial subgroups, 1) Black-White Biracials, 2) Latino-White Biracials, and 3) Asian-White Biracials, and four single-race classes, 4) single-race Whites, 5) single-race Blacks, 6) single-race Latinos, and 7) single-race Asians. Individuals included in the Biracial subgroup classes are identified as participants who marked only one box to indicate the race of their mother and only one box for the race of their father, and who indicated differing races between them and corresponding to the subgroup class in which they are enumerated. For example, Black-Whites Biracials are those who indicated only Black for their father and only White for their mother, or vice-versa. Single-race individuals indicated only one, shared race for both their mother and father.

	Sample Sizes	
	CMPS	**Nationcape**
Single-Race Classes		
Whites (n_W)	2,888	79,047
Blacks (n_B)	3,277	10,218
Latinos (n_L)	2,309	8,242
Asians (n_A)	2,879	5,110
Biracial Classes		
Black-Whites (n_{BW})	189	1,092
Latino-Whites (n_{LW})	493	3,132
Asian-Whites (n_{AW})	204	777

Table 1: Sample sizes for each dataset.

[10] Participants were asked in three separate items, "What do you consider [your; your mother's; your fathers] race or ethnicity? Mark one or more boxes." Allowable responses include: White; Hispanic or Latino; Black or African American; American Indian/Native American; Arab, Middle Eastern or North African; Native Hawaiian; Not Hawaiian, but other Pacific Islander.

[11] Item asks participants to, "Please indicate your race and the race of your biological parents if known. Mark one or more boxes." participants were presented with three columns of boxes, and asked to indicate among the following categories to describe their own and their parents' races: White, not-Hispanic; Black or African American; Hispanic or Latino; Asian or Asian American; American Indian or Alaska Native; Native Hawaiian or Pacific Islander; Middle Eastern or North African; Other; Not sure/Unknown.

[12] Nationscape included these items for a total of 22 weeks (though intermittently discontinued between waves 45 and 48, and waves 51 and 56).

It is important to note that class imbalance as a result of the disparate single-race subsample sizes in my data poses the risk of providing undue advantage in the classification accuracy of some machine learning algorithms compared to others. For example, more observations of single-race Whites might aid an algorithm's ability to learn what best predicts Whites while racial group classes with smaller sample sizes might be less accurately predicted. To account for class imbalance, I rely on random undersampling so that there are equal amounts of participants among each racial group class (Drummond & Holte, 2003; Weiss, 2004). For the CMPS, the undersampled size is 2,309 participants for each single-race class (undersampled to match the size of single-race Latinos). For Nationscape, the undersampled size is 2,760 to match the sample of single-race Asians).

2.3 Feature Set Selection

I include each and every feature available in my datasets which are related to political attitudes (e.g. policy preferences, public opinion, racialized attitudes, non-racialized attitudes, social attitudes, feeling thermometers, candidate evaluations, etc.). However, each participant must have an identical set of features, so I exclude attitudinal items which were not asked to each and every participant and with the exact same wording. This excludes items which were fielded to only certain participants via embedded experiments, items that imputed responses to previous questions, and in the case of Nationscape I only include attitudinal items which were available on each of the waves where the question about participants' parents' races were included. Categorical features are binarized and continuous features are rescaled to take values either or in between 0 and 1. Observations are also randomly shuffled.

3 Method

3.1 Superlearner

The specific ensemble machine learning classification model I use is superlearner (Polley et al., 2011; Samii et al., 2016; Van der Laan et al., 2007). Superlearner operates by training multiple different candidate machine learning algorithms and then applying a specific weight to the output of each model depending on their ability to accurately predict each binary classification. Since it is not typically possible to know beforehand whether or not a given candidate algorithm will be more accurate than another, combining the models helps to elide issues of model selection and has been consistently demonstrated to perform either as good or better than any single algorithm at predicting out-of-sample outcomes (Van Der Laan & Dudoit, 2003).

In description of the superlearner process, the first step is to randomly split the total dataset into V cross-validation folds, and for this study I use 10 folds. Multiple cross-validation folds are important here to account for potential overfitting concerns. Accordingly, the sample size of each fold which comprises the training set is $N - (N/V)$. The test set or hold-out set is $N_v = N/V$. Also note that the total ensemble model is comprised of $c = 1, ...C$ candidate algorithms. Specifically, I use five candidate algorithms in my ensemble model which are 1) logistic regression, 2) Lasso, 3) random forest, 4) support vector machines, and 5) neural networks.

After splitting the total dataset into V cross-validation folds, the next step is to fit C models on cross-validation folds 2 through V, therefore excluding the first fold which is our first-round testing set. The estimator function for a given candidate algorithm and a given cross-validation fold is described as $\hat{g}_j^{c,v}(\cdot)$. Next, we calculate predictions for the testing set ($v = 1$) and estimate predictive accuracy using the mean squared error (MSE) for each candidate algorithm. This process is then repeated iteratively (V−1 times), but each time using a different cross-validation fold as the testing set in the order of $v_i = 2, 3, \ldots, V$.

