Implicit Statistical Discrimination in Predictive Models

Devin G. Pope
The Wharton School
University of Pennsylvania

Justin R. Sydnor
Department of Economics
Weatherhead School of Management
Case Western Reserve University

September 2007

Working Paper # 2007-09-11

Risk Management and Decision Processes Center
The Wharton School, University of Pennsylvania
3730 Walnut Street, Jon Huntsman Hall, Suite 500
Philadelphia, PA, 19104
USA
Phone: 215-898-4589
Fax: 215-573-2130
http://opim.wharton.upenn.edu/risk/

CITATION AND REPRODUCTION

This document appears as Working Paper of the Wharton Risk Management and Decision Processes Center, The Wharton School of the University of Pennsylvania. Comments are welcome and may be directed to the authors.

This paper may be cited as: Devin G. Pope and Justin R. Sydnor, "**Implicit Statistical Discrimination in Predictive Models,**" Risk Management and Decision Processes Center, The Wharton School of the University of Pennsylvania, September 2007.

Implicit Statistical Discrimination in Predictive Models*

Devin G. Pope
The Wharton School
University of Pennsylvania

Justin R. Sydnor
Department of Economics
Weatherhead School of Management
Case Western Reserve University

May, 2007

Abstract

How should statistical profiling models be implemented when anti-discrimination policies prohibit basing predictions on characteristics such as race, gender, or age? Companies, schools, and social-program administrators typically address such concerns by simply excluding these sensitive characteristics from the models they estimate. However, other variables that may be correlated with these omitted characteristics – such as zip codes, credit scores, and job tenure – are routinely used and may serve as partial proxies for the excluded groups. We examine the importance of this issue for the federally mandated Worker Profiling and Reemployment Services system, in which states profile unemployment-insurance (UI) claimants and require workshop attendance from those who are predicted to be likely to exhaust their benefits. Using a large data set on UI claimants, we utilize a simple procedure to compare and contrast the approach commonly used by states with one that eliminates the ability for modeling variables to proxy for sensitive characteristics. In this way we can establish the degree to which the program outcomes are affected by modeling variables serving as proxies for the excluded characteristics. We find a significant effect, especially across racial groups, which we demonstrate is largely driven by the correlation between race and zip codes. Our benchmark results suggest that eliminating the influence of the sensitive characteristics on the predictive process would decrease the fraction of required workshop attendees that are black by roughly 25%. We address the question of predictive accuracy and discuss the relevance of these findings for other situations such as mortgage lending, insurance pricing, and college admissions.

Contact Pope at dpope@econ.berkeley.edu and Sydnor at justin.sydnor@case.edu.

*We thank Dan Black, David Card, Stefano DellaVigna, Jesse Leary, Jaren Pope, Matthew Rabin, Stephen Ross, Jesse Rothstein, Paige Skiba, Jeff Smith, and seminar participants at UC Berkeley for helpful comments and suggestions. We are also grateful to Dan Black and David Card for generously sharing data used for this project. All errors are our own.

In an increasing number of important economic settings decision-makers use statistical predictions when determining how to treat people. Examples include mortgage and credit-card lending, insurance pricing, college admissions, and social-benefit provision. Which individual characteristics should modelers use to make these predictions? The simple answer is any characteristic that improves predictive accuracy. In many settings, though, the use of a number of "socially unacceptable predictors" (SUPs) – such as race, gender, and age – in the predictive process is considered inappropriate statistical discrimination. Profilers typically respond by taking a "SUP-blind" approach, simply dropping the sensitive characteristics from their models.[1] However, at the same time they routinely use other variables – such as zip codes, credit scores, and job tenure – that may be highly correlated with these omitted characteristics. Because of this correlation, these variables may serve as partial proxies for the excluded categories. As a result, excluded variables such as race, gender, and age may continue to influence predictions by distorting the weight given to other variables relative to those variables' direct effect on the outcome of interest. We label the effect that this distortion has on predicted outcomes "implicit statistical discrimination."

From a conceptual point of view, implicit statistical discrimination is simply a reflection of the standard omitted variable bias problem. What makes it more interesting, however, is that the potential bias here does not result from unobserved variables, but is rather the result of a determined effort to ignore the influence of certain known characteristics. While a researcher interested in the causal effect of job tenure on unemployment would not exclude age from her model, this is a routine solution for real-world practitioners concerned about discrimination. In these practical applications, does the residual influence of SUPs introduced by the standard approach affect predictions in a meaningful way, and if so is there a realistic alternative?

[1] In the literature on discrimination, treating individuals differently on the explicit basis of race, gender, or other protected characteristics is known as "disparate treatment." What we call SUP-blindness means avoiding disparate treatment.

In this paper, we analyze the way unemployment-insurance claimants are treated on the basis of predicted benefit exhaustion, and find that in this setting the answer to both of these questions is "yes". Specifically, we look at the importance of implicit statistical discrimination in the federally mandated Worker Profiling and Reemployment Services (WPRS) system. Since 1993, individual states have been required by the federal government to develop and use profiling models that predict which unemployment-insurance (UI) claimants are most likely to exhaust their UI benefits. For those profiled as likely to exhaust, UI benefits are withheld unless the claimant attends a reemployment service workshop provided by the state.[2] In general, the WPRS system allows states to use any variables that help them predict the likelihood of benefit exhaustion. Examples include education, job tenure, occupation, and geographic area (e.g. zip codes). However, because of anti-discrimination laws, states are not allowed to base predictions on any of the following variables (SUPs): age, race, ethnic group, sex, color, national origin, disability, religion, political affiliation, and citizenship.

When implementing WPRS models, states typically estimate simple OLS regressions of benefit-exhaustion measures on individual characteristics (e.g. education, zip code) while omitting the SUPs. Our concern with this approach is that the weight given to zip codes may partially reflect differences in benefit exhaustion across racial groups, the weight given to job tenure may partially account for age, and so on. To examine the potential importance of these concerns, we use a simple estimation procedure, first introduced by Ross and Yinger (2002), which eliminates the ability for modeling variables to proxy for SUPs, while still maintaining SUP-blind predictions.[3] By comparing and contrasting the common approach with this alternative procedure, we can establish the degree to which the program outcomes are affected by correlations between modeling variables and excluded characteristics.

---

[2] In 2005, throughout the United States, 1.08 million UI claimants were required to attend reemployment workshops. Statistics from the WPRS system are available on the following Department of Labor website: http://ows.doleta.gov/unemploy/profile.asp. The statistic reported here includes all states except for New Mexico, for which data was unavailable.

[3] Ross and Yinger's work builds on earlier theoretical analysis by Lundberg (1991).

This approach involves first estimating a full model which includes all of the usual variables *as well as* all relevant SUPs. Including the SUPs in the estimation ensures that the coefficients on the other variables are orthogonal to their correlation with the sensitive characteristics. The key step in the process, however, is that only the coefficients from the non-sensitive predictors are used when producing the predicted values for each individual. The SUP coefficients are replaced with zeros.[4] Thus, this approach ensures that the weight given to each variable that is used to form predicted values reflects its estimated direct effect on the outcome of interest. Importantly, though, while the SUPs are used in the estimation, they are not used in forming predicted values, and thus the procedure maintains SUP-blindness.[5]

Using a data set that contains a 100% sample of UI claimants from the state of New Jersey between 1995 and 1997, we analyze the effect of these different approaches on a WPRS model similar to those being used by most states. We find that using the Ross-Yinger procedure instead of the typical approach significantly affects the SUP composition of the group of claimants required to attend reemployment service workshops. The largest effect is for the composition by race. For example, under a system in which the 10% of claimants with the highest predicted exhaustion rates are required to attend workshops, the unbiased procedure decreases the proportion of blacks sent by over 25%.[6] We show that this difference mainly arises due to the strong correlation between race and geographic controls (i.e. zip codes). In fact, we demonstrate that the bias on the zip-code coefficients in the standard approach results in claimants of any race being more likely to be assigned to a workshop if they come from areas with high black concentration relative to their counterparts from highly white areas. In total, at this 10% cutoff 14,444 claimants (24% of the 59,092 workshop attendees) would be affected by eliminating the residual influence of the sensitive characteristics. Extrapolating these findings to the nation as a

[4] As described below, in situations where levels are relevant (rather than just rank order), a simple renormalization of the mean is also needed.

[5] As we discuss below, this approach generates different predicted values than the standard approach only to the extent that the variables in the standard approach are in fact proxying for the SUPs. If no correlation exists between included variables and SUPs, or if the SUPs have no direct predictive effect on the measured outcome, there will be no difference in the two approaches.

