

NBER WORKING PAPER SERIES

TESTING THEORIES OF DISCRIMINATION:  
EVIDENCE FROM *WEAKEST LINK*

Steven D. Levitt

Working Paper 9449

<http://www.nber.org/papers/w9449>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

January 2003

I would like to thank Mark Duggan, Roland Fryer, Austan Goolsbee, and Doug MacAdam for helpful discussions on the topic, as well as Jeannette Levitt for providing the raw data. Carol Klein and Katherine Owens performed truly exemplary research assistance. Part of this paper was written while I was a fellow in residence at the Center for Advanced Studies in Behavioral Sciences at Stanford, CA. Financial support provided by the National Science Foundation. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

©2003 by Steven D. Levitt. All rights reserved. Short sections of text not to exceed two paragraphs, may be quoted without explicit permission provided that full credit including notice, is given to the source.

Testing Theories of Discrimination: Evidence from *Weakest Link*  
Steven D. Levitt  
NBER Working Paper No. 9449  
January 2003  
JEL No. J71

### **ABSTRACT**

In most settings, it is difficult to measure discrimination, and even more challenging to distinguish between competing theories of discrimination (taste-based versus information-based). Using contestant voting behavior on the television game show *Weakest Link*, one can in principle empirically address both of these questions. On the show, contestants answer questions and vote off other players, competing for a winner-take-all prize. In early rounds, strategic incentives encourage voting for the weakest competitors. In later rounds, the incentives reverse, and the strongest competitors become the logical target. Controlling for other observable characteristics including the number of correct answers thus far, both theories of discrimination predict that in early rounds, excess votes will be made against groups targeted for discrimination. In later rounds, however, taste-based models predict continued excess votes, whereas statistical discrimination predicts fewer votes against the target group. Empirically, I find some evidence of information-based discrimination towards Hispanics (i.e., other players perceive them as having low ability) and taste-based discrimination against older players (i.e., other players treat them with animus). There is little in the data to suggest discrimination against women and Blacks.

Steven D. Levitt  
Department of Economics  
University of Chicago  
1126 East 59th Street  
Chicago, IL 60637  
and NBER  
slevitt@midway.uchicago.edu

Discrimination features prominently in American history. The leading target of discrimination has been Blacks, whether in the form of slavery, Reconstruction, Jim Crow, or other guises. Asians have also been the target of severe discrimination historically, as have women. In modern society, the extent of discrimination is less clear. Some data suggest that discrimination may still be at work. For instance, large racial gaps remain on many social indicators such as income, education, and life expectancy (Blank 2001), and women continue to earn far less than men on average in the labor force. On the other hand, legislation now provides legal protection against discrimination on the basis of race, gender, and age. Along some dimensions, discrimination appears less pronounced. For instance, between 1970 and 1990 the number of Black-White interracial marriages more than tripled. In addition, women and minorities have increasingly come to occupy prominent positions in politics, business, and the media. Viewed in this light, the current racial and gender economic gaps may be the consequence of past discrimination, rather than a reflection of current discriminatory practices.

To the extent that discrimination persists, understanding the source of that discrimination is of first-order importance. There are two leading theories of discrimination (see Fryer 2001 for a detailed survey). The first theory is based on tastes, and originates with Becker (1957). In the taste-based story, some economic actors prefer not to interact with a particular class of people and are willing to pay a financial cost to avoid such interactions. The other leading explanation is based on incomplete information. The simplest information-based model involves one group having mistaken beliefs about another group's skill level and acting accordingly. That simple model, while perhaps a reasonable description of behavior, is not a very satisfying economic model because it implies that individuals are making systematic errors. A series of more sophisticated information-based statistical discrimination models circumvent that criticism (Phelps

1972, Arrow 1973, Aigner and Cain 1977, Lundberg and Startz 1983, Coate and Loury 1993, Fryer and Jackson 2002). In these models, individuals (typically employers) discriminate against particular groups either because (1) signals of ability are less informative within that group or, (2) in the presence of human capital investment, equilibria exist in which negative prior beliefs about members of a particular group become self-fulfilling. In models of statistical discrimination, economic actors have no animus (unlike taste-based models), but discriminatory outcomes nonetheless arise.

Measuring the extent of discrimination poses a difficult empirical challenge.<sup>1</sup> Self-report data are unlikely to accurately reflect attitudes if there is a perceived stigma attached to racist views.<sup>2</sup> A number of different approaches have been employed in an attempt to address this question. One method, known as the “audit study” uses matched pairs of individuals of different races who masquerade as consumers or job hunters.<sup>3</sup> Using this methodology, for instance, Ayres and Siegelman (1995) find that car dealerships attempt to charge higher prices to Blacks and women. Bertrand and Mullanaithan (2002) find that resumes carrying distinctively Black names are less likely to receive job interviews. A second approach to the empirical study of

---

<sup>1</sup> For a more complete survey of the empirical literature on discrimination than is presented here, see Altonji and Blank (1999).

<sup>2</sup> As one vivid example, Howell and Sims (1994) analyze the Louisiana Gubernatorial election in which White supremacist politician David Duke received far more votes than pre-election surveys and exit polls had predicted. Kuklinski et al. (1997), using a clever methodology designed to minimize self-reporting issues, find no evidence of discrimination towards Blacks outside the South, but substantial racism among White southerners.

<sup>3</sup> Heckman and Siegelman (1992) criticize audit studies on the grounds that the simulated discriminatory transactions that occur in audit studies would be unlikely to arise in a market economy, i.e. Blacks may not shop for cars at all-White car dealerships because they fear being discriminated against. If that is the case, then in the real world, the consequences of discrimination will be far less than suggested by the audit study.the

discrimination is to compare salaries to marginal products of labor for Blacks and Whites. These studies are most commonly performed for athletes, where salaries are known and “output” is relatively easy to quantify (e.g., Kahn 1991). Mixed evidence of discrimination is found in this literature, with some studies finding salary and customer discrimination against Blacks, particularly in basketball. Knowles, Persico, and Todd (2001), analyzing outcomes of drug searches following police stops, find no evidence of racial bias. In general, empirical tests have a difficult time distinguishing between taste-based and information-based models of discrimination.<sup>4</sup>

In this paper, I try to both measure the extent of discrimination and distinguish between competing theories using an unusual data source: contestant behavior on the television game show “Weakest Link.”<sup>5</sup> On this show, contestants answer trivia questions over a series of rounds, with one contestant eliminated each round based on the votes of the other contestants, until only two contestants remain. The last two contestants compete head-to-head for the winner-take-all prize. Because the prize money at stake is large (as much as \$190,000 on a single show), participants have powerful incentives to vote in a manner that maximizes their chance of winning.