To calculate the MSE for a given candidate learner c over the course of the iterated cross-validation fold and testing set pairs, we use:

$$\widehat{MSE}(c) = \frac{1}{V}\sum_{v=1}^{V}\Big[\frac{1}{n_v}\sum_{i:V_i=v}\{Y_i - \hat{g}_j^{c,v}(X_i|\mathbb{P}_n^{-v})\}^2\Big]$$

In the above notation, the bracketed middle term $\{Y_i - \hat{g}_j^{c,v}(X_i|\mathbb{P}_n^{-v})\}$ denotes the MSE for a single candidate algorithm c and crossfold. Specifically, Y_i denotes the true classification outcome as defined in the testing set, and $\hat{g}_j^{c,v}(X_i|\mathbb{P}_n^{-v})$ denotes the predicted outcome of candidate learner c given feature set $(X_i|\mathbb{P}_n^{-v})$. X_i here indicates the feature set matrix, and $\mathbb{P}_n^{-v}$ indicates the training data observations in which $V_i \neq v$. The outer sum simply averages across crossfolds.

Next, moving on from the discrete or single-learner estimation, the following works to minimize the MSE weighted average of all candidate learning algorithms, and does so by calculating ensemble weights as follows:

$$(w_j^{1*}, ..., w_j^{C*}) = argmin_{(w_j^1,...,w_j^C)}\frac{1}{N}\sum_{i=1}^{N}\Big[Y_i - \sum_{c=1}^{C} w_j^c \hat{g}_j^{c,v(i)}(X_i|\mathbb{P}_n^{-v})\Big]^2$$

$$\text{subject to} \sum_{c=1}^{C} w_j^c = 1 \text{ and } w_j^c \geq 0 \text{ for all } c$$

Finally, we apply weights to the predicted outputs of each candidate learner based on their performance (MSE), such that those weights operate as coefficients in a regression estimating out-of-sample predictions.

3.2 Logistic Regression

The first candidate algorithm[13] used in the ensemble model is logistic regression. Logistic regression is a popular model used to help determine the probability that a given binary outcome will occur, $(0 \leq P(y_i = 1) \leq 1)$ given a specific feature set X_i. Y denotes a vector of 0 or 1 outputs while X exists as a data matrix of input features where $X \in \mathbb{R}^{n \times p}$ and is an n by p matrix in which n is the number of observations in the data and p is the number of features. In linear regression, the probability that a given output value might occur can be calculated as a function of the features set X and a vector of coefficients β, as given by:

$$\pi_i = X_i \beta$$

π_i here represents the probability. However, the predicted outcome from linear regression can at times span outside of the 0 and 1 probability bounds. Therefore, logistic regression instead takes a functional form based on the Bernoulli distribution, $Y_i \sim Bernoulli(\pi_i)$, and employs the sigmoid function to calculate loss:

$$\pi_i = \frac{1}{1 + e^{-X_i \beta}}$$

3.3 Lasso

The second candidate algorithm is Lasso (Least Absolute Shrinkage and Selection Operator). Given that my data matrix includes a high number of features, Lasso is useful because it applies a shrinkage parameter λ which will penalize the impact of variables and exclude those which are not helpful for predicting Y. This process is helpful for accounting for over-fitting. Lasso uses the L_1 norm which uses the absolute value of magnitude of coefficients within the loss function which is:

[13] The following model descriptions are meant to provide only a brief familiarity with each model, and are by no means comprehensive.

$$\sum_{i=1}^{n}(Y_i - \sum_{j=i}^{p} X_{ij}\beta_j)^2 + \lambda \sum_{j=1}^{p} \beta_j$$

3.4 Random Forest

Next, I use Random Forest which has become an increasingly popular tool in social science for data problems involving large amounts of covariates. Random forest is based on the decision tree framework. Decision trees begin with a root node (a sample [or subsample] of observations) which undergoes a process called splitting to create two or more decision nodes. Splitting occurs according to the feature values which explain the most variation among the original node. This process is iterated repeatedly until each branch reaches its leaf or terminal node in which splitting will not improve the homogeneity of the terminal group. However, single tree models may be highly prone to overfitting. As a consequence, Random forest, which is considered an ensemble model in its own right, expands on the decision tree process by repeating itself through bootstrapping. The outcomes of each bootstrapped decision tree are then averaged via a process known as bagging to calculate the final result.

3.5 Support Vector Machines

Support Vector Machines (SVM) are another useful algorithm for predicting binary classification from a given set of outcomes and features. SVM operates by separating each datapoint on a multi-dimensional hyperplane. The hyperplane that represents the largest separation between each binary outcome will be used to classify them. Each hyperplane is chosen such that the distance between the nearest data point on each side is maximized. In a high-dimensional space, however, linear separators might not be very successful at discerning between different data classes. Non-linear kernel transformations are useful in these cases since they can apply a more flexible boundary to cleave data class boundaries. For this study I rely on the RBF kernel: (Hsu et al., 2003):

$$K(x_i, x_j) = exp\left(-\frac{||x_i - x_j||^2}{2\sigma^2}\right)$$

3.6 Neural Networks

Next, I use neural networks. Neural networks were originally designed with the goal of approximating the relationship between neurons in the nervous system. A neural network consists of three layers, 1) an input layer, 2) a hidden layer, and 3) an output layer. Each node is connected to one another across each layer by weighted edges, where higher weights assigned to each path are associated with more importance in terms of accurately predicting the final outcomes.