[6] The percent of workshop participants that are black would fall from 21.9% to 16.0%.

whole suggests that a significant number of UI claimants across the country (on the order of 250,000 annually) are affected by this issue.[7]

These results show that that the bias inherent in the standard attempt to comply with anti-discrimination laws affects who is profiled as a likely UI-benefit exhauster. However, this leaves at least two important questions unanswered. First, who benefits from eliminating this implicit statistical discrimination? Black et al. (2003a) show that the WPRS system decreases the mean weeks of UI benefit receipt among claimants who are required to attend a reemployment workshop. That would seem to suggest that these workshops are effective and beneficial. However, Black et al. conclude that a large part of this reduction is a direct result of claimants being notified of the requirement to attend a workshop, as opposed to benefits accrued from actually attending. This seems to suggest that claimants find the prospect of attending the workshop undesirable. In this study we remain agnostic about the benefits (or the lack thereof) of reemployment workshops. In fact, this is perhaps a benefit of studying the WPRS system rather than more contentious topics such as college admissions or mortgage lending. The purpose of this paper is to examine an important issue in the practical implementation of statistical models that exists regardless of who the winners and losers from a change in procedure might be.

The second question, which may be closely linked to the first, is whether policymakers should seek to eliminate this implicit statistical discrimination. There is a tradeoff here between fairness and predictive accuracy. As we demonstrate in our empirical analysis, because the standard approach is less constrained than the Ross-Yinger approach, implementing this alternative procedure decreases the predictive accuracy of the benefit-exhaustion model. It may be the case that policymakers wish to capture any available predictive power they can, as long as they are not treating people differently at the individual level on the basis of race, gender, or age. However, we see a distinctly slippery slope in this logic, because any advantage the standard

[7] 250,000 comes from 25% of the over 1 million workshop attendees throughout the U.S. annually.

approach has over the Ross-Yinger procedure in predictive accuracy is gained solely through the ability of the included variables to proxy for the excluded SUPs. If one simply wants the most accurate SUP-blind predictions, the optimal modeling strategy should be to use proxies that are extremely correlated (though not quite perfectly correlated) with characteristics like race. We suspect few would argue in favor of such a method. Furthermore, we demonstrate that while it improves overall predictive accuracy, the biasing of the non-sensitive coefficients under the standard approach reduces predictive accuracy within SUP groups.[8] Ultimately, while we cannot resolve the debate between fairness and accuracy, we would suggest that at the very least it is a debate worth having.

This analysis suggests a way to reframe the debates about which variables are appropriate to use in statistical models, such as those that surround the use of credit scoring in insurance, zip codes in mortgage lending, or of SAT scores in college admissions. These debates generally discuss just two options. Either these variables should not be allowed in profiling models because they may proxy for sensitive characteristics or they should be included because of their predictive value.[9] Our analysis suggests that given anti-discrimination concerns there exists a middle ground. Using the proposed method in this paper, much of the predictive power that comes from variables such as zip codes can still be used while the predictive power that comes from the correlation between zip codes and sensitive characteristics is removed. Indeed, in our analysis of the Worker Profiling system, we document that while the proposed method lowers overall predictive accuracy relative to the standard approach, implementing the proposed procedure *improves* overall predictive accuracy meaningfully when the alternative involves banning the use of zip codes outright.

[8] Since there are differences in benefit exhaustion by race, allowing zip codes to partially proxy for racial composition improves predictions overall because it helps identify those who are more likely to be minorities. If one looks at predictive accuracy within a racial group, however, this skewing of weights given to different zip-codes decreases predictive accuracy.

[9] One prominent example of such debates arose in the passage of California's Proposition 103, which limited the variables that automobile insurers could use to determine their pricing schedules. While it might be true that local conditions (e.g. zip code) affect insurance losses directly, many argued that there was no way to guarantee that these variables were not serving as proxies for sensitive characteristics.

In the next section we discuss how our analysis relates to the existing literatures on discrimination and affirmative action. The remainder of the paper then proceeds as follows: Section II outlines the conceptual framework of profiling in the presence of SUPs for the OLS case. Section III describes the worker profiling data and Section IV presents the results of our analysis. Finally, Section V concludes.

I. Related Literature

The general idea of statistical discrimination has long been a topic of interest for economists. See, for instance, Arrow (1972), Arrow (1973), Phelps (1972), Aigner and Cain (1977), Borjas and Goldberg (1978), and Schwab (1986).[10] However, by far the most relevant paper for our study is Ross and Yinger (2002). As noted above, the alternative procedure for generating predicted values that we use was (to our knowledge) first introduced by Ross and Yinger in their book *The Color of Credit*. They use this approach to look at discrimination in mortgage lending. They argue that not accounting for minority status when predicting loan default may lead to biased predicted values in a way that harms minorities. However, most likely because their data set lacks key variables that may be highly correlated with minority status, such as zip code and credit score, they find no sizable effects of the methodology change.[11]

Ross and Yinger's study was partially motivated by the theoretical work of Lundberg (1991), who examined the effectiveness of equal opportunity laws in labor markets that exhibit statistical discrimination. Lundberg considers how policy makers can enforce these laws when firms use measures such as test scores or height as partial proxies for race or gender. While we use the phrase "implicit statistical discrimination", both Lundberg and Ross and Yinger relate these concerns to what is known in legal settings as disparate-impact discrimination. Most of this

[10] Unlike taste-based discrimination (Becker 1957), statistical discrimination is information based and does not imply any animus on the part of the decision maker against a particular group.
[11] Ross and Yinger (2002) also present simulations that suggest that the use of credit scores in standard models could significantly harm minorities.

litigation centers around labor-market practices, and courts have ruled that basing hiring criterion on individual characteristics that do not influence job performance but correlate with membership in protected groups is illegal.[12] It is less clear whether implicit statistical discrimination arising from the use of variables that are both directly related to performance and highly correlated with SUPs (as we see in the WPRS system) would be illegal under disparate-impact rulings.

Disparate-impact discrimination is not the same as the more frequently discussed case of disparate-treatment discrimination. Disparate-treatment discrimination refers to situations in which otherwise identical individuals are treated differently due to their membership in a protected group.[13] There is a large literature surrounding tests of disparate treatment. Many studies try to establish whether, after controlling for all information that decision makers say they are using, variables such as race or gender are able to predict their decisions. A classic example of this type of test is Tootle's (1996) study of mortgage redlining in Boston. It is worth noting that under implicit statistical discrimination otherwise identical individuals are treated the same regardless of group membership. Therefore, while these sorts of tests may be appropriate for establishing the existence of disparate-treatment discrimination, they could not identify the sort of implicit statistical discrimination we find in the WPRS system.

This paper also relates to some of the literature on affirmative action in labor markets and college admissions. For instance, both Chan and Eyster (2003) and Fryer, Loury, and Yuret (2003) discuss how universities may react to bans on affirmative action by altering the weights they give to various admissions criteria in a way that increases diversity. Both papers suggest that quota-based, affirmative-action policies will be more efficient than such a distorted admissions process. Similarly in Section IV, we use a quota method to illustrate the within-group inefficiencies that are caused by using biased SAP coefficients. Another relevant study is

[12] See Ross (2005) and Ross & Yinger (2002), especially pages 313-324, for a discussion of relevant court cases and a review of the legality of disparate-impact discrimination. Also see the Supreme Court ruling in Griggs v. Duke Power Co. (401 U.S. 424 [1971]) for an early ruling on disparate-impact discrimination.
[13] Disparate-treatment discrimination can arise because of either taste-based discrimination or statistical discrimination.

Rothstein (2005), who finds that including demographic and socioeconomic characteristics in a regression of college success on SAT and high-school GPA reduces the importance of the SAT relative to high-school GPA. Thus if universities base admissions on statistical models without accounting for demographics, the importance given to SAT scores relative to GPA in the admission process may be greater than is warranted by its *ceteris paribus* direct effect on college achievement.

II. Conceptual Framework

Our analysis begins by assuming a fully specified baseline model. This model may contain both socially acceptable and unacceptable predictors. Assume a data generating process satisfying the standard OLS assumptions:

$$(1) \qquad y_i = \alpha + \beta X_i^{SAP} + \theta X_i^{SUP} + \varepsilon_i ,$$

where y_i is the observed outcome for person *i* who has "socially unacceptable predictor" (SAP) characteristics X_i^{SAP} and SUP characteristics X_i^{SUP}.[14] The outcome y_i could be any outcome of interest such as whether individual *i* exhausted his/her UI benefits, defaulted on his/her mortgage loan, got in a car accident, or the GPA he/she received in college. Estimating this baseline model, which we label the *full method*, provides the following OLS predicted values

$$(2) \qquad \hat{y}_i^{full} = \hat{\alpha}^{full} + \hat{\beta}^{full} X_i^{SAP} + \hat{\theta}^{full} X_i^{SUP} .$$

By their very definition, the use of the SUPs in generating $\hat{y}_i^{full}$ is considered unacceptable. The common approach to dealing with this issue is to simply omit X_i^{SUP} from the estimation of Equation (1). The resulting predicted values from this *common method* are

[14] While not specifically treated in this paper, the methods analyzed here could be extended to include more complicated models (e.g. interactions between SAPs and SUPs).