The total prize money is an increasing function of the number of questions that are answered correctly over the course of the program. As a consequence, there is a strong incentive

---

<sup>4</sup> Altonji and Pierret (2001) present a clever methodology for ferreting out statistical discrimination using information about how wages change with employment tenure. They find little evidence of statistical discrimination on the basis of race.

<sup>5</sup> After the first draft of this paper was written, it came to my attention that another group of researchers were independently collecting *Weakest Link* data to answer questions about discrimination (Antonovics, Arcidiacono, and Walsh 2002). Their research also concludes that there is little evidence of discrimination towards Blacks and women. They do not analyze discrimination towards Hispanics or the elderly.

See Gertner (1993) and Metrick (1995) for other economic analyses using game shows as a vehicle.

to vote off the least skilled players in the early rounds.<sup>6</sup> Both taste-based and statistical discrimination theories would predict that discrimination will manifest itself as an increased propensity to vote off the group that is discriminated against early in the game, conditional on other observables such as the number of questions answered correctly up to that point in the show, or the individual's education (which may predict future skill in answering questions).<sup>7</sup>

As the end of the show nears, however, strategic incentives switch. The value of building the prize pool becomes outweighed by the question of whether a contestant can beat another contestant in a head-to-head challenge. Ideally, one would like to be competing against a low-skill opponent in the final round to increase one's chance of winning the final prize. In the late rounds, the two theories discussed above offer different predictions about how discriminators will behave. In the taste-based models, those who discriminate towards a group do so not because they think that group is less talented, but rather, because they do not like them. Thus, one would expect that a taste-based discriminator would continue to disproportionately vote against the target group in the late rounds. In information-based models, on the other hand, discriminators perceive the other group as less qualified. Thus, in the late rounds, such discriminators will *avoid* voting for members of the group that is viewed as less qualified, in order to raise the probability that the final

---

<sup>6</sup> The set of information about other competitors is fairly limited. In addition to observing the performance of each contestant in answering questions up to that point in the game, players also announce their age, occupation, and home town. The contestants can also see one another, and therefore can observe race, gender, and other visual cues, such as degree of physical attractiveness.

<sup>7</sup> Statistical discrimination models arising as a consequence of minorities having noisier signals will not necessarily generate this pattern. In those models, noisy signals result in minorities being less likely to be perceived as being in the tail of the distribution (both upper and lower). As a consequence, minorities may be *less* likely to be voted off in early rounds in such a model.

opponent will be from that group, increasing the perceived likelihood that the discriminator will win the prize. Consequently, voting behavior in the late rounds provides a unique opportunity to empirically distinguish between the two competing models.<sup>8</sup>

There are a number of important caveats concerning the applicability of these game show patterns to everyday behavior. First, the setting examined is not a market. Unlike in the real world, the contestants have little choice with respect to whom they are interacting. Second, the individuals who appear on this game show are highly selected, both with respect to who applies to be on the show and who is chosen to be on the show. Little is known about the precise nature of the selection process. It is clear, however, that those who appear on the show are not representative of the underlying populations. For instance, the education levels of Whites, Blacks, and Hispanics appearing the show are nearly identical, whereas there are big differences in mean education levels for these groups in the overall population. Third, the voting is taking place in front of a televised audience. Just as racist individuals lie about whether they voted for David Duke in an exit poll, they may be loath to broadcast racist views on a television game show. Finally, the decision about whom to vote for depends not just on one's own views, but also about one's beliefs as to how other contestants will vote. Race may serve as a focal point in a setting in which multiple equilibria exist, exaggerating the amount of measured discrimination. Thus, someone who is not racist, but believes others are racist, will have incentives to vote for Blacks in the early rounds.

With those important caveats in mind, the empirical findings reveal little systematic

---

<sup>8</sup>In either model of discrimination, there is another potential factor at work: losing to someone in the target group may be especially embarrassing, increasing the incentive to eliminate members of that group towards the end of the game.

evidence of discrimination towards Blacks, Asians, or women. The number of votes cast against these groups are similar to the numbers for White males. Hispanics and the elderly, however, do appear to face systematic discrimination on the show. These two groups consistently receive more votes (which is a bad thing) early in the show than other contestants, even after controlling for other factors such as education level and performance in answering questions up to that point in the game. The data are consistent with an information-based discrimination story for Hispanics. In the final round, Hispanics are significantly *less* likely to receive votes – a reversal from earlier rounds. For the elderly, the discrimination appears to be taste based. Even in the final round, the elderly receive significantly more votes than other contestants, controlling for relevant factors. There is also some tendency in the data for individuals to vote less frequently for members of their own group. For instance, women vote more frequently for men and vice-versa. Blacks tend to vote less for other Blacks than would be expected. The exception to this rule is the elderly, who are *more* likely to vote for other elderly than are contestants in general.

The remainder of the paper is organized as follows. Section II provides greater background on the television show *Weakest Link*. Section III discusses in more detail the strategic considerations players face. Section IV presents the empirical results. Section V concludes.

## Section II: Background on *Weakest Link*

*Weakest Link* is a television game show in which contestants compete against one another to obtain a winner-take-all prize. Each round, contestants take turns answering trivia questions, with the goal to build a “chain” of answers. The more consecutive correct answers given, the greater the prize. Prize money builds slowly for the first few correct answers and much more



quickly when many consecutive correct answer have been assembled.<sup>9</sup> Before hearing a question, the contestant has the opportunity to “bank” the money in the current chain. If the contestant chooses to “bank,” then the money assembled in the current chain is added to the final prize pool and the team starts over in building a chain. An incorrect answer also causes the chain-building to start over, but all of the money assembled in the current chain is lost, rather than added to the final prize money.<sup>10</sup> At the conclusion of each round, contestants secretly record the player who they would like to vote off. There is no communication between players after the conclusion of questions and prior to the casting of the votes. The votes, along with the identities of the individuals casting the votes, are then revealed. The competitor who receives the most votes is eliminated from the show (the “weakest link”), no longer answering questions and ineligible for the final prize. In the case of a tie, the contestant who had the most correct answers that round (the “strongest link”) determines which of the players with the most votes will be eliminated. This process continues until only two contestants remain, at which point they play one more round as a team trying to raise the final prize (the only round in which no player is voted off), and then compete head-to-head in a final round to determine who keeps the prize. The winner is the contestant that answers the most questions correctly in the final round. If the players are tied at the end of the final round, questions are asked until one player provides a correct answer and the other

---

<sup>9</sup> On the prime time show, the first correct answer is worth \$1,000, the next answer \$2,500, and the progression continues: \$5,000, \$10,000, \$25,000, \$50,000, \$75,000, \$125,000. On the day time show, the value of answers is as follows: \$250, \$500, \$1,000, \$2,500, \$5,000, \$12,500.