Neural networks operates in that the weights corresponding to each edge connecting each node are optimized during the training period. At each neuron, the weighted sum of the inputs is calculated via:

$$Y = \sum (weight * input) + bias$$

3.7 Empirical Strategy

My empirical approach begins by training three separate ensemble superlearner models to predict binary racial group classification among three different single-race class pairings. Specifically, the first model is trained on a training set subsample including only single-race Blacks and single-race Whites $\hat{g}_{B+W}(\cdot)$, the second includes only single-race Latinos and single-race Whites $\hat{g}_{L+W}(\cdot)$, and the third includes only single-race Asians and single-race Whites $\hat{g}_{A+W}(\cdot)$. In each model, $y_i = 0$ denotes when a participant belongs to the non-White class, and $y_i = 1$ denotes when a participant belongs to the single-race White class. Within each training set there are equal numbers of Whites and non-Whites.

Next, I take these three trained superlearner models and apply them to Minority-White Biracial subgroup data. Specifically, I first apply the model trained to predict single-race Black or single-race White class membership to a feature set matrix of Black-White Biracials: $\hat{g}_{B+W}(X_{BW})$. I then apply the model trained to predict single-race Latino or single-race White class membership onto a subsample of Latino-White Biracials $\hat{g}_{L+W}(X_{LW})$, and third I use the model trained to predict single-race Asian or single-race White class membership onto a subsample of Asian-White Biracials $\hat{g}_{A+W}(X_{AW})$.

My quantity of interest is the mean probability with which each Biracial subgroup is categorized as single-race White. If a Biracial subgroup has a mean probability for White classification higher than 0.5, I infer that they are more similar to Whites than to their single-race minority group counterparts in terms of their political attitudes.

3.7.1 Assumptions and Potential Bias

The approach outlined above makes several assumptions which are important to address. First, it is important to acknowledge that this process is not purporting to conclusively grasp the essence of what it means to be a human being or a member of any racial group in general. Rather, this process endeavors merely to capture in some form the underlying differences between participants who have identified with certain racial groups through certain survey items, and based only on the specific features available in each model.

Moreover, it is important to acknowledge that unknown bias may affect my results. Most presciently, my empirical strategy may be biased if there are substantially more features which are informative for predicting classification with one racial group as compared to another. Case in point, it is important to keep in mind that social science surveys typically include more items which are central to the experience of Blacks than Asians. Therefore, it is possible that machine learning algorithms might be more successful at discerning between Black and White class membership than say, Asians and Whites.

3.7.2 Hypotheses

Hypothesis 1: The following analysis assesses two primary questions. First, I explore whether certain Biracial subgroups are more similar to single-race Whites or to their single-race minority counterparts. If a Biracial subgroup is more likely on average to be classified as single-race White than with their single-race minority counterpart group, then we would infer that the color lines surrounding that Biracial subgroup are blurring and that the corresponding minority group is headed toward equity with Whites. However, if a Biracial subgroup is more likely to be classified as single-race minority, this would signal the enduring power of color lines which relegate the minority group to a subjugated status.

My explicit expectations are borne from the theory of Black exceptionalism which argues that the color line separating Blacks is the most salient and enduring racial group boundary, capturing recent minorities such as Latinos and Asians only temporarily before they assimilate with Whites (Leslie & Sears, 2022; Sears & Savalei, 2006). Black-White Biracials should be more likely to be categorized as single-race Black, while Latino-White Biracials and Asian-White Biracials should be more likely to be categorzied as single-race White than with their single-race minority counterparts.

Specifically, I expect the following:

- H1a: $P(White|X_{BW}) < 0.5$; Black-Whites will be more similar to their minority racial group than to Whites.
- H1b: $P(White|X_{LW}) \geq 0.5$; Latino-Whites will be more similar to Whites than to their minority racial group.

- H1c: $P(White|X_{AW}) \geq 0.5$; Asian-Whites will be more similar to Whites than to their minority racial group.

Hypothesis 2: Second, I explore the relative positioning of each racial minority group vis-à-vis Whites to assess the differential degree to which some minority groups show signs of political assimilation with Whites than do others. For example, the racial hierarchy is currently characterized as having Whites on top, followed by Asians, Latinos, and then Blacks on the bottom, respectively. This characterization of the hierarchy suggests that Asians would be the most similar to Whites, followed by Latinos, and then Blacks. To test these assumptions, I compete the relative distancing between Whites and each minority group against one another to determine their theoretical placement on a stratified racial hierarchy.

Compared to Whites, I expect the following:

- H2a: $P(White|X_{BW}) < P(White|X_{LW})$; Blacks are more distanced from Whites than Latinos
- H2b: $P(White|X_{BW}) < P(White|X_{AW})$; Blacks are more distanced from Whites than Asians
- H2c: $P(White|X_{LW}) < P(White|X_{AW})$; Latinos are more distanced from Whites than Asians

4 Results

4.1 Assessing Model Efficacy

First, I report findings which are important for gauging whether or not the ensemble models were successful at their task of accurately predicting members of one racial group class or another. Below, I report their accuracy at predicting single-race outcomes. Specifically, I train the three ensemble models $\hat{g}_{B+W}(\cdot)$, $\hat{g}_{L+W}(\cdot)$, and $\hat{g}_{A+W}(\cdot)$ using a training set containing 80% of my total data (randomly withheld, class-balanced and shuffled). I then—in this model assessment scenario—apply these trained ensemble models to subsamples of single-race groups which make up the remaining 20% randomly withheld testing set.