(3) $$\hat{y}^{common} = \hat{\alpha}^{common} + \hat{\beta}^{common} X_i^{SAP}.$$

To the extent that $\hat{\theta}^{full}$ does not equal the zero vector and that there is correlation between X_i^{SAP} and X_i^{SUP}, $\hat{\beta}^{common}$ is of course biased. The direction of this bias depends on both the sign of the elements of $\hat{\theta}^{full}$ and the nature of the correlation between X_i^{SAP} and X_i^{SUP} (Greene 2003, pp. 334-337).

The *proposed method* avoids this omitted-variable-bias problem by estimating Equation (1) to obtain the unbiased coefficients on the SAPs ($\hat{\beta}^{full}$) from Equation (2). However, as opposed to Equation (2) which includes $\hat{\theta}^{full} X_i^{SUP}$ in the generation of the predicted values, the *proposed method* results in the following predicted values:

(4) $$\hat{y}_i^{proposed} = \hat{\alpha}^{proposed} + \hat{\beta}^{full} X_i^{SAP},$$

where $\hat{\alpha}^{proposed}$ is obtained by regressing ($y_i - \hat{\beta}^{full} X_i^{SAP}$) on a constant.

To summarize, the *full method* is a case of explicit statistical discrimination and would fall under the legal heading of disparate-treatment discrimination. Because both Equations (3) and (4) have the desired benefit that people with the same SAP characteristics receive the same predictive values from the model, both the *common* and *proposed methods* are "race-blind" or more generally "SUP-blind". That is, neither is a case of disparate-treatment discrimination. However, since the *common method* allows the SAPs to partially proxy for the SUPs, it captures some of the predictive power of the SUPs. Thus the *common method* is a case of implicit statistical discrimination. The *proposed method,* on the other hand, avoids this implicit statistical discrimination by basing prediction only on the direct effect of the SAPs.

An additional complication to this framework exists if there are SAPs which have a causal impact on y_i, but are unobserved to the researcher. As a somewhat contrived example, imagine that height (a SAP) helps to determine how quickly a worker is hired, and therefore the fraction of UI benefits exhausted. Furthermore assume that there is a correlation between height and (say) race, but that height is unobservable to the researcher. Thus, part of the direct affect attributed to race in the *full-method* regression would actually be a reflection of race proxying for height. In such a situation, the *common method* allows zip-code factors to (partially) proxy for height, through the correlation between race and zip code. As long as the only variable (SAP or SUP) that was directly correlated to height was race, the *proposed method* would eliminate the predictive power of height rather than the (inappropriate) predictive influence of race.

While this is certainly a valid concern, we feel that the *proposed method* is likely the more appropriate and cautious approach in the face of concerns about discrimination. In the context of unemployment duration, given the literature on discrimination in labor markets (e.g. Bertrand and Mullainathan, 2004), it is entirely feasible that race, age, and citizenship status are directly causal to benefit exhaustion. Even if this is not the case, and these variables are proxying for some other unobserved characteristic, it is not clear whether that unobserved characteristic would be considered socially acceptable for use if it were observable. Ultimately, it is impossible to prove that the effect of the SUPs on benefit exhaustion is causal, but it is possible to show that the effect is not causal. Thus, in this situation, we would argue that the logical default is to use the *proposed method* and place the burden on the profiler to find additional SAPs if it is believed that the SUPs do not have causal effects. To return to our (contrived) example, if one thinks race proxies for height, the solution is to find data on height.

III. Data

We compare these different methods in order to analyze potential discrimination in a worker profiling model of the type mandated by the U.S. Worker Profiling and Reemployment Services (WPRS) system. In 1993, the Unemployment Compensation Amendments (P.L. 103 - 152) to the Social Security Act required all states to develop and implement statistical profiling models of unemployment insurance (UI) claimants. The stated goal was to identify "who will be likely to exhaust regular compensation and will need job search assistance services." Those individuals who are predicted to be most likely to exhaust their benefits must attend reemployment service classes/workshops in order to continue receiving UI benefits. The percentage of claimants who are sent to the classes/workshops and the type of classes/workshops that are taught vary by state and depend in part on capacity constraints. Similar worker profiling and reemployment service systems are in effect in Australia, Canada, and the United Kingdom (OECD 1998).[15]

The Department of Labor gave suggestions to states on potentially useful modeling variables and also specified certain variables that according to Section 188 of the Workforce Investment Act of 1998 are not to be used in the predictive process. Specifically, it was determined that "a worker profiling system is not permitted to produce results which discriminate against groups of people… For this reason, the following variables may not be used in the worker profiling: age, race, ethnic group, sex, color, national origin, disability, religion, political affiliation, and citizenship" (Wandner and Messenger 1999). We thus classify these variables as SUPs. While states currently do not include any of these variables in their predictive models, most of them do have data for age, race, sex, and citizenship.

In a comprehensive report, Black et al. (2003b) provided recommendations to states on how to develop and implement a Worker Profiling model. While states vary in the complexity of the models that they use, most states have implemented a profiling model similar to what Black et al. proposed (Card 2003). We use Black et al.'s recommendations as the foundation for our baseline model.

Black et al. make a few key recommendations. First, they suggest that the model should be estimated using OLS rather than alternatives such as logit or Tobit. Second, they argue that the dependent variable should be the fraction of potential benefits that an

[15] Further detail regarding the Worker Profiling and Reemployment Services program in the U.S. can be found in Balducchi (1996) and Wandner (1997).

individual exhausts rather than a discrete measure of benefit exhaustion. Furthermore, they argue that richer models do better, although there is an obvious tradeoff between tractability and predictive accuracy. In the end, the model that Black et al. describe as their preferred model includes the following predictors: Education dummies, job tenure, job tenure squared, 1-digit occupation codes, pre-unemployment wages, local-unemployment-office fixed effects, welfare measures (food stamps, public transit use, etc.), and measures of previous UI benefit exhaustion if applicable.

We analyze a data set that contains a 100% sample of UI claimants from the state of New Jersey between 1995 and 1997.[16] We include as SAPs in our baseline model all variables that the New Jersey data has available that were suggested by Black et al. In the end, we are able to include education dummies, job tenure, job tenure squared, 1-digit industry codes, pre-unemployment wages, and local-unemployment-office fixed effects in the model.[17] The New Jersey data set does not include measures for welfare or previous UI benefit exhaustion. While clearly imperfect, these variables are similar to those used by many states and should provide a reasonable illustration of the potential implicit statistical discrimination inherent in these models. The New Jersey data also contains information on the race, gender, citizenship, and age of the UI claimants. These are the SUPs that we analyze.

We begin with 732,980 UI claimant observations between 1995 and 1997. We drop 39,923 observations (5.4%) that have data problems or extreme values.[18] Following the

[16] See Card and Levine (2000) for a more detailed description of the data.
[17] As studied by Card and Levine (2000), New Jersey had a special program during a portion of our sample, which allowed claimants to receive extra weeks of benefits. We include fixed effects denoting whether a claimant was eligible for this program in our analysis.
[18] Observations are dropped if they are missing an industry code, gender, age < 18, age > 80, average weekly wage in top 1% or bottom 1%, or if the race was undefined (not defined as black (non-Hispanic), white (non-Hispanic), or Hispanic)

directions of Black et al. and the Department of Labor, we also drop 102,133 individuals (13.9%) that indicated they were union members.[19]

Summary statistics for the remaining 590,924 observations are displayed in Table 1. The Table lists measures of benefit exhaustion, SUP characteristics, and SAP characteristics. The SUP classifications are race (black (non-Hispanic), white (non-Hispanic), Hispanic), age (<30, 30-45, >45), gender, and U.S. citizenship.[20] Looking at the dependent variable measure (fraction of benefits exhausted), UI claimants on average exhausted 74% of their available benefits. This exhaustion rate appears to vary across race and age groups and to a lesser extent gender and citizenship. For instance, on average whites, blacks, and Hispanics exhausted 72%, 78%, and 76% of available benefits respectively. Table 1 also provides the first glimpse of correlations that exist between SUP and SAP characteristics. For example, there are high correlations between tenure and age, between pre-unemployment wages and race, and between education and citizenship. While not reported due to space constraints, some of the largest correlations occur between SUP characteristics (especially race) and local-office controls.

IV. Results

Changes in regression coefficients. Following the specification suggested by Black et al. (2003b), Table 2 presents the regressions using the fraction of potential benefits exhausted as the dependent variable. The first column provides the estimated coefficients from the OLS regression using both the SAPs and the SUPs as predictors (*full method,* Equation (1)). Of particular interest, we find that the SUPs do have an effect on benefit exhaustion outcomes even after controlling for the SAPs in the model. Females, older

[19] Black et al. actually drop claimants who find work through a union hiring hall or who have a known recall date. Since these variables are not available in our dataset we drop all union workers.