<sup>10</sup> The tradeoff involved in the decision to “bank” is that an additional correct answer has a bigger payoff if one does not “bank,” but the cost of an incorrect answer is also greater. Examining contestant behavior regarding the “bank” decision is of potential interest, but my data set does not contain information on who banked.

an incorrect answer. The time devoted to answering questions decreases each round.

Two different versions of the show are included in the data set used in this paper. The program originally aired as a one-hour prime time show involving eight competitors, eight rounds (seven in which one player is eliminated, and the penultimate round in which no players are voted off), and a theoretical maximum total prize of \$1 million. Later, the show was transformed into a syndicated 30 minute daytime show with six competitors, six rounds, and a theoretical maximum total prize of \$75,000.<sup>11</sup> In practice, the total prize money earned is well below the theoretical maximum. The median payout is roughly \$80,000 in the prime time version, and \$10,000 in the day time version.

The data were collected by video recording the televised programs and transcribing the results. The data set includes almost every prime time show except for those involving celebrities (e.g. a Brady Bunch reunion show) and virtually every day time show aired prior to January 2003.<sup>12</sup> There are a total of 25 prime time shows and 136 day time shows. Each prime time show yields 8 person-level observations, 33 person-votes (eight in the first round, seven in the second round, ..., 3 in the 6<sup>th</sup> round). Each day time show yields 6 person-level observations, 18 person-votes. Thus, in total, the data set includes 1,016 person-level observations and 3,273 person-votes.

---

<sup>11</sup> Toward the end of the sample, the round in which the last two remaining competitors play cooperatively to build the pool was eliminated. Because there is no voting in that round, it is not used in my analysis, so the change in format is immaterial.

<sup>12</sup> The prime time show no longer airs. Data collection is ongoing on the day time show and future versions of this paper will be updated to reflect the additional data.

A few shows are missing due to human error regarding recording of the shows, but data for every show that was successfully recorded are in the sample. There is no reason to believe that the missing shows would exhibit any systematic bias since the outcomes of the shows are not known at the point when the failure to record the show occurs.

Table 1 presents summary statistics for the data set. The top panel of the panel reports data at the contestant level. Roughly 20 percent of the contestants are Black. Asians and Hispanics represent a much smaller fraction of the players (2.9 and 2.2 percent respectively).<sup>13</sup> Players are almost evenly distributed between males and females. Approximately 7 percent of the contestants are fifty years or older (the mean age is 34). Contestant self-reports of occupation were crudely categorized according to the likely education level of the players: high school, college, professional school, doctorate, still in school, and unknown (assigned to those contestants whose occupation is missing, unemployed, retired without specifying an earlier occupation, or stay-at-home parent). Roughly one-third of the contestants perform jobs that require no more than a high-school education (although many of them may nonetheless have higher levels of education). Over forty percent are classified as college educated, with an additional eight percent having professional degrees (e.g. law) and another 3.6 percent with doctorates (PhD or MD).

The second panel of Table 1 reports game statistics at the level of the contestant-round. These statistics are presented for the show as a whole, as well as for early, middle, and late rounds of each program.<sup>14</sup> Overall, each contestant is asked between two and three questions per round, with that number rising in later rounds. The success rate for answering questions is slightly above 60 percent.

---

<sup>13</sup> Because of some ambiguity in visual assessments of race, the category labeled “Hispanic” is actually a catch-all category for non-Black, non-White, non-Asian contestants which includes a small number of Native Americans and Pacific Islanders, as well as Hispanics.

The producers of the show appear to have made a concerted effort to include minorities on each show. Of 161 shows in the sample, only four lack a minority contestant.

<sup>14</sup> Early rounds correspond to the first two rounds of the prime time show (one hour long) and the first round of the day time show (30 minutes long). Middle rounds correspond to rounds 3-5 of the prime time show and rounds 2 and 3 of the day time show. Final round corresponds to round 6 of the prime time show and round 4 of the day time show.

Also reported in Table 1 are statistics reporting the deviation of that player from other contestants on the same show in the percent of questions answered correctly this round and cumulatively up until that point in the show, as well as the number of votes for a contestant. The means of these variables are not themselves interesting. By definition the mean deviation from other players on the episode is equal to zero. Similarly, since each player casts one vote in each round, the mean number of votes must be equal to one. The standard deviations of these variables, however, will be of use in interpreting the regression estimates.

### Section III: Strategic considerations of contestants

The primary strategic decision that a player faces is for whom to vote.<sup>15</sup> The complexity of the situation precludes a formal model. Within each round there are as many as eight contestants and as many as seven possible targets for whom to vote. The optimal action for any one player depends critically on the beliefs about how other players will vote. There is nothing *a priori* that suggests that pure strategy equilibria should be the norm. Most critically, voting behavior in every round except the first will be a function not only of observable characteristics and performance thus far in the game, but also of past voting. Given that players have perhaps 30 seconds to determine their votes after the completion of each round, it is implausible that they could perform a rigorous optimization anyway. Rather, they must rely on rules of thumb or subjective criteria in determining their votes.

Absent a formal model, it is nonetheless possible to map out four broad considerations that arise in considering how to vote. The first consideration relates to building the size of the prize

---

<sup>15</sup> As noted earlier, the question of whether or not to “bank” also poses a question of decision making under uncertainty, but not one that is addressed in this paper.

pool. Players have a collective interest in having as many correct answers as possible in order to generate the biggest prize. Thus, one would like to vote off players who give correct answers with low probability. The minimum expected value of a correct answer, assuming optimal banking strategies, is the lowest payoff for a correct answer: \$1,000 in the prime time show and \$125 in the day time show. Back-of-the-envelope calculations suggest that the actual expected value of a correct answer will be at least three times higher, since rewards rise substantially when consecutive correct responses are tallied.<sup>16</sup> Players who survive to the final round are on average asked approximately 24 questions total. Thus, differences in ability can have a substantial impact on the final pool: thousands of dollars in the day time game and tens of thousands of dollars in the prime time game. The incentive to vote off poor players because of this channel is greatest in the earliest rounds since later in the game there are fewer questions remaining to be answered, so the opportunity cost of having a bad player alive decreases proportionally.

