

TESTING THEORIES OF DISCRIMINATION: EVIDENCE FROM *WEAKEST LINK**

STEVEN D. LEVITT
University of Chicago

ABSTRACT

Contestant voting behavior on the television game show *Weakest Link* provides an unusual opportunity to distinguish between taste-based and information-based theories of discrimination. 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 characteristics, 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. Although players are voting strategically, evidence of discrimination is limited. There is little in the data to suggest discrimination against women and blacks. I find some patterns consistent with information-based discrimination toward Hispanics (other players perceive them as having low ability) and taste-based discrimination against older players (other players treat them with animus).

I. INTRODUCTION

THERE are two leading theories of discrimination. The first theory is based on tastes and originates with Gary Becker.¹ 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 price 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

* I would like to thank Mark Duggan, Roland Fryer, Austan Goolsbee, Doug MacAdam, and an anonymous referee for helpful discussions on the topic, as well as Jeannette Levitt for providing the raw data. Carol Klein, Monica Li, 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, California. Financial support was provided by the National Science Foundation.

¹ See Gary S. Becker, *The Economics of Discrimination* (1957).

[*Journal of Law and Economics*, vol. XLVII (October 2004)]

© 2004 by The University of Chicago. All rights reserved. 0022-2186/2004/4702-0014\$01.50

models circumvent that criticism.² In these models, individuals (typically employers) discriminate against particular groups because either (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.³

Measuring the extent of discrimination poses a difficult empirical challenge.⁴ Self-reported data are unlikely to accurately reflect attitudes if there is a perceived stigma attached to racist views.⁵ 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.⁶ Using this methodology, for instance, Ian Ayres and Peter Siegelman find that car dealerships attempt to charge higher prices to blacks and women.⁷ Marianne Bertrand and Sendhil Mullanaithan find that applicants who submit resumes

² See Edmund Phelps, *A Statistical Theory of Racism and Sexism*, 62 *Am. Econ. Rev.* 659 (1972); Kenneth Arrow, *The Theory of Discrimination*, in *Discrimination in Labor Markets* 91 (Orley Ashenfelter & Albert Rees eds. 1973); Dennis Aigner & Glen Cain, *Statistical Theories of Discrimination in the Labor Market*, 30 *Indus. & Lab. Rel. Rev.* 175 (1977); Shelly Lundberg & Richard Startz, *Private Discrimination and Social Intervention in Competitive Markets*, 73 *Am. Econ. Rev.* 340 (1983); Stephen Coate & Glenn Loury, *Will Affirmative-Action Policies Eliminate Negative Stereotypes?* 83 *Am. Econ. Rev.* 1220 (1993); Roland Fryer & Matt Jackson, *Categorical Cognition: A Psychological Model of Categories and Identification in Decision Making* (unpublished manuscript, Harvard Univ., Dep’t Econ. 2002).

³ See Roland Fryer, *Economists’ Models of Discrimination: An Analytic Survey* (unpublished manuscript, Harvard Univ., Dep’t Econ. 2001).

⁴ For a more complete survey of the empirical literature on discrimination than is presented here, see Joseph Altonji & Rebecca Blank, *Race and Gender in the Labor Market*, in 3C *Handbook of Labor Economics* 3143 (David Card & Orley Ashenfelter eds. 1999).

⁵ As one vivid example, Susan Howell & Robert Sims, *Survey Research and Racially Charged Elections: The Case of David Duke in Louisiana*, 16 *Pol. Behav.* 219 (1994), analyzes the Louisiana gubernatorial election in which white supremacist politician David Duke received far more votes than preelection surveys and exit polls had predicted. James Kuklinski *et al.*, *Racial Prejudice and Attitudes towards Affirmative Action*, 41 *Am. J. Pol. Sci.* 402 (1997), uses a clever methodology designed to minimize self-reporting issues and finds no evidence of discrimination toward blacks outside the South but substantial racism among white southerners.