Table 2 below presents an assessment of each superlearner models' ability to accurately predict racial group class among each single-race minority and single-race White racial group pair. 'Risk' here refers to the mean squared error for each candidate machine learning algorithm as it was trained to predict racial group class. 'Coef' denotes the weight applied to the outputs of these models and which increases according to the degree with which each specific algorithm accurately predicted racial group class.

"Ensemble Accuracy" presents the proportion of observations from each single-race minority/single-race White subset pair which were accurately predicted by the ensemble model.

Overall, the ensemble models were very successful as accurately classifying each single-race class outcome. Across all models, K-SVM is the most accurate individual candidate algorithm, and so garners the largest coefficient which emphasizes its predicted outputs within the ensemble models. In general, random forest provided the second most efficient predictions. A second notable trend is that the models trained using the CMPS dataset are more accurate than those using the Nationscape dataset. While Nationscape has a much larger number of observations, the CMPS feature set, X_{CMPS}, is vastly larger which allows the highly flexible machine learning algorithms more information with which to identify and differentiate each observation according to their known racial group class. Moreover, the importance of flexible machine learning models can also be seen in their high propensity to more accurately predict outcomes than logistic regression which has been the standard approach for most existing studies. Across both the CMPS and Nationscape models, logit tends to be the least accurate in their predictions, which suggests that models which are better able to account for complex and non-linear functional forms are ultimately better at discerning between racial groups.

Table 2 also provides some information which may be of substantive importance regarding the theoretical distance between racial groups or the nature of racial group boundaries. Specifically, both the CMPS and Nationscape models have more success at accurately predicting racial group class in the models containing only single-race Black and single-race White observations, $\hat{g}_{B+W}(\cdot)$, as compared to the other two single-race White/single-race minority paired models. This may also be evidence that single-race Blacks are more different or distanced from single-race Whites in terms of their political attitudes than are other single-race minority groups.

	ML Algorithm	Subset: **Blacks and Whites** *Risk*	*Coef.*	Subset: **Latinos and Whites** *Risk*	*Coef.*	Subset: **Asians and Whites** *Risk*	*Coef.*
CMPS	Logit	0.021	0.017	0.023	0.003	0.014	0.044
	LASSO	0.009	0.000	0.016	0.122	0.010	0.164
	Random Forest	0.013	0.101	0.021	0.134	0.014	0.135
	K-SVM	0.008	0.857	0.015	0.721	0.009	0.613
	NNET	0.013	0.025	0.024	0.021	0.014	0.045
	Ensemble Accuracy:	**0.951**		**0.886**		**0.922**	
Nationscape	Logit	0.026	0.045	0.043	0.302	0.040	0.041
	LASSO	0.026	0.000	0.043	0.000	0.039	0.000
	Random Forest	0.026	0.245	0.044	0.282	0.039	0.280
	K-SVM	0.024	0.682	0.043	0.345	0.037	0.679
	NNET	0.030	0.028	0.049	0.072	0.049	0.000
	Ensemble Accuracy:	**0.832**		**0.737**		**0.797**	

Table 2: Assessing accuracy of ensemble machine learning models at predicting single-race classes.

4.2 Main Results: Are Biracials More Similar to Minorities or to Whites?

Figure 1 below presents the results of the trained superlearner ensemble machine learning models as applied to Biracial subsamples. The y-axis of Figure 1 denotes the Biracial subgroup included in each model, and the x-axis denotes the mean probability (error bars denote 95% confidence intervals) of being categorized as single-race White. Values closer to the right side of the x-axis denote higher probabilities that members of the given Biracial subgroups would be classified as single-race White, and values trending toward the left-side indicate decreasing probabilities of being classified as single-race White (increasing probabilities of being classified with the corresponding single-race Minority group). Values at the halfway mark on the x-axis indicate an equal probability on average of each observation being categorized either as single-race White or with their single-race minority group.

My first set of hypotheses predict that Black-White Biracials will be more often categorized as single-race Black than as single-race White, while Latino-White Biracials and Asian-White Biracials will be more often categorized as single-race White than with their single-race minority counterpart groups (i.e. single-race Latino and single-race Asian). Such results would signify that Black-White Biracials are the only of these three Minority-White Biracial subgroups which are more similar to their minority counterparts than to Whites, which would augur a future of blurred racial group boundaries between Whites and non-Black minorities.

The results largely confirm my expectations, but with some caveats in terms of Latino-Whites. First, Black-Whites are indeed much more likely to be categorized as single-race Black than as single-race White (support for H1a). This result substantiates that the descendants of Black-White intermixing remain much more similar to their minority counterparts than to Whites, as expected by the one-drop rule (Davis, 1991). Black-Whites on average have a probability of only 15.5% of being categorized as White in the CMPS, and 30.2% in Nationscape. Oppositely, Asian-Whites are the least likely to be categorized with their single-race minority counterparts (i.e single-race Asian), and are in fact highly likely to be classified as single-race White (support for H1c). Specifically, Asian-Whites have a high, 69.2% probability of being classified as single-race White in the CMPS model, and a 72.8% probability in the Nationscape model. In terms of Latino-Whites, the CMPS and Nationscape offer conflicting results. In the CMPS, Latino-Whites are more likely to be categorized as single-race Latinos than as White (33.0% probability of single-race White categorization), but are more likely to be categorized as single-race White in the Nationscape model (59.0%). While Black-Whites are more similar to their single-race minority group and Asian-Whites are more similar to Whites, the results for Latino-Whites are inconclusive (inconclusive for H1b).