[20] The race designations are those available in the data. The data includes exact age, however, to ease the interpretation of the results and after examining the non-parametric relationship between age and benefit exhaustion, we chose to use three discrete age groups in the analysis.

individuals, US citizens, blacks and Hispanics are all more likely to use up their UI benefits than their equivalent counterparts. While we include all industry and local-unemployment-office dummies in the regression, due to space constraints we only present the coefficient values on two industry dummies and four local-unemployment-offices dummies. To aid the discussion that follows, we highlight those controls that are most correlated with race. Specifically, we report the coefficient on public administration, which is the industry that has the highest percentage of black claimants, and construction, which is the industry that has the highest percentage of white claimants. Similarly, we present the coefficients for local-unemployment office 10 and 12, which have the highest percentage of black UI claimants, and office 21 and 26, which have the highest percentage of white UI claimants. In brackets next to each of these offices and industries we indicate the racial breakdown.[21]

The second column of Table 2 presents the coefficient values that are generated by regressing fraction of potential benefits exhausted on only the SAPs (*common method*). Column 3 provides the difference in the coefficients between the *full* and *common method* with bootstrapped standard errors reported in parenthesis to facilitate tests of significance (almost all coefficient differences are statistically significant at conventional levels).[22] The changes in the SAP coefficients between the *full method* and the *common method* reflect the correlations between the SAPs and omitted SUPs. For example, in the *full method* the coefficient on tenure is negative and statistically significant. However, once the SUPs are dropped from the model, the coefficient on tenure switches sign and is indistinguishable from zero. This is likely the result of the positive correlation between a claimant's age and his/her tenure with previous employer. Indeed, the *full method* reveals that all else equal claimants over the age of 45 exhaust on average 5.6% more of their available benefits than do those under the age of 30.

[21] The local-unemployment office base-group (office #17) and the industry base group (services) were chosen to be groups that had racial compositions most similar to the makeup of the overall population.
[22] In our bootstrap routine we used 100% samples with replacement. The regression coefficients for both the *full* and the *common method* were recorded for each of 200 repetitions.

Some of the biggest effects of such omitted variable bias are seen in the changes to the coefficients for the local-unemployment-offices controls. For instance, the predicted fraction of benefits exhausted by claimants from local office #10 (which has the highest concentration of black claimants) rises by 2.8 percentage points when SUP controls are dropped from the model. In addition, the coefficient for local office #21 (which has the highest concentration of white claimants) falls by 0.7 percentage points. The combined effect of these two changes (3.5%) indicates that in the *common method,* these highly segregated local unemployment office controls are able to absorb a significant fraction of the black-white difference that was found when using the *full method* (5.4%). In contrast, the *proposed method* uses the same SAP coefficients as the *full method* and simply eliminates the indicated 5.4% black-white difference during the generation of the final predicted values.

Changes in workshop composition. We are interested in how the different coefficients on the SAPs might affect which claimants are required to attend reemployment service classes. Specifically we ask what the composition of the class (by race, gender, age, citizenship) would be under the alternative profiling methods. In Table 3, for each method we consider two hypothetical cutoff rules: 90% and 80%.[23] Under these rules, the top 10% and 20% of predicted exhausters, respectively, would be required to attend a workshop. Claimants are ranked by their predicted benefit exhaustion for each of the three profiling methods discussed in this paper. The table then compares the SUP makeup of the top 10% or 20% of predicted exhausters using these rankings.

Looking at the 90%-cutoff rule, if the *full method* is used (which includes explicit controls for racial categories), 30.5% of the class is white, 41.3% is black, and 28.2% is Hispanic. In the *common method* (which drops the racial controls from the estimation), the class composition is 45.2% white, 21.9% black, and 32.9% Hispanic. Finally, in the *proposed method* (which includes race in the estimation, but not in the predicted values), the makeup is

[23] While these cutoff rules are used simply as examples, in 2005, New Jersey required 16.9% of all UI claimants to attend workshops.

47.6% white, 16.0% black, and 36.4% Hispanic. It is worth noting that in the absence of omitted variable bias, the predictions from the *common* and *proposed method*s would be identical. Table 3 reveals, however, that the fractions of the required workshop participants that are white and Hispanic are lowered by 2.4% and 3.5%, respectively, under the *common method* relative to the proposed alternative. Blacks, on the other hand, make up 5.9% more of the workshop under the *common method*.[24] Thus, using the *proposed method*, instead of the *common method*, generates more than a 25% reduction in the fraction of the class that is black.[25] A bootstrap routine reveals that these differences are significant at the 1% level.

For gender, age, and citizenship the differences between the workshop composition under the *common* and *proposed method*s are in the directions that one would predict given the omitted variable bias inherent in the *common method*. Females, older claimants, and U.S. citizens all represent a higher fraction of the workshop when using the *common* rather than the *proposed method*. The bootstrapped standard errors reveal that in all but one case these differences are statistically significant at the 5% level, with most significant at the 1% level. In general, however, the difference between the two methods is larger for breakdowns by race than age, gender, or citizenship.

All of these results are similar for the 80%-cutoff rule. One might worry, however, that the results may vary depending on the chosen cutoff rule. Figure 1 illustrates the breakdowns from Table 3 for any possible assignment rule. For each SUP characteristic, the graphs show the fraction of the workshop with that characteristic at each cutoff level. The results highlighted in Table 3 hold across a wide range of potential cutoff rules and again the largest differences between the *proposed* and *common method*s are found across racial categories.

[24] Given that the coefficient on the Hispanic dummy in Table 2 suggests that being Hispanic increases the probability of exhaustion, it may be surprising to find that Hispanics are more likely to get sent to attend workshops under the *proposed method*. If the regressions only included Hispanics and whites, then Hispanics would be sent to classes less often under the *proposed method*. However, this effect is being overpowered by the much larger change in the coefficients on highly black zip codes, since blacks are relatively more likely to exhaust benefits than both whites *and* Hispanics.

[25] (21.9% - 16.0%)/21.95 = 0.269

Effects within a SUP group. Thus far we have highlighted how implementing the *proposed method* rather than the *common method* might change the composition of a reemployment services workshop. Focusing solely on these results, however, understates the total number of people that would be affected by such a change. The differences between the two methods are manifested not only between, but within SUP groups as well. Because in the *common method* the SAP coefficients are allowed to proxy for the omitted SUPs, claimants who have SAP characteristics that easily identify them as being a member of a particular SUP group are treated differently than otherwise identical claimants. For example, relative to the *proposed method*, under the *common method* black claimants who live in areas with a high black concentration are assigned higher predicted exhaustion levels relative to otherwise similar blacks who live in high white concentration areas. Thus, while the *common method* sends more blacks in general than the *proposed method*, it sends especially more blacks whose SAPs "identify" them as likely being black. The same effects also exist for whites and Hispanics who happen to live in highly concentrated black areas. Figure 2 shows this clear pattern by comparing the difference in the percent of each racial group that is required to attend a workshop under a 90%-cutoff rule between the *common* and *proposed methods* across local offices with varying levels of black-claimant concentration.

Table 4 looks at this issue by analyzing the effects of the *common* versus *proposed methods* under a hypothetical 90%-assignment rule for each of the 36 distinct SUP combinations.[26] The table provides the total number of claimants for each of the possible SUP combinations. The subsequent four columns indicate how many of these individuals are assigned to attend a workshop under just the *common method*, just the *proposed method*, under both methods, and in neither method. Overall, both methods send the same total number of people (59,092), but they differ in how many are sent from each SUP combination. For instance, under the 90%-cutoff rule, the *proposed method* assigns 3,739 claimants from the first SUP group (US citizen, white, male, young) to the workshop while the *common method* sends

[26] (3 races) X (2 genders) X (3 age groups) X (2 citizenships) = 36

3,280. Thus, there are 459 more claimants assigned from this SUP combination under the *proposed method*.

Breaking these numbers down further, only 3,198 of these claimants are assigned under both systems, which implies that 623 claimants are actually affected by the change. That is, even though the *common method* sends fewer claimants from this group, 82 claimants are sent under the *common method* that would not have been sent under the *proposed*. The shaded numbers in the table represent the total number of people who are affected by such within-group ranking differences. To place these numbers in perspective, note that since the *proposed method* has the same within-SUP-group rankings as the *full method*, when comparing the *proposed* to the *full* all shaded numbers would be zero.[27] If the *proposed method* were to send more than the *full*, no one would be sent under the *full method* but not also the *proposed*. In total, for the 90%-cutoff rule, 14,444 claimants would be affected by eliminating the implicit statistical discrimination through implementing the proposed methodology change. This represents 24% of the total number of claimants required to attend a workshop under this cutoff rule.