⁶ James Heckman & Peter Siegelman, *The Urban Institute Audit Studies: Their Methods and Findings*, in *Clear and Convincing Evidence: Measurement of Discrimination in America* 187 (Michael Fix & Raymond Struyk eds. 1992), criticizes audit studies on the grounds that the simulated discriminatory transactions that occur in audit studies would be unlikely to arise in a market economy. For example, blacks may not shop for cars at all-white car dealerships because they fear discrimination. If that is the case, then in the real world, the consequences of discrimination will be far less than suggested by the audit study.

⁷ Ian Ayres & Peter Siegelman, *Race and Gender Discrimination in Bargaining for a New Car*, 85 *Am. Econ. Rev.* 304 (1995).

that carry distinctively black names are less likely to receive job interviews.⁸ A second approach to the empirical study of discrimination is to compare salaries with 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.⁹ Mixed evidence of discrimination is found in this literature, with some studies finding salary and customer discrimination against blacks, particularly in basketball. John Knowles, Nicola Persico, and Petra Todd, 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.¹⁰

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*.¹¹ On this show, contestants answer trivia questions over a series of rounds, with one contestant eliminated each round on the basis of 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,

⁸ Marianne Bertrand & Sendhil Mullanaithan, 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, Univ. Chicago, Grad. Sch. Bus. 2002).

⁹ Lawrence Kahn, Discrimination in Professional Sports: A Survey of the Literature, 44 *Indus. & Lab. Rel. Rev.* 395 (1991).

¹⁰ John Knowles, Nicola Persico, & Petra Todd, Racial Bias in Motor Vehicle Searches: Theory and Evidence, 109 *J. Pol. Econ.* 203 (2001). Also, Joseph Altonji & Charles Pierret, Employer Learning and Statistical Discrimination, 116 *Q. J. Econ.* 293 (2001), presents a clever methodology for ferreting out statistical discrimination using information about how wages change with employment tenure. They find little evidence of statistical discrimination based on race.

¹¹ 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 (Kate Antonovics, Peter Arcidiancono, & Randall Walsh, Race and Gender Discrimination in the *Weakest Link* (unpublished manuscript, Duke Univ., Dep’t Econ. 2004)). Their research also concludes that there is little evidence of discrimination toward blacks and women. They do not analyze discrimination toward Hispanics or the elderly. Two other unpublished manuscripts also analyze data from *Weakest Link* to answer very different questions. Philippe Fevrier & Laurent Linnemer, Equilibrium Selection: Payoff or Risk Dominance? The Case of the “Weakest Link” (unpublished manuscript, Univ. Chicago, Dep’t Econ. 2002), examines how contestants coordinate on equilibria on the show. Marco Haan *et al.*, The Weakest Link: A Field Experiment in Rational Decision Making (unpublished manuscript, Univ. Groningen 2002), considers the decision of contestants to “bank” money on the show. See Robert Gertner, Game Shows and Economic Behavior: Risk Taking on “Card Sharks,” 108 *Q. J. Econ.* 507 (1993), and Andrew Metrick, A Natural Experiment in “Jeopardy!” 85 *Am. Econ. Rev.* 240 (1995), for other economic analyses using game shows as a vehicle.

there is a strong incentive to vote off the least-skilled players in the early rounds.¹² 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 that are answered correctly up to that point in the show or the individual's education (which may predict future skill in answering questions).¹³

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-skilled opponent in the final round to increase one's chance of winning the final prize. This change in behavior is observed quite clearly in the data.

In the late rounds, the two theories that were discussed above offer different predictions about how discriminators will behave. In the taste-based models, those who discriminate against a group do so not because they think that group is less talented but rather because they do not like the group. Thus, one would expect that a taste-based discriminator would continue to vote disproportionately 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 opponent will be from that group, which increases the perceived likelihood that the discriminator will win the prize. Consequently, voting behavior in the late rounds provides a unique opportunity to distinguish empirically between the two competing models.¹⁴

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 in with whom they interact. Second, the individuals who appear on this game show are highly selected, with respect to both who applies to be on the show and who is chosen to be on the show. Little is known about the precise

¹² 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 hometown. The contestants can also see one another and therefore can observe race, gender, and other visual cues, such as degree of physical attractiveness.