The second set of hypotheses predicted that the relative similarity between each Minority-White subgroup and single-race Whites would follow the pattern scholars generally use to characterize the modern racial hierarchy (Sidanius & Pratto, 2001): Black-Whites should be the most dissimilar from Whites, followed by Latino-Whites and Asian-Whites, respectively. Across both datasets, I observe strong evidence that this is indeed the case. While Latino-Whites in the CMPS are shown to be more similar to single-race Latinos than to Whites, and in Nationscape are shown to be more similar to Whites, in both datasets the pattern occurs such that the mean probability of being categorized as White is the smallest for Black-Whites, the second smallest for Latino-Whites, and the highest for Asian-White Biracials, such that $P(White|X_{BW}) < P(White|X_{LW}) < P(White|X_{AW})$.

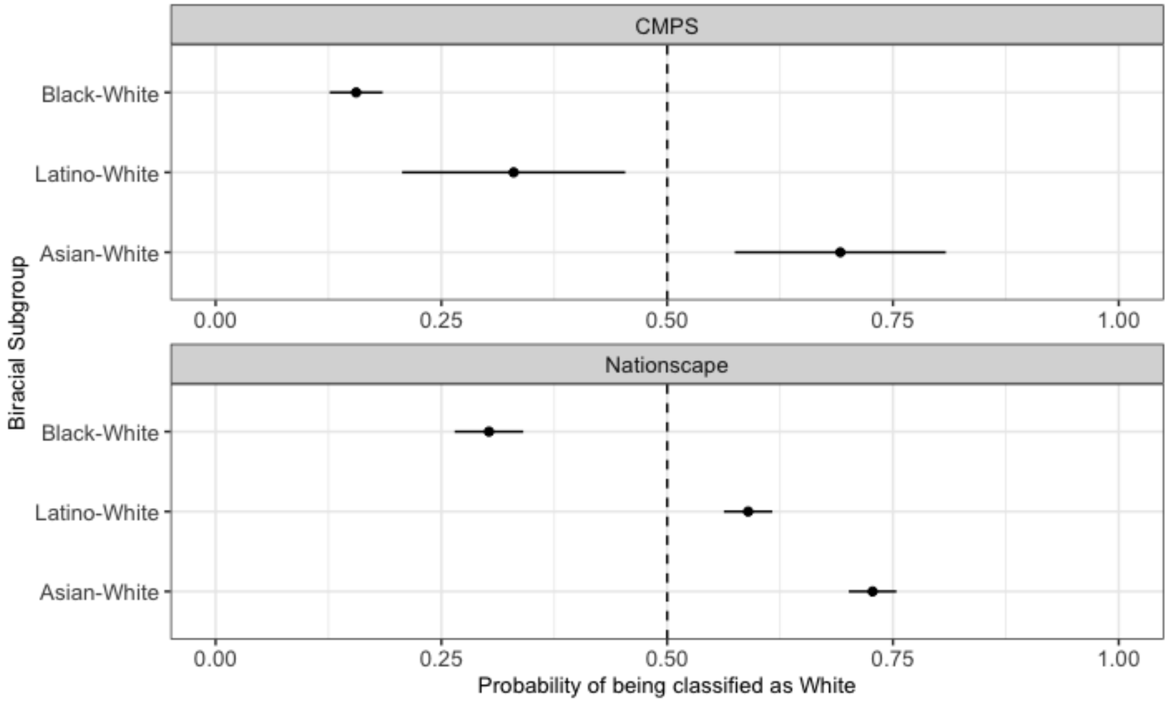

Figure 1: Ensemble machine learning results estimating whether minority-White Biracials are more similar to single-race Whites to to their single-race minority counterparts.

4.3 Feature Importance

To provide insight into which features were most important for differentiating between each racial group class, I explore feature importance. As noted above, it has become common practice when comparing racial groups to do so using one feature at a time. However, this method may be highly susceptible to selection issues such as picking and choosing the features which best support one's theory. Contrarily, a benefit of this machine learning analysis is that I can transparently display which features are most important. To determine feature importance, I rely on Lasso. Lasso uses the $L1$ regularization which applies a penalty equal to the absolute value of the magnitude of each features' coefficient. Essentially, features which are not meaningful predictors of racial group class are automatically excluded. Lasso was not the most accurate candidate model at predicting outcomes, however, its convenience as a feature selection tool may allow for ease of interpretation for readers and provide helpful insight into which features are considered most important for distinguishing between one racial group and another.

Table 3 below presents the top 10 most important features for distinguishing between single-race Whites and each single-race minority group class as indicated in the column headers. Once again, all feature values are rescaled to range between 0 and 1, and the binary outcome of $y_i = 1$ indicates categorization as single-race White, while $y_i = 0$ denotes categorization as single-race minority. Features are ordered in terms of the absolute value of their coefficient, ("$Coef$"), with higher values indicating higher correlation with the probability that an observation would be categorized as single-race White.