The second consideration when voting is the desire to eliminate strong players in order to maximize the likelihood that you will win in the head-to-head contest in the final round. On the prime time show, the final round consists of five questions per player. A contestant with a 20 percentage point edge in answering questions (e.g. 70 percent correct answers versus an opponent

---

<sup>16</sup>The precise value of a correct answer varies depending on whether the preceding and following questions are answered correctly and how frequently players “bank.” Assuming that players are risk neutral (see Metrick 1995) and bank money optimally, a lower bound on the expected value of a correct answer can be arrived under the assumption contestants never bank: they simply answer questions until they complete the chain (at which time the money is automatically banked) or until a wrong answer is given. Assuming that the average probability of a correct answer is .60, the expected value of a correct answer is roughly \$3,500 in the prime time show and \$900 in the day time show. For the prime time show, the payoff for eight straight correct answers is \$125,000, and the probability that any one answer is pivotal is  $1/6^7$ ; for the day time show, the payoff to six straight correct answers is \$12,500, and the probability that any one answer is pivotal is  $1/6^5$ . If banking is done optimally, a risk-neutral player will only bank if it increases the expected value of the pool, making this a lower bound.

with 50 percent correct) will win the final round about 80 percent of the time. On the day time version, the final round is only three questions which introduces more randomness, but the skilled player continues to win the great majority of the time. Moreover, empirically, performance in the early rounds of the show is a powerful predictor of who will answer questions correctly in the final round: the player with the higher fraction of questions correct going into the final round wins 64 percent of the time. Thus, towards the end of the game, the desire to eliminate strong competitors becomes paramount in importance.

The third consideration relates to the impact that your vote has on other's voting behavior towards you. In a given round, it is dangerous to vote against players who have answered many questions correctly and thus will be the "strongest link" and cast the deciding vote in the case of a tie. Given the speed with which questions are asked, as well as some apparent randomness in who is declared the strongest link on the show, it is difficult for contestants to know precisely who the strongest link is in any one round.<sup>17</sup> Thus, in general there is incentive to shy away from voting for players who have done well in this round, aside from the first two considerations discussed in the preceding paragraphs. In addition, in describing their motivations for why players vote for one another, the fact that another contestant voted for them in a previous round is frequently mentioned. Therefore, to the extent that one is successful in voting for players who receive many votes, the contestant increases his or her likelihood of survival. This is both because the target of the votes is less likely to remain alive in the next round, and because even if the target survives, if he or she received multiple votes, retribution cannot be delivered to all of those who targeted him or her in

---

<sup>17</sup> In the frequent cases in which multiple competitors tie for having the most correct answers in a given round, the next criterion used is the percent of answers correct in that round (i.e. players who were asked more questions are penalized). When a tie remains, the allocation rule used is unclear.

the previous round. Relative to the other two channels, it is less clear how this final consideration varies by round. It is likely, however, that this set of concerns are of reduced importance in the final round of voting, since there will be no later opportunity for the party voted against to exact revenge. In light of the tie-breaking power of the strongest link, incentive to avoid the player who answered the most questions correctly persists in the final round (but only for the two contestants who are not the strongest link in the last round).

The final consideration is the pressure applied by the show's host to vote off the weakest link. Before each vote, the host exhorts the players to eliminate the weakest link. Players who vote for strong competitors rather than the weakest link are often singled out by the host in the repartee that follows each vote. That sort of attention is unwanted because it provides a possible focal point for other players to coordinate their votes in the next round. This message is most pronounced just before the final vote, in which the host, acknowledging the importance of the second consideration above in the late rounds, implores the contestants to "have the courage" to vote off the weakest link.

Given the informality of the discussion above, the empirical predictions one can derive are somewhat circumscribed. Ignoring for the moment the third and fourth consideration (repercussions of one's vote and attempts at moral suasion), the prize-building channel unambiguously weakens over time and the weak-final-round-opponent channel unambiguously grows over time. Both of these factors, therefore, point to an increasing tendency to vote against stronger players over time. At least in the final round, the third consideration (repercussion of one's vote) also would suggest an increased likelihood of voting against a strong player. Only the fourth consideration serves to moderate this tendency. Thus, it seems plausible that in the early rounds, the optimal strategy is slanted towards voting off weak players, whereas at the end (at

least in relative terms), the incentives to vote off strong players are greater.

The issue of whether weak or strong players are the target is critical to differentiating taste-based and information-based models of discrimination. In a taste-based model, the discriminatory behavior is not predicated on beliefs about the target of discrimination's talent level. Thus, one would expect a tendency to vote against the discriminated group throughout all rounds. In an information-based model, on the other hand, the discrimination arises from beliefs regarding the talent level of the discriminated group. The discriminator believes that the target group has lower ability, controlling for observable characteristics, and thus should be more likely to vote against that group in early rounds, but *less* likely to vote against that group towards the end, when the desire to face a weak final round opponent becomes paramount.

#### Section IV: Empirical Results

Table 2 presents raw data on the number of votes received as a function of contestant characteristics in early, middle, and final rounds of the game. For the average player, one vote would be the expectation. Receiving a vote is a bad outcome, because the player with the most votes is eliminated. Comparing males and females in the first two columns, the numbers are similar across gender in all three rounds, with women weakly less likely to receive votes at each stage. Thus, there is no evidence in these numbers of discrimination against women on the show. Comparing across races in columns 3-6, Blacks have a slightly elevated rate of vote receipt in the early rounds, and slightly depressed rate in the middle and final rounds. In no instance, however, are Black and White rates statistically distinguishable. Asians are the one group that stand out in terms of how proficient they are in answer questions (Asians answer correctly approximately 65 percent of the time, four percentage points higher than any other racial group in the sample), so one



might expect that in the raw data they would be voted off less frequently in the early and middle rounds, but more at the end. That pattern, however, does not appear.<sup>18</sup> Hispanics receive high rates of votes in the early and middle rounds, but very low rates in the final round, potentially consistent with information-based models of discrimination. It is important to note, however, that there are a small number of observations for Hispanics, making the estimates imprecise. Contestants aged 50 and over receive excess votes at all three stages of the game, suggesting the possibility of taste-based discrimination.