Interpretation. While these results clearly illustrate that a large number of UI claimants are affected by the question of whether to remove any residual predictive influence of SUPs, it is less clear how to weigh the winners and losers of the alternative procedures. Black et al. (2003a) show that the WPRS system decreases the mean weeks of UI benefit receipt among marginal claimants who are required to attend a reemployment workshop. However, they conclude that a large part of this reduction is a direct result of claimants being notified of the requirement to attend a workshop as opposed to benefits accrued from actually attending. This seems to suggest that while tax payers may benefit from reduced UI-benefit expenditures with this program, the marginal claimants find the prospect of

[27]. While the *proposed method* deviates from the *full method*'s overall rankings, it maintains the same ranking order for claimants that have the same SUP characteristics (i.e. X_i^{SUP}). In contrast, the *common method* allows the SAP coefficients to change and thus alters the rankings even within a set of SUP characteristics.

attending the workshop undesirable. In this study we remain agnostic about the benefits (or the lack thereof) of reemployment workshops. Our interpretation of anti-discrimination policies is that they aim for what we call SUP-blindness, and that the question of how much implicit statistical discrimination affects this program should by and large be abstracted from the question of whether removing these effects benefits (say) those in highly white neighborhoods or those in highly black neighborhoods.

Effects on predictive accuracy. Perhaps the more important issue that who would win or lose, is what the different effects of the alternative methods on predictive accuracy would be. Because the estimation in the *common method* is less constrained than the *proposed method*, the *common method* obviously achieves more overall predictive power than the *proposed method*. For instance, the average fraction of benefits actually exhausted by the 7,222 claimants assigned to the class in Table 4 under the *common method* but not the *proposed* was 80.39%. Those assigned under the *proposed* but not the *common method*, however, exhausted on average 80.25%. Thus if the goal is to identify those most likely to exhaust their benefits, the *common method* overall performed slightly better.

However, this increase in predictive power is gained solely from allowing the SAPs to proxy for the SUPs, thereby permitting variables which have been deemed to be socially unacceptable to influence the predictive process. Moreover, within a SUP group combination, using biased SAP coefficients to generate rankings causes distortions in the predictions. We present one way to demonstrate this point. First, to ground the comparison we hold fixed the number of people assigned to the workshop under the *common method* for each SUP group combination. This establishes a quota for each SUP combination. Then we compare the actual exhaustion rates of those sent under the *common method* to those that would have been sent under such a quota using the *proposed method*. The *proposed method* would exchange 5,020 claimants assigned in the *common method* with other individuals that have the same set of SUPs. We label these exchanges "switches," and present in Table 4 the number of total "switches" made for each SUP combination. The last two columns give the average fraction of benefits exhausted for "switchers" coming from the *common method* and those chosen to replace them by the *proposed method*. Because the *proposed method* has the same "accurate" within-SUP-combination rankings as the *full method*, we would anticipate such switching to improve predictive accuracy. Indeed, as expected the fraction of benefits actually exhausted by the *common method* "switchers" is lower than those from the *proposed method* – 79.7% versus 81.9%.

Implications for variable selection. Given that most of the implicit statistical discrimination arises from the correlation between race and many of the local unemployment offices, as an alternative to the *proposed method* one might consider eliminating these geographic controls from the statistical model (similar to arguments often made against the use of zip codes in insurance). This alternative solution would not only create potential problems with the remaining SAPs, but would also likely reduce the overall predictive accuracy of the model. For instance, if one were to implement the *common method* without controlling for local unemployment office, the actual fraction of benefits exhausted among the top 10% of predicted exhauster would be 79.3%. This compares to 83.2% for the *proposed method*, which maintains the geographic controls.

V. Conclusion

We address the question of how to eliminate the influence of socially unacceptable predictors in profiling models. Using data on unemployment insurance claimants, we show that simply omitting these sensitive characteristics from WPRS models, as is commonly done by states, does not eliminate their influence on predicted outcomes. Instead, the predicted exhaustion values attributed to individuals seem to be influenced to a considerable degree by variation in the racial composition of zip codes. We both employ in our analysis and argue in favor of the use in practical applications of a simple procedure that eliminates the distortions that are caused by allowing socially acceptable variables to proxy for the excluded predictors.

Future research may address the importance of implicit statistical discrimination in other markets where it is likely that variables such as race, gender, and age continue to influence predictions even though anti-discrimination policies prohibit their use. Credit scores, zip codes, and SAT scores are three predictors that have been identified as variables which potentially serve as proxies for socially unacceptable predictors. The approach in this paper provides a framework for identifying the extent to which these contentious variables derive their predictive power from their correlations to sensitive characteristics, and as such should prove useful for anyone interested in these debates.

References

Aigner, Dennis and Cain, Glen. "Statistical Theories of Discrimination in Labor Markets." *Industrial and Labor Relations Review*, 1977, *30*(2), pp. 175-87.

Arrow, Kenneth. "Some Mathematical Models of Race Discrimination in the Labor Market," in Anthony Pascal, ed., *Racial Discrimination in Economic Life.* Lexington: Lexington Books, 1972.

Arrow, Kenneth. "The Theory of Discrimination," in Orley Ashenfelter and Albert Rees, eds., *Discrimination in Labor Markets.* Princeton: Princeton University Press, 1973.

Balducchi, David, ed. *Worker Profiling and Reemployment Services (WPRS) Systems: National Colloquium, Selected Papers and Materials.* Washington, DC: Employment and Training Administration, U.S. Department of Labor, 1996.

Becker, Gary. *The Economics of Discrimination.* Chicago: University of Chicago Press, 1957.

Bertrand, Marianne and Mullainathan, Sendhil, "Are Emily and Brendan More Employable than Latoya and Tyrone? Evidence on Racial Discrimination in the Labor Market from a Large Randomized Experiment," *American Economic Review,* 2004, *94(4)*, pp 991-1013.

Black, Dan A., Smith, Jeffrey A., Berger, Mark C., and Noel, Brett J. "Is the Threat of Reemployment Services More Effective than the Services Themselves? Evidence from Random Assignment in the UI System." *American Economic Review*, 2003, *93*(4), pp 1313-27.

Black, Dan A., Smith, Jeffrey A., Plesca, Miana, and Shannon, Suzanne. "Profiling UI Claimants to Allocate Reemployment Services: Evidence and Recommendations for States." Employment and Training Administration, Department of Labor, 2003.

Borjas, George and Goldberg, Matthew. "Biased Screening and Discrimination in the Labor Market." *American Economic Review*, 1978, *68*(5), pp. 918-22.

Card, David and Levine, Phillip B. "Extended Benefits and the Duration of UI Spells: Evidence from the New Jersey Extended Benefits Program." *Journal of Public Economics,* 2000, 7*8*(1), pp. 107-38.

Card, David. "A Survey of State Worker Profiling Models." Unpublished Paper, 2003.

Chan, Jimmy and Eyster, Erik. "Does Banning Affirmative Action Lower College Student Quality?" *American Economic Review*, 2003, *93*(3):858-72.

Durlauf, Steven. "Assessing Racial Profiling." Unpublished Paper, 2004.

Fryer, Roland, Loury, Glenn C., and Yuret, Tolga. "Color-Blind Affirmative Action." NBER Working Paper # 10103, 2003.

Greene, William H. *Econometeric Analysis*, 5th ed. New Jersey: Prentice Hall, 2003.

Lundberg, Shelly J. "The Enforcement of Equal Opportunity Laws Under Imperfect Information: Affirmative Action and Alternatives." *Quarterly Journal of Economics*, 1991, *106*(1), pp. 309-26.

OECD. *Early Identification of Jobseekers at Risk of Long-term Unemployment*. Paris: Organization for Economic Co-Operation and Development, 1998.

Phelps, Edmund. "The Statistical Theory of Racism and Sexism." *American Economic Review*, 1972, *62*(4), pp. 659-61.

Risse, Mathias and Zeckhauser, Richard. "Racial Profiling." *Philosophy and Public Affairs*, 2004, *32*(2), pp. 131-70.

Ross, Stephen. "The Continuing Practice and Impact of Discrimination." University of Connecticut, Economics Working Paper 2005-19.

Ross, Stephen and Yinger, John. *The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement.* Cambridge: MIT Press, 2002.

Rothstein, Jesse. "SAT Scores, High Schools, and Collegiate Performance Predictions." Unpublished Paper, 2005.

Schwab, Stewart. "Is Statistical Discrimination Efficient." *American Economic Review*, 1986, *76*(1), pp. 228-234.

Tootell, Geoffrey M. "Redlining in Boston: Do Mortgage Lenders Discriminate Against Neighborhoods?" *Quarterly Journal of Economics,* 1996, *111*(4), pp. 1049-79.

Wandner, Stephen. "Early Reemployment for Dislocated Workers in the United States." *International Social Security Review,* 1997, *50*(4), pp. 95-112.