The first general pattern evident in Table 3, as might be expected, is that explicitly racialized features tend to garner the most polarized responses from each participant, which makes them the most powerful predictors of racial group class. However, far more interesting are the many important features which are not explicitly racialized or which are not linked to the specific single-race pairs included in the model. For example, one highly surprising result is that among all racial group pairs in the CMPS data, participants' beliefs about how individuals who are Middle Eastern/North African (MENA) should be racially categorized weigh heavily on the total ensemble models. White participants show a great affinity for inclusion in that perceiving MENA individuals as White is one of the strongest predictors of White classification.

Also interesting are that non-Black participants' attitudes toward Blacks often weigh heavily on their racial classification. For example, in the CMPS dataset the belief that Blacks are violent is the fifth overall most important feature for predicting Latino racial group membership (as opposed to classification as White). This is quite surprising given that the feature set contained over 400 different variables, dozens of which ask specifically about Latino-interested attitudes. Similarly, in the Nationscape data participants who are more favorable of Blacks and who hold less racial resentment toward Blacks are more likely to be classified as White rather than as Latino. However, perceiving higher rates of racial discrimination against Blacks is a strong predictor of being categorized as Latino. And, for the model predicting Asian or White classification (in Nationscape) support for intimacy between Blacks and Whites is an important predictor of classification as Asian.

Finally, non-racialized features also have a large impact on racial classifications. Gendered attitudes in particular tend to be highly important. In the CMPS, supporting the idea that half of elected officials should be women is a major predictor of being categorized as White rather than as Latino or as Asian. In the Nationscape models, perceiving that a lot of discrimination exists again men in this country is positively associated with identification as White rather than as minority (Black or Latino), and both Latinos and Asians are more likely to agree with the sentiment that women are less capable of thinking logically than men are (as compared to Whites).

	CMPS					
	Blacks and Whites		**Latinos and Whites**		**Asians and Whites**	
Order of Importance	*Feature Name*	*Coef.*	*Feature Name*	*Coef.*	*Feature Name*	*Coef.*
1st	Racial Context Preference: Blacks	-2.23	Linked Fate: Latinos	-2.96	Linked Fate: Asians	-5.16
2nd	Linked Fate: Blacks	-2.00	Racial Context Preference: Latinos	-2.09	Linked Fate: Whites	3.25
3rd	Linked Fate: Whites	1.90	Linked Fate: Whites	2.06	Racial Context Preference: Asians	-2.19
4th	Support: Police See Me As Thug	-1.87	Support: Latinos are Threatening	1.68	Support: Asians are Threatening	1.46
5th	Support: MENA are White	1.27	Support: Blacks are Violent	-1.05	Support: MENA are White	1.45
6th	Support: Whites are Underrepresented	1.20	Support: Police Treat Me Worse	-0.81	Support: My Network Works Together	1.10
7th	Support: MENA are Black	-1.19	Support: Whites are Threatening	-0.80	Support: Discrimination Affected Life	-1.02
8th	Racial Context Preference: Whites	1.00	Support: MENA are White	0.77	Support: Whites are Threatening	-0.97
9th	Racial Identity Importance	-1.00	Racial Context Preference: Whites	0.72	Support: More Women Elected Reps.	0.93
10th	Support: God Rewards Faithful	-0.95	Support: More Women Elected Reps.	0.64	Support: Preserve White Heritage	0.87
			Nationscape			
1st	Democrat (PID7)	-2.55	Favorability: Latinos	-2.55	Favorability: Asians	-2.83
2nd	Discrimination Against Blacks	-2.01	Favorability: Whites	1.45	Discrimination Against Asians	-1.99
3rd	Discrimination Against Men	-1.83	Discrimination Against Whites	1.34	Support: Women are Less Logical	-1.76
4th	Discrimination Against Whites	1.71	Support: Women Are Less Logical	-0.99	Discrimination Against Whites	1.68
5th	Racial Resentment (Index)	1.55	Discrimination Against Men	-0.83	Favorability: Whites	1.54
6th	Favorability: Blacks	-1.10	Favorability: Blacks	0.83	Political Interest	1.20
7th	Ideology (Liberal)	1.04	Favorability: Republicans	-0.74	Support: Own a Gun	1.11
8th	Favorability: Police	1.02	Democrat (PID7)	-0.71	Support: Blacks and Whites Dating	-1.09
9th	Favorability: Whites	0.98	Racial Resentment: Irish, Italians, etc.	-0.68	Democrat (PID7)	-0.83
10th	Favorability: Obama	-0.75	Discrimination Against Blacks	-0.58	Favorability: Latinos	0.75

Table 3: Feature Importance: Top 10 most important features given Lasso feature selection.

5 Discussion

My theoretical objective has been to examine the current nature of racial group boundaries and to make inferences about how they may change in the future. I do so by comparing the political attitudes of Biracials to those of their single-race counterparts. I ask the questions: what are Biracials like in America who have one White parent and one minority parent? Do they tend to think about political issues in a manner similar to how Whites view politics, or do they more so resemble minorities? If Biracials are the canaries in the coal mine of intergroup assimilation, evidence that Minority-White Biracials have overlapping political goals with Whites might foreshadow a future where color lines are less salient tomorrow than they are today. On the other hand, if Minority-White Biracials continue to both view themselves and be treated by others as minorities, the social cleavages associated with subjugation and discrimination may be likely to endure.