The simple statistics in Table 2 may be misleading if contestants systematically differ along other observable dimensions such as their occupation or the skill with which they answer questions. Regression estimates of the following form are therefore estimated:

$$(1) \quad Vote_{cre} = X_{ce}'\beta + Z_{cre}'\gamma + \alpha_r + \eta_{cre}$$

where  $c$ ,  $r$ , and  $e$  index contestants, rounds, and episodes respectively.  $Vote$  is the number of votes a contestant receives. The specifications include a set of fixed contestant characteristics  $X$  (such as race, gender, age, educational level, region of residence), as well as variables that change by round (like measures of how well the player answered questions in current and previous rounds). In addition, fixed effects for each round are also included, with prime time rounds and day time rounds treated as different (i.e., there is a dummy for the first round of a prime time show and a separate dummy for the first round of a day time show).<sup>19</sup>

---

<sup>18</sup> The percent of correct answers is very similar for all other races. Men are 2 percent more likely to provide correct answers on average than women. Contestants age 50 and over correctly answer 2 percent fewer of their questions than the typical player.

<sup>19</sup> The results reported are from ordinary least squares. Negative binomial specifications which incorporate the fact that the dependent variable is a count variable, yield virtually identical results.

The results of the regressions are reported in Table 3, with early, middle, and final rounds analyzed separately. Specifications with a minimal set of covariates are presented in the odd columns; fuller specifications are shown in the even columns. Before turning to the coefficients of primary interest (gender, race, and age), it is important to first note that the conjectures regarding player strategy are largely confirmed in the data. In early and middle rounds, poor performance leads to substantial increases in votes. The inclusion of squared terms capturing non-linearities in how voting responds to correct answers makes it difficult to interpret the raw coefficients on these variables; evaluating the impact on votes of a player being one standard deviation above or below the mean aids in the interpretation. In the early rounds, a player one standard deviation below the mean in correct answers this round receives approximately one vote more than a contestant who is one standard deviation above the mean – a large impact given that the mean number of votes received per round is one. In the middle rounds, that same comparison yields a gap in votes of .82. In addition, in the middle rounds a contestant one standard deviation below the mean on *cumulative* performance prior to this round garners an extra .25 votes relative to a player one standard deviation above the mean. In the final round, however, the picture is very different. Slightly more votes are received if one performs poorly this round (about .20, consistent with the incentive not to vote for the strongest link in the final round because of the tie-breaking power). The tendency to punish bad cumulative performance disappears. There is no difference in the number of votes received for a player one standard deviation above and below the mean in *cumulative* performance up to the final round. Given the powerful incentive to vote off strong players at the end, it is surprising that good players in past rounds receive the same number of final-round votes (instead of more votes). Either the players succumb to the host's admonitions to vote for the weakest player, the contestants are not very skilled at determining who the most

successful player has been to date, or they (incorrectly) believe past performance is not a good predictor of future success.

The other strategic variable in the regression, the cumulative number of opponents voted for who are still alive, captures the “revenge” motive. This variable enters with the expected sign in all three parts of the game and is statistically significant in the middle and final rounds. Each extra opponent how you unsuccessfully tried to vote off in the past translates into between .09 and .27 extra votes against you in the current round.

The coefficients on education (not presented in the table) provide further confirmation of the shift in voting behavior in the final round. The only education category that consistently differs from the others is doctors. Controlling for other factors in the regressions, players with a doctorate receive substantially fewer votes than any other education category in early rounds (about .25 less per round), but receive *more* votes than any other category in the final round (about .15 votes more on average). This result is consistent with players perceiving that these highly educated contestants are strong players, even after controlling for observed performance up to this time.<sup>20</sup> Overall, the models are not particularly successful in explaining voting patterns, with  $R^2$  values ranging from .013 to .387.

For the most part, the regression results with respect to gender, race, and age mirror the findings that appeared in the raw data. In all specifications, women receive weakly fewer votes than men (the omitted category), although in no case is the difference across gender statistically significant. The differences between Blacks and Whites (the excluded group) fail to follow a systematic pattern and are never statistically significant. The Asian coefficients continue to be

---

<sup>20</sup> This pattern is consistent with statistical discrimination *in favor* of those with doctorates.

inconsistent with either model of discrimination. In the early and middle rounds, both Hispanics and older contestants carry positive coefficients. Because of large standard errors, the coefficients are statistically different than zero only in the middle rounds. These estimates are, however, substantively quite large: an otherwise average contestant attracts a 30-40 percent increase in votes by virtue of being Hispanic or old.

Perhaps the most interesting finding in the paper, once again mirroring the raw data, is the stark difference in the votes received by Hispanics and the old in the final round. While caution is warranted given the imprecision of the estimates, it is nonetheless intriguing that Hispanics carry an economically large coefficient of  $-.35$  in the final round. Although this estimate is not statistically distinguishable from zero at the  $.05$  level, it is statistically different from the coefficient on Hispanic in the middle rounds, suggesting a reversal in voting behavior towards Hispanics over the course of the game. These parameter estimates are consistent with statistical discrimination towards Hispanics: other players have low expectations about the skill of Hispanic competitors, even controlling for occupation, age, and performance during the game. In stark contrast, older players continue to attract excess votes in the final round. The persistent punishment of older players throughout the course of the game is consistent with a taste-based model of discrimination.<sup>21</sup>

Table 4 explores the sensitivity of these conclusions to various permutations of the data. Only the coefficients on the gender, race, and age variables are reported in Table 4. Results are again shown separately for early, middle, and final rounds. The first column simply reports the

---

<sup>21</sup> It certainly does not appear from the results of the head-to-head competition that competitors have private information about the likely success of Hispanics and old players in the final round. Hispanics win three of their five head-to-head opportunities, whereas old contestants are just one for seven in winning the final round.

baseline results presented in columns 2, 4, and 6 of the preceding table. Columns 2 and 3 of Table 4 divide the sample into day time and prime time shows. The same general pattern of results appears in both sets of shows, although individual coefficient estimates become imprecise, especially for the prime time show which has fewer observations. Column 4 eliminates approximately 20 percent of the shows in which explicit collusion may have been present, as evidenced by blocks of players voting for exactly the same progression of competitors in the initial rounds of the program.<sup>22</sup> The similarity of voting may, however, simply have arisen by chance, making this a crude indicator of explicit collusion. There are few systematic differences relative to the baseline estimates. The final two columns of Table 4 allow the impact of race and age to vary by gender. The omitted category is White males. The coefficient reported in the female row corresponds to White females. Imprecise estimates make it difficult to draw strong inferences, but the results do little to change the basic conclusion that there is little evidence of discrimination against women, Blacks, or Asians, but potential discrimination against Hispanics and older players.