Wandner, Stephen and Messenger, Jon, eds. "Worker Profiling and Reemployment Services Policy Workgroup: Final Report and Recommendations." Washington, DC: U.S. Department of Labor, 1999.

Figure 1. Composition of Reemployment Services Class by Percentile Cutoff of Predicted Exhaustion

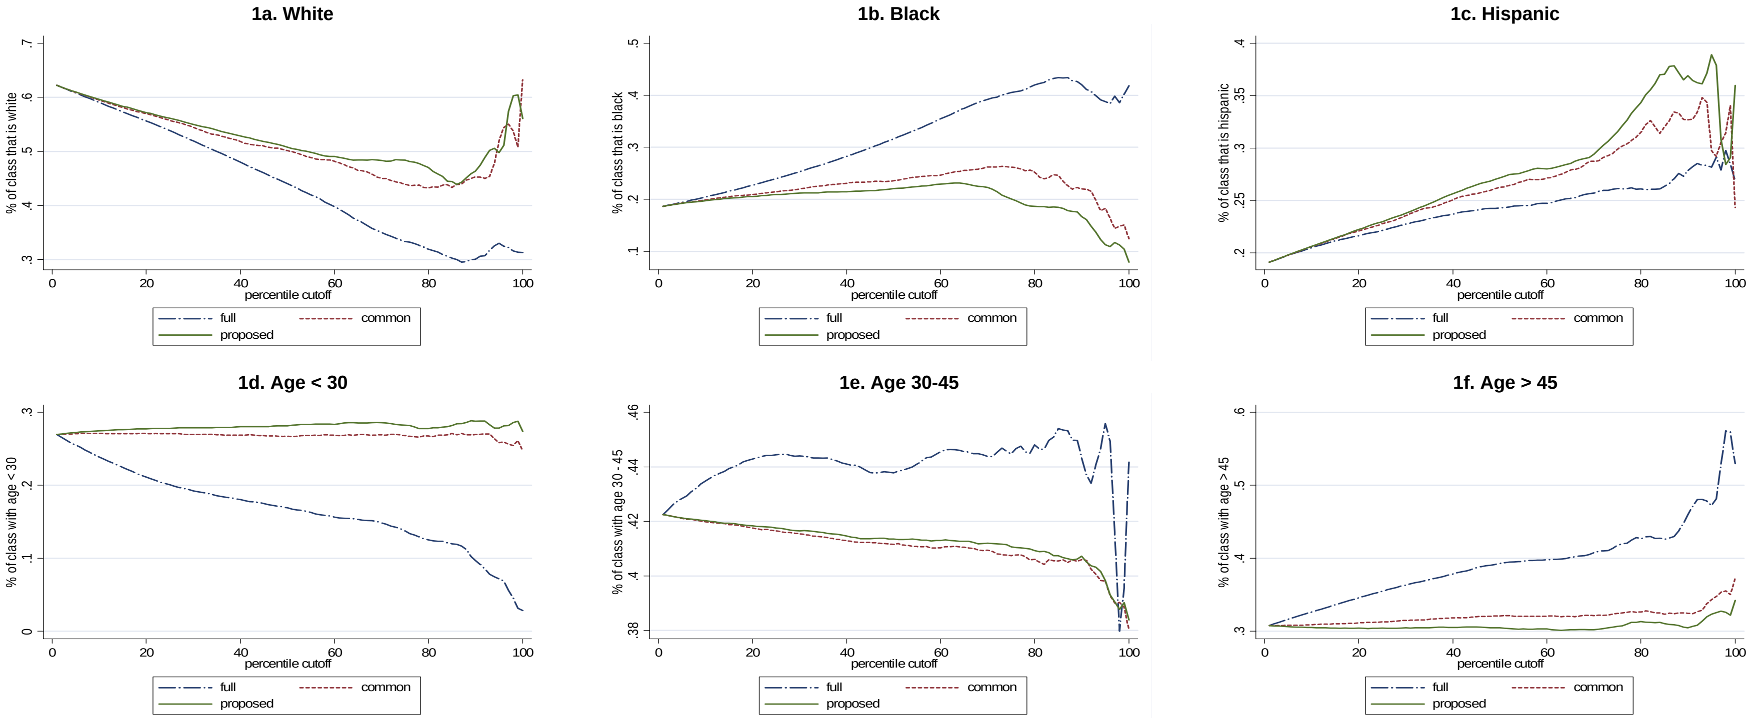

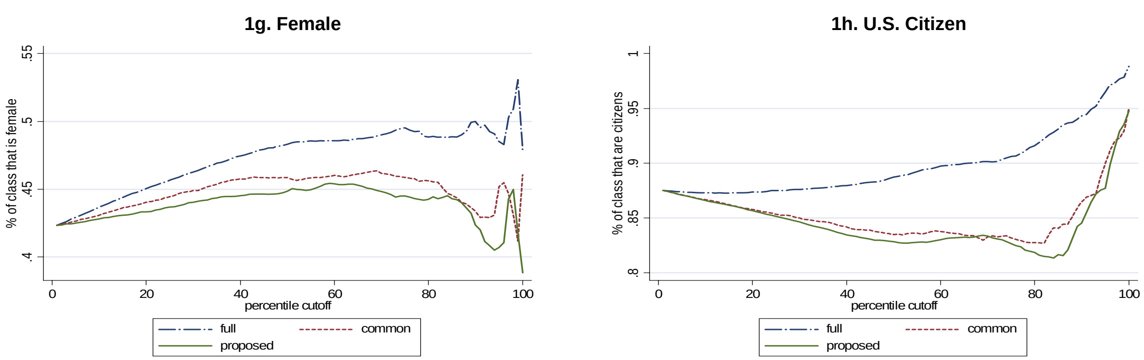

Notes: For any predicted exhaustion cutoff, the figures illustrate the composition of the people required to attend a reemployment services class. For instance if the predicted exhaustion cutoff is at the 80th percentile, the people with the highest 20% of predicted exhaustion are required to attend a workshop. The subfigures present a curve for each of the three methods discussed in the text (*full*, *common*, and *proposed*). For example, Figure 1f shows that when the predicted exhaustion cutoff is set at the 60th percentile, using the *full method* 40% of the class is older than 45.

Figure 2. Difference in the Percent of Each Racial Group Required to Attend Workshop Between the Common and Proposed Methods by Local Office (90%-cutoff Rule)