The first major contribution of this study is its base findings. With more than 400 different features related to political attitudes and two very large, nationally-representative datasets at my disposal, I demonstrate that Black-White Biracials are vastly more similar to single-race Blacks in their politics than they are to single-race Whites. Oppositely, Asian-White Biracials are much more similar to single-race Whites than they are to single-race Asians. Latino-Whites, however, lay somewhere in between their White and minority counterparts. In the CMPS dataset, Latino-Whites are predicted to be more similar to single-race Latinos than to Whites (though more similar to Whites than are Black-Whites). However, in the Nationscape data Latino-Whites are more proximal to their single-race White counterparts.

These findings serve as an important update to a long legacy of social theories characterizing the nature and future of intergroup relations and the racial hierarchy. For example, frameworks such as social dominance theory (Sidanius & Pratto, 1999), group position theory (Blumer, 1958; Bobo, 1999, 2004), and classic assimilation (Gordon, 1964) have helped us understand how individuals orient themselves into different social groups based on varying degrees of prosperity, privilege, and power. Some of these theories have predicted that racial boundaries will remain stable such that racial minorities will always exist as a subjugated group underneath Whites (Sidanius & Pratto, 2001). Others expect racial divisions to wane gradually and eventually disappear (Gordon, 1964). The results from this study, however, best align with theories of Black exceptionalism. Specifically, Black exceptionalism argues that the color line separating Blacks from all others will endure as America's true force of social and racial division (Leslie & Sears, 2022; Sears & Savalei, 2006). Newer immigrant groups seem to have disrupted this dichotomous structure only temporarily as their descendants appear poised toward assimilation with Whites. The findings demonstrated here that Asian-White Biracials (perhaps Latino-Whites also) are more similar to Whites in their politics suggests that these non-Black Biracial groups feel a closeness and shared destiny with Whites, the dominant racial group in America's hierarchy. This finding is compounded by the realization that 52% of U.S. born Latinos and a whopping 72% of U.S. born Asians are intermarrying outside of their minority racial group (Lee & Bean, 2012), and almost always with Whites. Though rates of interracial marriage have increased somewhat for Blacks (18% outmarry), currents demographic trends demonstrate quite transparently that Asians and Latinos may be on a pathway toward assimilation with Whites, or may perhaps rearticulate what it means to be part of the hegemonic group in the United states all together (See also Alba, 2020).

Methodologically, this study bridges together two varieties of scholarship—the study of racial politics and modern methods of machine learning—which are seldom used together despite the great advances they have made simultaneously in recent years. Specifically, I trained three different supervised binary classification ensemble machine learning models to predict among three different racial group class pairings: 1) Blacks and Whites, 2) Latinos and Whites, and 3) Asians and Whites. Then, rather than test these trained models on testing sets consisting of single-race participants, I applied these models to Biracial subgroup samples corresponding to each racial group pair: 1) Black-White Biracials, 2) Latino-White Biracials, 3) Asian-White Biracials. The mean probability with which each Biracial subgroup was predicted to be single-race White, rather than as a member of their single-race minority group, paints an empirical portrait of the theoretical distance between racial groups and foreshadows assimilative trends that may contour racial group boundaries in our near future.

This empirical approach improves on previous studies in a number of ways. First, existing research compares Biracials to their single-race counterparts using one or a handful of features at a time via linear (or partially linear) regression models (Alba, 2020; Davenport, 2018; Hochschild et al., 2012; Leslie & Sears, 2022; Masuoka, 2008). However, the ability to pick and choose which features to include in one's study runs the risk of prioritizing features which fit scholars' preexisting expectations—even if unintentionally. Selecting on the dependent variable here may bias our understanding of the true characteristics of Biracials as they compare to single-race groups. Instead, my empirical approach elides such risks by including, somewhat indiscriminately (though with certain a prior criteria standards), each and every possible variable which can be used to compare the political attitudinal characteristics of one group to another. This approach is consistent with a formidable literature which touts the advantages machine learning models pose for wading through through large and complex data structures to flexibly deduce functional forms and to accurately explain outcomes (Ferri-Garcıa & Rueda, 2020; Grimmer et al., 2021; Pirracchio et al., 2015). Ultimately, it is important that scholars investigating important topics such as race and racial identity avail themselves of the most appropriate and modern methods in order to make clear and precise inferenced in the future.

References

Abrajano, M., & Alvarez, R. M. (2010). *New faces, new voices: The hispanic electorate in america.* Princeton University Press.

Alba, R. (2020). The great demographic illusion. In *The great demographic illusion.* Princeton University Press.

Alba, R., Beck, B., & Basaran Sahin, D. (2018). The us mainstream expands–again. *Journal of Ethnic and Migration Studies*, *44*(1), 99–117.

Bennett, P. R. (2011). The social position of multiracial groups in the united states: Evidence from residential segregation. *Ethnic and Racial Studies*, *34*(4), 707–729.

Blumer, H. (1958). Race prejudice as a sense of group position. *Pacific sociological review*, *1*(1), 3–7.