An important consideration in whether one interprets the voting pattern data above as discriminatory or not, depends on whether these traits (such as Hispanic or old) are indeed predictive of future performance, conditional on other observable characteristics. Table 5 reports results relating player attributes to success in answering questions. The first three columns show results for early, middle, and final rounds respectively. In the last column, the dependent variable is an indicator corresponding to whether a contestant wins the head-to-head finale, conditional on

---

<sup>22</sup> While explicit collusion against minorities is not at odds with either model of discrimination, the mechanism underlying explicit collusion is very different from that of individual-level discrimination, so it is useful to try to distinguish the two.

making it that far. All the regressions control for education level and region of residence, as well as the cumulative percentage of correct answers given up to the current point in the game and that value squared. In no case is the coefficient on Hispanic statistically significant in Table 5, although it is negative in three of the four columns. Interestingly, the coefficient for those over age 50 is negative in all four columns, and very large and statistically significant in the last column. Controlling for other factors, old players do extremely poorly in the head-to-head finale. That result makes it even *more* striking that old players are voted off more than expected in the final round of voting – basic strategy would predict the opposite. Few other coefficients are statistically significant in this regression, except for past performance, which is highly predictive of future performance.

The analysis above, by focusing on overall voting patterns, may overlook discrimination by particular sub-groups towards other sub-groups. For instance, discrimination by Whites against Blacks might be disguised if Blacks rarely vote for one another. To examine that hypothesis further, Table 6 breaks down players into three dichotomous sets of groups by race, gender, and age.<sup>23</sup> The voting patterns within these groups are reported in the table. Rows correspond to the group casting votes; columns are the group being voted for. The odd columns of the table present the actual frequency with which votes are cast. The even columns reflect the predicted frequency under the null hypothesis that votes are cast randomly. Entries on the diagonal, which are of greatest interest, are boxed in bold.

The results with respect to race once again show little evidence of discrimination. Non-Blacks are slightly more likely to vote for Blacks than would be predicted by chance, but this

---

<sup>23</sup> Only rarely do Asians and Hispanics have a chance to vote for members of their own group, so they have been combined with Whites in this analysis.

difference is not statistically significant. Blacks vote for other Blacks somewhat less than would be expected (.208 vs .242), but again the difference is substantively small and not statistically significant. There is little evidence of discrimination against Blacks by males or females, the old or the young.

There is stronger evidence with respect to gender. Women are about 6 percentage points more likely to vote for men; men are 3 percentage points more likely to vote for women. In both cases, one can reject the null hypothesis of no gender bias at or near the .05 level. Because the contestants are almost equally balanced between men and women, however, the two biases largely counterbalance one another and there is little aggregate evidence of discrimination. If, however, the relevant group was dominated by men (e.g. high-ranking managers in many companies), gender bias of this sort could have an important negative impact on women.

Finally, with respect to voting against the old, both young and old contestants are more likely to vote against old competitors. Indeed, old players show an even stronger propensity to vote against other old contestants, although this coefficient is imprecisely estimated due to the small number of cases in which old players have the opportunity to face one another.

#### Section IV: Conclusions

Using the unique institutional set-up of *Weakest Link*, this paper tests for the presence and type of discrimination. Perhaps surprisingly, no evidence of discrimination towards Blacks or women is found, whereas there is substantively large magnitudes of observed discrimination towards Hispanics and the elderly. The data are consistent with statistical discrimination toward Hispanics and taste-based discrimination toward the old. There is also evidence that women tend to vote more frequently for men and vice-versa. It is important to emphasize, however, that

characteristics such as race, gender, and age do not appear to be the primary determinants of voting behavior.

Given the highly stylized nature of the interactions on this television show, one must use extraordinary caution in trying to draw general conclusions from these results. Indeed, one could imagine that the absence of observed discrimination towards Blacks in this artificial context might arise precisely *because* of the presence of real-world discrimination towards Blacks which has sensitized Americans to the importance of not appearing outwardly racist, regardless of inward beliefs. At some conscious or unconscious level, contestants may shy away from targeting Blacks on a nationally televised program. In contrast, players may be less concerned about appearing to target Hispanics and the elderly. Ideally, one would like to isolate real-world settings in which the strategic incentives flip as they do on *Weakest Link* to provide a more readily generalizable test of competing theories of discrimination.



Table 1: Summary Statistics

Contestant-level variables (N=1,016)	Overall	Early rounds	Middle rounds	Final round
Black	.198 (.399)	----	----	----
Hispanic	.022 (.146)	----	----	----
Asian	.029 (.167)	----	----	----
Female	.503 (.500)	----	----	----
Old (age>49)	.070 (.257)	----	----	----
Education=high school	.344 (.475)	----	----	----
Education=college	.421 (.494)	----	----	----
Education=post-college professional school	.084 (.277)	----	----	----
Education=doctorate	.036 (.187)	----	----	----
Education=Still in school	.076 (.265)	----	----	----
Education=Uncertain	.038 (.192)	----	----	----
On primetime show?	.197 (.398)	----	----	----
Contestant-round-level variables (N=3,273)	Overall	Early rounds	Middle rounds	Final round
Questions asked	2.62 (.62)	2.49 (.56)	2.63 (.58)	2.93 (.74)
Questions correct	1.66 (.91)	1.63 (.89)	1.64 (.90)	1.78 (1.01)
Percent correct this round (deviation from episode avg.)	.00 (.28)	.00 (.30)	.00 (.27)	.00 (.26)
Cumulative % correct (dev. from episode avg.)	.00 (.16)	.00 (.10)	.00 (.21)	.00 (.12)
Votes for contestant	1.00 (1.20)	1.00 (1.43)	1.00 (1.11)	1.00 (.76)
Cumulative opponents voted against who are still alive	.40 (.66)	.06 (.24)	.53 (.69)	.79 (.84)
Number of observations	3,273	1,191	1,599	483

Notes: The top panel of the table presents contestant-level variables; the bottom panel presents contestant-round level variables. Standard deviations are in parentheses. Early rounds correspond to the first two rounds of the prime time show (one hour long) and the first round of the day time show (30 minutes long). Middle rounds correspond to rounds 3-5 of the prime time show and rounds 2 and 3 of the day time show. Final round corresponds to round 6 of the prime time show and round 4 of the day time show. Average votes are equal to one in all rounds because each player casts one vote per round. The cumulative percent correct variable captures the deviation in percent correct for this player up to this point in the game, relative to the mean percent correct for all players on that episode.