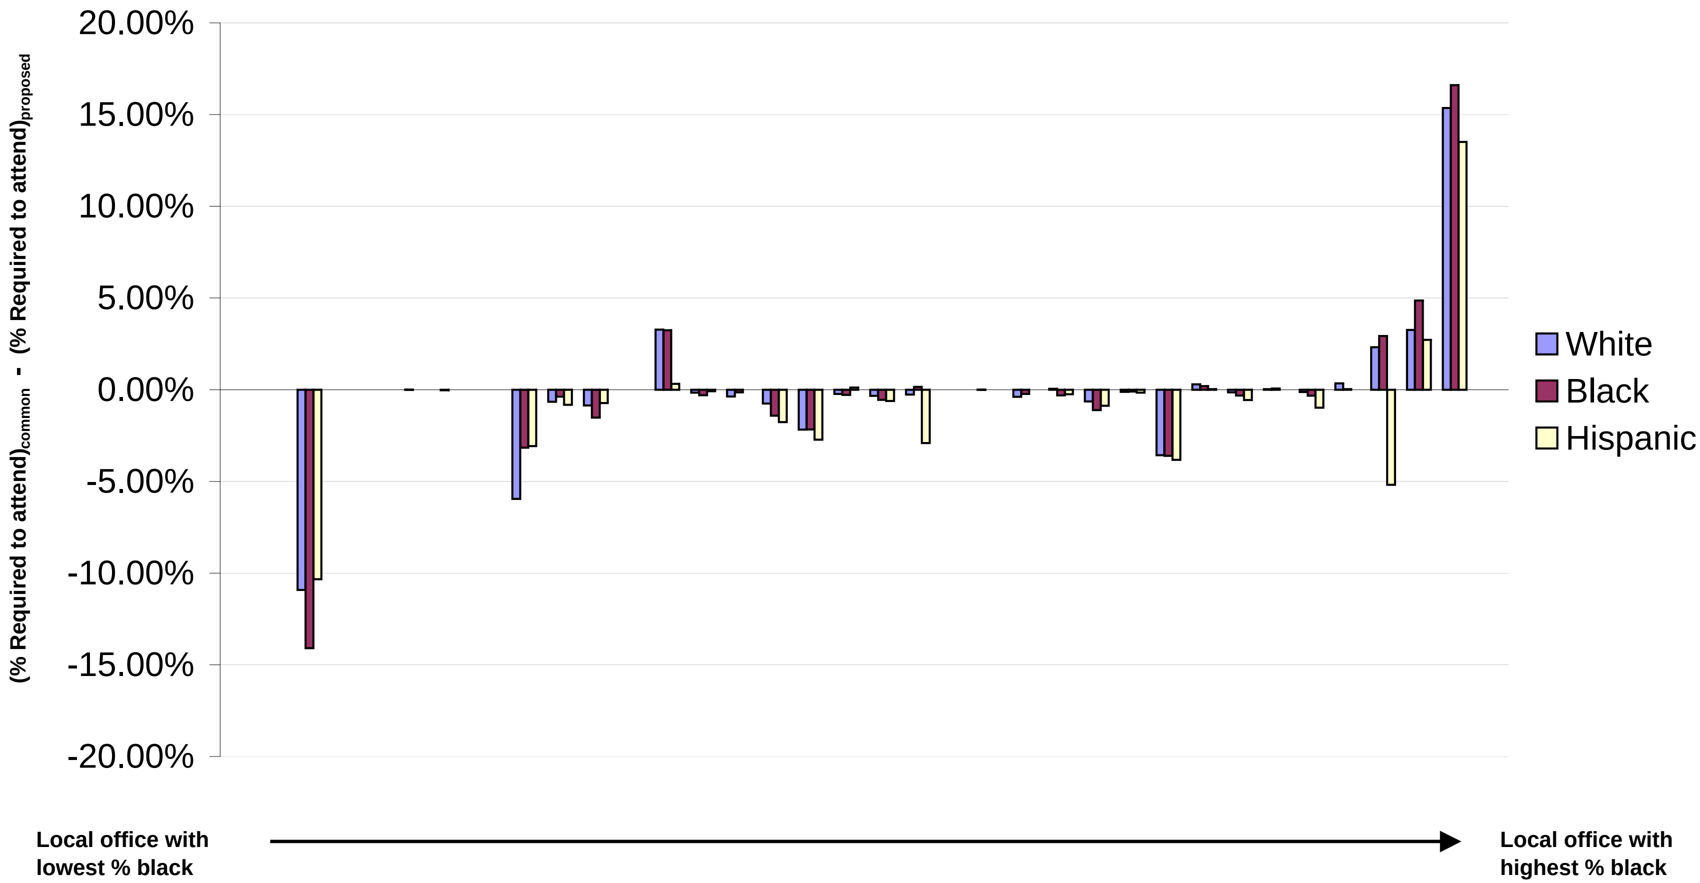

Notes: The x-axis indexes the local offices in order of the concentration of claimants from that office that are black. It is not scaled evenly by concentration black. For each racial group in each office, the fraction of claimants mandated to attend a workshop under a 90%-cutoff rule was calculated for each profiling method. The y-axis gives the difference in this measure between the *common* and the *proposed methods*.

Table 1. Summary Statistics

Variable	Full Sample	Race: White	Race: Black	Race: Hispanic	Age: < 30	Age: 30-45	Age: > 45	Gender: Female	Gender: Male	Citizenship: US	Citizenship: Non Us
Benefit Exhaustion Measures											
Fraction of benefits exhausted	0.74	0.72	0.78	0.76	0.71	0.74	0.76	0.75	0.73	0.74	0.73
Fraction who exhausted completely	0.48	0.44	0.56	0.50	0.42	0.48	0.52	0.50	0.46	0.48	0.45
SUPs											
Race											
White	0.62	1.00	0.00	0.00	0.56	0.60	0.70	0.63	0.62	0.68	0.24
Black	0.19	0.00	1.00	0.00	0.22	0.20	0.14	0.20	0.18	0.20	0.13
Hispanic	0.19	0.00	0.00	1.00	0.22	0.20	0.16	0.18	0.20	0.13	0.63
Age											
Age < 30	0.27	0.24	0.31	0.32	1.00	0.00	0.00	0.24	0.29	0.27	0.26
Age 30 - 45	0.42	0.41	0.46	0.43	0.00	1.00	0.00	0.41	0.43	0.42	0.47
Age > 45	0.31	0.35	0.23	0.25	0.00	0.00	1.00	0.35	0.28	0.31	0.27
Gender											
Female	0.42	0.43	0.44	0.40	0.38	0.41	0.48	1.00	0.00	0.43	0.40
Male	0.58	0.57	0.56	0.60	0.62	0.59	0.52	0.00	1.00	0.57	0.60
Citizenship											
US Citizen	0.88	0.95	0.92	0.59	0.88	0.86	0.89	0.88	0.87	1.00	0.00
Non US Citizen	0.12	0.05	0.08	0.41	0.12	0.14	0.11	0.12	0.13	0.00	1.00
SAPs											
Employment Information											
Pre-UI weekly wage ($100s)	4.9	5.5	4.1	3.7	3.9	5.2	5.5	4.2	5.4	5.1	3.9
Tenure (years)	0.95	1.06	0.84	0.71	0.48	0.86	1.50	1.00	0.90	0.99	0.69
Education											
Less than high school	0.19	0.13	0.17	0.43	0.17	0.17	0.25	0.17	0.21	0.15	0.48
High school diploma	0.46	0.47	0.48	0.38	0.47	0.46	0.44	0.46	0.45	0.48	0.31
Some college	0.21	0.22	0.25	0.13	0.25	0.22	0.16	0.23	0.19	0.22	0.13
College degree	0.14	0.18	0.09	0.06	0.11	0.15	0.14	0.14	0.14	0.15	0.08
Graduate degree	0.002	0.003	0.001	0.001	0.001	0.002	0.003	0.002	0.002	0.002	0.001
Industry											
Agriculture	0.05	0.04	0.02	0.09	0.07	0.05	0.02	0.01	0.07	0.04	0.06
Mining	0.001	0.001	0.0003	0.0004	0.001	0.001	0.001	0.0002	0.002	0.001	0.0002
Construction	0.09	0.11	0.04	0.06	0.1	0.1	0.07	0.02	0.14	0.09	0.11
Manufacturing	0.18	0.14	0.14	0.32	0.13	0.18	0.22	0.18	0.17	0.15	0.35
Transportation	0.07	0.07	0.09	0.05	0.06	0.07	0.08	0.06	0.08	0.08	0.04
Wholesale trade	0.08	0.09	0.07	0.09	0.08	0.09	0.08	0.08	0.09	0.08	0.09
Retail trade	0.16	0.17	0.15	0.11	0.19	0.14	0.14	0.17	0.15	0.16	0.11
Finance, insurance, real estate	0.06	0.06	0.06	0.03	0.05	0.06	0.06	0.08	0.04	0.06	0.02
Public administration	0.03	0.03	0.05	0.01	0.02	0.02	0.04	0.04	0.02	0.03	0.01
Services	0.29	0.28	0.39	0.22	0.29	0.29	0.28	0.36	0.23	0.3	0.21
Observations	590,924	367,794	110,239	112,891	159,327	249,693	181,904	250,156	340,768	517,185	73,739

Notes: The data includes the maximum number of weeks of benefits for which the claimant was eligible. The fraction of benefits exhausted for each claimant is computed by dividing the number of weeks of benefits used by the total number available. Graduate degrees include masters, professional, and doctoral degrees. The industry groups are based on the 1-digit, US Standard Industrial Classification (SIC) industry codes from1987.

Table 2: Predicting Exhaustion

Dependent Variable: Fraction of Benefits Exhausted			
Variable [%white, %black, %Hispanic]	**Full Model**	**Full Model (omitting SUPs)**	**Difference**
Black	0.0538**		
	(0.0012)		
Hispanic	0.0220**		
	(0.0014)		
30 < Age < 45	0.0392**		
	(0.0010)		
45 < Age	0.0559**		
	(0.0012)		
Female	0.0077**		
	(0.0009)		
U.S. Citizen	0.0340**		
	(0.0014)		
Prior weekly wage ($100s)	-0.0001	0.0002	-0.0003**
	(0.0002)	(0.0002)	(0.0001)
Tenure (yrs)	-0.0055**	0.0003	-0.0058**
	(0.0005)	(0.0005)	(0.0001)
Tenure squared	0.0007**	0.0004**	0.0003**
	(0.0001)	(0.0001)	(0.0000)
Less than HS	-0.0011	-0.0013	0.0002
	(0.0012)	(0.0012)	(0.0003)
Some college	-0.0088**	-0.0112**	0.0024**
	(0.0011)	(0.0011)	(0.0001)
College degree	-0.0349**	-0.0395**	0.0046**
	(0.0014)	(0.0014)	(0.0002)
Masters degree	-0.0541**	-0.0569**	0.0028**
	(0.0095)	(0.0095)	(0.0007)
Construction [79, 9, 12]	-0.0190**	-0.0359**	0.0169**
	(0.0017)	(0.0016)	(0.0005)
Public admin. [58, 35, 7]	-0.0221**	-0.0159**	-0.0062**
	(0.0026)	(0.0026)	(0.0003)
Local office #10 [9, 85, 6]	0.0555**	0.0843**	-0.0288**
	(0.0032)	(0.0031)	(0.0008)
Local office #12 [27, 41, 32]	0.0572**	0.0674**	-0.0102**
	(0.0029)	(0.0029)	(0.0005)
Local office #21 [97, 1, 2]	-0.0541**	-0.0609**	0.0068**
	(0.0040)	(0.0040)	(0.0004)
Local office #26 [97, 1, 1]	-0.0969**	-0.1016**	0.0047**
	(0.0073)	(0.0073)	(0.0006)
Other industry codes	X	X	
Other local offices	X	X	
Observations	590924	590924	
R-squared	0.034	0.026	

Notes: Coefficient values and robust standard errors are presented from the OLS regression using fraction of benefits exhausted as the dependent variable. White is the racial base-group. Age < 30 is the age basegroup. Having a high school diploma is the education basegroup. All local offices (#17 = basegroup) and all industry groups (services = basegroup) are included in the regression, but only six coefficients are presented. The brackets next to these variables presents the percent of white, black, and Hispanic claimants, respectively, for which each local office or industry group is composed. The third column provides the difference between the coefficient estimates. The standard errors on the differences were found using a bootstrap routine with 200 repititions.