Bobo, L. D. (1999). Prejudice as group position: Microfoundations of a sociological approach to racism and race relations. *Journal of Social Issues*, *55*(3), 445–472.

Bobo, L. D. (2004). Inequalities that endure? racial ideology, american politics, and the peculiar role of the social sciences. *The changing terrain of race and ethnicity*, 13–42.

Bratter, J. L. (2018). Multiracial identification and racial gaps: A work in progress. *The Annals of the American Academy of Political and Social Science*, *677*(1), 69–80.

Bratter, J. L., & Mason, C. (2016). Multiracial americans, health patterns, and health policy: Assessment and recommendations for ways forward. *Race policy and multiracial Americans*, 155–171.

Davenport, L. D. (2016). Beyond black and white: Biracial attitudes in contemporary us politics. *American Political Science Review*, *110*(1), 52–67.

Davenport, L. D. (2018). *Politics beyond black and white: Biracial identity and attitudes in america.* Cambridge University Press.

Davenport, L. D., Iyengar, S., & Westwood, S. J. (2022). Racial identity, group consciousness, and attitudes: A framework for assessing multiracial self-classification. *American Journal of Political Science*, *66*(3), 570–586.

Davis, F. J. (1991). *Who is black?: One nation's definition*. Penn State Press.

Does, S., Leslie, G. J., Bell, A. N., Yamaguchi-Pedroza, K., Shalbi, K. M., & Shih, M. (2021). The harms of racial miscategorization: Comparing multiracial individuals' well-being in the continental us versus hawai 'i. *Cultural diversity and ethnic minority psychology*.

Ferri-Garcıa, R., & Rueda, M. d. M. (2020). Propensity score adjustment using machine learning classification algorithms to control selection bias in online surveys. *PloS one*, *15*(4), e0231500.

Gordon, M. M. (1964). *Assimilation in american life: The role of race, religion, and national origins*. Oxford University Press on Demand.

Grimmer, J., Roberts, M. E., & Stewart, B. M. (2021). Machine learning for social science: An agnostic approach. *Annual Review of Political Science*, *24*, 395–419.

Hochschild, J. L., Weaver, V. M., & Burch, T. R. (2012). Creating a new racial order. In *Creating a new racial order*. Princeton University Press.

Hsu, C.-W., Chang, C.-C., Lin, C.-J., et al. (2003). A practical guide to support vector classification.

Lee, J., & Bean, F. D. (2012). A postracial society or a diversity paradox?: Race, immigration, and multiraciality in the twenty-first century. *Du Bois Review: Social Science Research on Race*, *9*(2), 419–437.

Leslie, G. J., & Sears, D. O. (2022). The heaviest drop of blood: Black exceptionalism among multiracials. *Political Psychology*.

Masuoka, N. (2008). Political attitudes and ideologies of multiracial americans: The implications of mixed race in the united states. *Political Research Quarterly*, *61*(2), 253–267.

Masuoka, N. (2017). *Multiracial identity and racial politics in the united states*. Oxford University Press.

Mihoko Doyle, J., & Kao, G. (2007). Are racial identities of multiracials stable? changing self-identification among single and multiple race individuals. *Social psychology quarterly*, *70*(4), 405–423.

Miller, A. H., Gurin, P., Gurin, G., & Malanchuk, O. (1981). Group consciousness and political participation. *American journal of political science*, 494–511.

Miller, B., Rocks, S., Catalina, S., Zemaitis, N., Daniels, K., & Londono, J. (2019). The missing link in contemporary health disparities research: A profile of the mental and self-rated health of multiracial young adults. *Health Sociology Review*, *28*(2), 209–227.

Park, R. E. (1928). Human migration and the marginal man. *American journal of sociology*, *33*(6), 881–893.

Pirracchio, R., Petersen, M. L., & Van Der Laan, M. (2015). Improving propensity score estimators' robustness to model misspecification using super learner. *American journal of epidemiology*, *181*(2), 108–119.

Polley, E. C., Rose, S., & Laan, M. J. (2011). Super learning. In *Targeted learning* (pp. 43–66). Springer.

Samii, C., Paler, L., & Daly, S. Z. (2016). Retrospective causal inference with machine learning ensembles: An application to anti-recidivism policies in colombia. *Political Analysis*, *24*(4), 434–456.

Sears, D. O., & Savalei, V. (2006). The political color line in america: Many "peoples of color" or black exceptionalism? *Political Psychology*, *27*(6), 895–924.

Sidanius, J., & Pratto, F. (2001). *Social dominance: An intergroup theory of social hierarchy and oppression*. Cambridge University Press.

Sims, J. P. (2016). Reevaluation of the influence of appearance and reflected appraisals for mixed-race identity: The role of consistent inconsistent racial perception. *Sociology of Race and Ethnicity*, *2*(4), 569–583.

Stonequist, E. V. (1935). The problem of the marginal man. *American journal of sociology*, *41*(1), 1–12.

Tate, K. (1994). *From protest to politics: The new black voters in american elections*. Harvard University Press.

Van Der Laan, M. J., & Dudoit, S. (2003). Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples.

Van der Laan, M. J., Polley, E. C., & Hubbard, A. E. (2007). Super learner. *Statistical applications in genetics and molecular biology*, *6*(1).