Table 2: Raw Data on Votes Received by Race, Gender, and Age of Contestants

Round	Contestant characteristics:						
	Male	Female	White	Black	Asian	Hispanic	Age 50+
Early rounds	1.05 (.06) [N=590]	.95 (.06) [N=601]	.98 (.05) [N=897]	1.10 (.10) [N=229]	.62 (.17) [N=37]	1.29 (.29) [N=28]	1.32 (.18) [N=82]
Middle rounds	1.00 (.04) [N=775]	1.00 (.04) [N=824]	1.00 (.03) [N=1217]	.93 (.06) [N=300]	1.24 (.18) [N=50]	1.38 (.26) [N=32]	1.27 (.13) [N=92]
Final round	1.02 (.05) [N=230]	.98 (.05) [N=253]	1.02 (.04) [N=370]	.97 (.08) [N=93]	.83 (.25) [N=13]	.71 (.29) [N=7]	1.50 (.16) [N=24]

Notes: Values in the table are the mean number of votes cast for contestants in the race and gender category named, by rounds of the game. On average, players will receive one vote per round. Standard errors are in parentheses. The number of observations is presented in brackets. Early rounds correspond to the first two rounds of the prime time show (one hour long) and the first round of the day time show (30 minutes long). Middle rounds correspond to rounds 3-5 of the prime time show and rounds 2 and 3 of the day time show. Final round corresponds to round 6 of the prime time show and round 4 of the day time show.

Table 3: Regression Analysis of Votes Received

Variable	Early rounds		Middle rounds		Final round	
	(1)	(2)	(3)	(4)	(5)	(6)
Female	-.09 (.09)	-.07 (.08)	.00 (.06)	-.04 (.05)	-.06 (.08)	-.10 (.08)
Black	.13 (.11)	.07 (.09)	-.05 (.08)	-.00 (.06)	-.01 (.09)	-.01 (.09)
Asian	-.34 (.17)	-.20 (.13)	.28 (.19)	.33 (.15)	-.17 (.24)	-.13 (.26)
Hispanic	.32 (.29)	.17 (.24)	.40 (.25)	.41 (.19)	-.34 (.28)	-.35 (.24)
Age 50+	.34 (.19)	.17 (.17)	.30 (.13)	.30 (.11)	.49 (.16)	.42 (.16)
% correct this round (deviation from other players)	----	-2.44 (.11)	----	-1.71 (.09)	----	-.47 (.16)
Squared % correct this round (deviation from other players)	----	2.53 (.30)	----	.73 (.25)	----	.24 (.37)
Cumulative % correct (deviation from other players)	----	-.90 (.38)	----	-.70 (.14)	----	-.33 (.33)
Squared cumulative % correct (deviation from other players)	----	1.37 (.88)	----	.59 (.33)	----	2.77 (1.47)
Cumulative opponents voted against who are still alive	----	.12 (.29)	----	.27 (.04)	----	.09 (.04)
R-squared	.013	.387	.010	.257	.043	.089
Education and Region dummies included?	No	Yes	No	Yes	No	Yes
Number of obs.	1,191	1,191	1,599	1,599	483	483

Notes to table 3: The dependent variable in all columns is the number of votes received by a contestant in a given round. The unit of observation is a contestant-round. Estimation is done with ordinary least squares. Standard errors, clustered by episode and round, are in parentheses. All regressions include an exhaustive set of interactions controlling for round\*show-length interactions. The even columns include state of residence fixed effects and occupation dummies, but these coefficients are not reported in the table. Early rounds correspond to the first two rounds of the prime time show (one hour long) and the first round of the day time show (30 minutes long). Middle rounds correspond to rounds 3-5 of the prime time show and rounds 2 and 3 of the day time show.

Final round corresponds to round 6 of the prime time show and round 4 of the day time show.

Table 4: Sensitivity Analysis of the Key Coefficients

Variable	Baseline	Day time shows only	Prime time shows only	Drop if suspect explicit collusion	Separate coefficients by gender	
Early rounds	(1)	(2)	(3)	(4)	(5) Male	(6) Female
Female	-.07 (.08)	-.07 (.09)	-.05 (.14)	-.09 (.08)	----	.01 (.09)
Black	.07 (.09)	.14 (.12)	-.13 (.13)	.02 (.09)	.20 (.12)	-.08 (.13)
Asian	-.20 (.13)	-.24 (.14)	-.06 (.22)	-.33 (.12)	.00 (.21)	-.39 (.15)
Hispanic	.17 (.24)	.24 (.42)	.12 (.28)	.13 (.24)	.41 (.35)	.00 (.34)
Old	.17 (.17)	.30 (.20)	-.19 (.26)	.19 (.19)	.23 (.24)	.14 (.22)
Middle rounds					Male	Female
Female	-.04 (.05)	-.02 (.06)	.03 (.10)	-.00 (.06)	----	-.08 (.06)
Black	-.00 (.00)	-.05 (.07)	.07 (.17)	-.03 (.07)	-.06 (.09)	-.01 (.09)
Asian	.33 (.15)	.38 (.19)	.08 (.31)	.23 (.15)	.14 (.23)	.40 (.20)
Hispanic	.41 (.19)	.24 (.21)	.61 (.55)	.53 (.21)	.35 (.30)	.38 (.26)
Old	.30 (.11)	.33 (.12)	.24 (.37)	.22 (.12)	.32 (.15)	.28 (.17)
Final round					Male	Female
Female	-.10 (.08)	-.12 (.08)	.24 (.21)	-.07 (.10)	----	-.10 (.09)
Black	-.01 (.09)	-.01 (.10)	-.11 (.27)	-.06 (.10)	.01 (.13)	-.14 (.14)
Asian	-.13 (.26)	-.33 (.33)	.09 (.45)	-.39 (.26)	-.11 (.32)	-.28 (.40)
Hispanic	-.35 (.24)	-.18 (.29)	-1.40 (.32)	-.22 (.26)	-.19 (.36)	-.67 (.26)
Old	.42 (.16)	.50 (.15)	.05 (.76)	.55 (.16)	.19 (.26)	.62 (.17)

Notes: The baseline specification in column 1 corresponds to the results reported in columns 2, 4, and 6 of Table 3. All other columns adopt identical specifications, except for the differences noted. Columns 2 and 3 of this table divide the sample into prime time and day time shows. Column 4 drops roughly 20 percent of the sample in which there is a possibility that explicit collusion occurred between some set of the players. Columns 5 and 6 show results allowing the effects of race and age to vary by gender. The omitted category in these regressions is White males. The coefficient reported for female in column 6 is for White females. Standard errors, clustered by

episode and round, are in parentheses.

Table 5: Do Observable Player Characteristics Predict Future Success in Answering Questions?