***** significant at 5%; ** significant at 1%

Table 3. Composition of Claimants Required to Attend a Reemployment Workshop

		90%-cutoff Rule					80%-cutoff Rule				
					Difference (Common - Proposed)					Difference (Common - Proposed)	
	Full Sample	Full	Common	Proposed	Mean	SE+	Full	Common	Proposed	Mean	SE+
Race											
% White (not Hispanic)	62.2	30.5	45.2	47.6	-2.4	0.6**	31.6	43.4	46.4	-3.0	0.4**
% Black (not Hispanic)	18.7	41.3	21.9	16.0	5.9	0.5**	42.3	24.5	18.6	5.9	0.6**
% Hispanic	19.1	28.2	32.9	36.4	-3.5	0.8**	26.1	32.1	35.0	-2.9	0.4**
Gender											
% Male	57.7	50.3	57.0	58.1	-1.1	0.3**	51.1	54.5	55.7	-1.2	0.2**
% Female	42.3	49.7	43.0	41.9	1.1	0.3***	48.9	45.5	44.3	1.2	0.2**
Age											
% < 30	27.0	9.1	26.9	28.8	-1.9	0.2**	12.4	26.8	27.9	-1.1	0.2**
% 30 - 45	42.3	43.8	40.5	40.5	0.0	0.1	44.7	40.5	40.9	-0.4	0.1**
% 45 +	30.8	47.1	32.6	30.7	1.9	0.2**	42.9	32.7	31.2	1.5	0.2**
Citizenship											
% U.S. Citizen	87.5	94.5	86.8	85.5	1.3	0.5*	91.9	82.8	81.7	1.1	0.2**
% Not U.S. Citizen	12.5	5.5	13.2	14.5	-1.3	0.5*	8.1	17.2	18.3	-1.1	0.2**

Notes: For each cutoff rule and profiling method, the values in the table give the proportion of the reemployment workshop that have a given SUP characteristic.

+ Standard errors are generated using a bootstrap routine with 200 repetitions and 100% samples with replacement.

* significant at 5%; ** significant at 1%

Table 4. Comparing Proposed vs. Common Methods Within SUP Groups

SUP Matrix					Sent to Class by Method (90% Cutoff)				Fraction of Benefits Exhausted by "Switchers"		
Citizenship	Race	Gender	Age	Total Number	Only Common	Only Proposed	Both	Neither	Total "Switches"	Only Common	Only Proposed
U.S.	White	Male	Age < 30	54,787	**82**	541	3,198	50,966	82	0.7452	0.7652
U.S.	White	Male	Age 30 - 45	82,876	**266**	605	5,083	76,922	266	0.7736	0.8063
U.S.	White	Male	Age > 45	62,289	461	**386**	4,433	57,009	437	0.8109	0.8321
U.S.	White	Female	Age < 30	31,249	**72**	437	2,258	28,482	72	0.7690	0.7090
U.S.	White	Female	Age 30 - 45	59,696	**264**	527	4,308	54,597	264	0.8256	0.8344
U.S.	White	Female	Age > 45	58,962	555	**451**	4,873	53,083	525	0.8342	0.8343
U.S.	Black	Male	Age < 30	16,797	446	**187**	1,412	14,752	288	0.8185	0.8077
U.S.	Black	Male	Age 30 - 45	26,424	806	**222**	2,040	23,356	471	0.8424	0.8428
U.S.	Black	Male	Age > 45	12,507	495	**55**	830	11,127	254	0.8203	0.8468
U.S.	Black	Female	Age < 30	14,866	456	**135**	1,286	12,989	285	0.7891	0.8067
U.S.	Black	Female	Age 30 - 45	19,279	782	**101**	1,488	16,908	366	0.8055	0.8362
U.S.	Black	Female	Age > 45	10,996	621	**32**	811	9,532	213	0.7447	0.8196
U.S.	Hispanic	Male	Age < 30	14,550	**101**	378	3,357	10,714	101	0.7693	0.7803
U.S.	Hispanic	Male	Age 30 - 45	16,852	**147**	322	3,751	12,632	147	0.8491	0.8596
U.S.	Hispanic	Male	Age > 45	9,649	**133**	190	2,317	7,009	133	0.7955	0.8196
U.S.	Hispanic	Female	Age < 30	7,878	**65**	235	1,273	6,305	65	0.8513	0.8455
U.S.	Hispanic	Female	Age 30 - 45	9,882	**136**	218	1,426	8,102	136	0.8125	0.8216
U.S.	Hispanic	Female	Age > 45	7,646	**177**	206	1,143	6,120	177	0.7541	0.8098
Non U.S.	White	Male	Age < 30	2,350	**8**	22	72	2,248	8	0.8083	0.9038
Non U.S.	White	Male	Age 30 - 45	5,358	**21**	48	210	5,079	21	0.8031	0.7717
Non U.S.	White	Male	Age > 45	3,636	31	**27**	167	3,411	28	0.8634	0.8605
Non U.S.	White	Female	Age < 30	1,098	**8**	17	69	1,004	8	0.8380	0.9142
Non U.S.	White	Female	Age 30 - 45	2,741	**21**	27	129	2,564	21	0.6304	0.8873
Non U.S.	White	Female	Age > 45	2,752	**27**	28	124	2,573	27	0.8661	0.8802
Non U.S.	Black	Male	Age < 30	1,513	85	**20**	116	1,292	45	0.7951	0.7341
Non U.S.	Black	Male	Age 30 - 45	2,686	135	**25**	227	2,299	67	0.7505	0.8157
Non U.S.	Black	Male	Age > 45	1,343	85	**17**	111	1,130	49	0.7006	0.7972
Non U.S.	Black	Female	Age < 30	1,100	80	**15**	85	920	32	0.7029	0.6895
Non U.S.	Black	Female	Age 30 - 45	1,858	193	**17**	157	1,491	60	0.5648	0.7688
Non U.S.	Black	Female	Age > 45	870	121	**5**	111	633	30	0.6573	0.7328
Non U.S.	Hispanic	Male	Age < 30	8,201	**35**	334	855	6,977	35	0.8587	0.7689
Non U.S.	Hispanic	Male	Age 30 - 45	13,074	**86**	441	1,332	11,215	86	0.6897	0.7685
Non U.S.	Hispanic	Male	Age > 45	5,876	**55**	194	745	4,882	55	0.6765	0.7550
Non U.S.	Hispanic	Female	Age < 30	4,938	**35**	219	510	4,174	35	0.7234	0.7983
Non U.S.	Hispanic	Female	Age 30 - 45	8,967	**70**	339	896	7,662	70	0.7780	0.7880
Non U.S.	Hispanic	Female	Age > 45	5,378	**61**	199	667	4,451	61	0.7538	0.8208
Sum				590,924	7,222	7,222	51,870	524,610	5,020		
Average										0.7966	0.819
										(0.0043)	(0.0041)

Notes: The table lists the total number of claimants within each of 36 distinct SUP-group combinations. The subsequent four columns indicate how many of these individuals are assigned to attend a workshop under just the *common method*, just the *proposed method*, under both methods, and in neither method given a 90%-cutoff rule. The shaded numbers show the number of individuals sent from one of the two methods even though that method sends fewer claimants from that SUP combination overall. We define "switches" as follows: For each SUP combination, we establish a quota equal to the number of claimants sent to a workshop under the *common method.* The number of "switches" is the number of claimants that would be replaced using the *proposed method* with such a quota. The last two columns give the average fraction of benefits exhausted for “switchers” coming from the *common method* and those chosen to replace them by the *proposed method*. Parantheses denote standard errors.