Variable	Dependent variable: Percent of answers correct this round			Dependent variable: Winner in head-to-head finale
	Early rounds	Middle rounds	Final round	Head-to-head
Female	-.006 (.018)	-.006 (.016)	-.055 (.027)	-.067 (.054)
Black	-.018 (.023)	.005 (.020)	-.004 (.036)	-.032 (.070)
Asian	.066 (.055)	-.014 (.044)	.133 (.062)	-.204 (.171)
Hispanic	-.054 (.066)	.008 (.060)	-.144 (.162)	-.020 (.195)
Age 50+	-.045 (.041)	-.002 (.035)	-.059 (.069)	-.390 (.173)
Cumulative % correct (deviation from other players)	-.011 (.116)	.064 (.043)	.215 (.128)	2.033 (.293)
Squared cumulative % correct (deviation from other players)	-.506 (.242)	-.241 (.138)	.506 (.645)	-.370 (2.292)
R-squared	.024	.016	.058	.172
Education and Region dummies included?	Yes	Yes	Yes	Yes
Number of obs.	1,191	1,599	483	322

Notes: The dependent variable is named at the head of the column. Columns 1-3 correspond to the percent of correct answers in a given round as a function of observable characteristics and percentage of questions answered correctly thus far in the show. The final column corresponds to an indicator variable for whether a contestant wins, conditional on making it to the final round. Standard errors, clustered by episode and round, are in parentheses.

Table 6: Propensity to Vote for Contestants of a Particular Race, Gender, or Age  
(By Contestant's Own Race, Gender, and Age)

Characteristic of the contestant that is voting	Likelihood of voting for a contestant with the following characteristic: (standard error in parentheses)					
	Black		Female		Old	
	Actual	Predicted if vote randomly	Actual	Predicted if vote randomly	Actual	Predicted if vote randomly
Black	.208 (.029) N=207	.242	.547 (.021) N=510	.522	.288 (.034) N=153	.240
Not Black	.285 (.011) N=1,713	.279	.486 (.010) N=2,209	.506	.327 (.018) N=553	.248
Female	.278 (.014) N=1,016	.275	.372 (.013) N=1,396	.430	.324 (.022) N=370	.249
Male	.275 (.014) N=904	.275	.630 (.015) N=1,323	.592	.312 (.023) N=336	.249
Old	.250 (.039) N=128	.286	.479 (.037) N=169	.523	.412 (.073) N=34	.255
Not old	.278 (.010) N=1,792	.274	.499 (.010) N=2,550	.508	.314 (.016) N=672	.246

Notes: Values in the table are the actual and predicted probabilities that votes cast by contestants of a particular group identified by rows of the table will be for contestants of the group identified in columns of the table. Predicted probabilities assume that players vote randomly among all eligible competitors. The values in the table exclude the final round of voting because strategic incentives are reversed in the final round. Standard errors, in parentheses, are computed under the null hypothesis of random voting. Cases in which all of the remaining competitors are in the target group, or none



of the remaining competitors are in the target group are excluded from the calculations.

## References

- Aigner, Dennis, and Glen Cain, 1977, "Statistical Theories of Discrimination in the Labor Market," *Industrial and Labor Relations Review* 30(January): 175-187.
- Altonji, Joseph, and Rebecca Blank, 1999, "Race and Gender in the Labor Market," in *Handbook of Labor Economics*, Vol. 3c, edited by David Card and Orley Ashenfelter.
- Altonji, Joseph, and Charles Pierret, 2001, "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics* 116 (February): 293-312.
- Antonovics, Kate, Peter Arcidiancono, Randall Walsh, 2002, "Race and Gender Discrimination in the Weakest Link," Incomplete and unpublished draft.
- Arrow, Kenneth, 1973, "The Theory of Discrimination," in Discrimination in Labor Markets, edited by Orley Ashenfelter and Albert Rees. Princeton, NJ: Princeton University Press
- Ayres, Ian, and Peter Siegelman, 1995, "Race and Gender Discrimination in Bargaining for a New Car," *American Economic Review*, Vol. 85, No. 3. (Jun., 1995), pp. 304-321.
- Becker, Gary, 1957, The Economics of Discrimination, Chicago: University of Chicago Press.
- Bertrand, Marianne, and Sendhil Mullanaithan, 2002, "Are Emily and Brendan More Employable than Latoya and Tyrone? Evidence on Racial Discrimination in the Labor Market from a Large Randomized Experiment," unpublished manuscript.
- Blank, Rebecca, 2001, "An Overview of Trends in Social and Economic Well-being, by Race," in America Becoming: Racial Trends and their Consequences, Volume 1, edited by Neil Smelser and William Julius Wilson, and Faith Mitchell, Washington, DC: National Academy of Sciences Press, pp. 21-39.
- Coate, Stephen, and Glenn Loury, 1993, "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* vol. 83, n5 (December 1993): 1220-40.
- Fryer, Roland, 2001, "Economists' Models of Discrimination: An Analytic Survey," Unpublished manuscript, University of Chicago Department of Economics.
- Fryer, Roland, and Matt Jackson, 2002, "Categorical Cognition: A Psychological Model of Categories and Identification in Decision Making," Unpublished manuscript.
- Gertner, Robert, 1993, "Game Shows and Economic Behavior: Risk Taking on 'Card Sharks'," *Quarterly Journal of Economics* 108(2): 507-522.
- Heckman, James, and Peter Siegelman, 1992, "The Urban Institute Audit Studies: Their Methods and Findings," in Clear and convincing evidence: Measurement of discrimination in America, edited by Micheal Fix and Raymond Struyk, pp. 187-258,

- Washington, DC: Urban Institute.
- Howell, Susan, and Robert Sims, 1994, "Survey Research and Racially Charged Elections: The Case of David Duke in Louisiana," *Political Behavior* 16(June): 219-236.
- Kahn, Lawrence, 1991, "Discrimination in Professional Sports: A Survey of the Literature," *Industrial & Labor Relations Review*, 44(April): 395-418.
- Knowles, John, Nicola Persico, and Petra Todd, 2001, "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy* 109 (February):203-229.
- Kuklinski, James, Paul M. Sniderman, Kathleen Knight, Thomas Piazza, Philip E. Tetlock, Gordon R. Lawrence, and Barbara Mellers, 1997, "Racial Prejudice and Attitudes towards Affirmative Action," *American Journal of Political Science* 41(April): 402-419.
- Lundberg, Shelly, and Richard Startz, 1983, "Private Discrimination and Social Intervention in Competitive Markets," *American Economic Review* 73 (3): 340-347.
- Metrick, Andrew, 1995, "A Natural Experiment in "Jeopardy!" *American Economic Review* Vol. 85, No. 1. (Mar., 1995), pp. 240-253.
- Phelps, Edmund, 1972, "A Statistical Theory of Racism and Sexism," *American Economic Review* 62(4): 659-661.