¹³ 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.

¹⁴ 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 toward the end of the game.

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, contestants tend to be young and well educated. 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 for whom to vote depends not just on one's own views but also on one's beliefs as to how other contestants will vote. Empirical results show that players attempt to punish competitors who voted against them in earlier rounds. A contestant would ideally like to cast his or her vote for the player who will receive the most votes and be eliminated, even if that player would not otherwise be the contestant's voting choice. Consequently, for instance, someone who is not racist but who 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 evidence of discrimination toward blacks, Asians, or women. The number of votes cast against these groups is similar to the number 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) than other contestants early in the show, 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.¹⁵ 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.

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

¹⁵ John List, *Friend or Foe: A Natural Experiment of the Prisoner's Dilemma* (unpublished manuscript, Univ. Maryland, Dep't Econ. 2002), also finds evidence of taste-based discrimination toward older players in a very different game show setting.

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 answers have been assembled.¹⁶ 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.¹⁷ At the conclusion of each round, contestants secretly and simultaneously 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 who cast the votes, are then revealed. The competitor who receives the most votes is eliminated from the show (the “weakest link”), can no longer answer questions, and is 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 who 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 an incorrect answer. The time that is 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 1-hour prime-time show with eight competitors, eight rounds (seven in which one player is eliminated and one 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 \$75,000.¹⁸ In practice, the total prize money earned

¹⁶ 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 daytime show, the value of answers is as follows: \$250, \$500, \$1,000, \$2,500, \$5,000, \$12,500.

¹⁷ The trade-off 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 decision to bank is of potential interest, but my data set does not contain information on who banked.

¹⁸ 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.

is well below the theoretical maximum. The median payout is roughly \$80,000 in the prime-time version and \$10,000 in the daytime version.

Contestants are chosen for the shows through open auditions that are conducted around the country. Only a small fraction of those auditioning for the show are accepted. Interviews with past contestants suggest that the criteria for being chosen for the show include not only skill at answering questions but also having idiosyncracies that will make for entertaining television (for example, an unusual job, such as exotic dancer).¹⁹ Like other game shows (see, for example, *List 2002*), contestants are disproportionately young adults.

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 that involve celebrities (for example, a *Brady Bunch* reunion show) and virtually every daytime show aired prior to January 2003. There are a total of 25 prime-time shows and 136 daytime shows. Each prime-time show yields eight person-level observations, 33 person-votes (eight in the first round, seven in the second round, . . . , three in the sixth round). Each daytime show yields six person-level observations and 18 person-votes. Thus, in total, the data set includes 1,016 person-level observations and 3,273 person-votes.

Table 1 presents summary statistics for the data set, which highlight the nonrepresentativeness of the sample along many dimensions. The top part of the table reports data at the contestant level. Blacks are overrepresented on the show (roughly 20 percent of the contestants are black), but Asians and Hispanics are underrepresented (2.9 and 2.2 percent, respectively).²⁰ The scarcity of Asians and Hispanics is particularly notable given that a disproportionate share of contestants are from California, where the show is taped. 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. Players are almost evenly distributed between males and females. Contestants are heavily skewed toward the younger portion of the adult age distribution. Only 7 percent of the contestants are 50 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

¹⁹ For interviews with contestants, see <http://www.chris-lambert.com/LINK/linkinterview.html>.

²⁰ Because of some ambiguity in visual assessments of race, the category labeled "Hispanic" is actually a catch-all category for nonblack, nonwhite, non-Asian contestants, which includes a small number of Native Americans and Pacific Islanders as well as Hispanics.

TABLE 1
SUMMARY STATISTICS

Variables	Overall	Early Rounds	Middle Rounds	Final Round
Contestant level ($N = 1,016$):				
Black	.198 (.399)
Hispanic	.022 (.146)
Asian	.029 (.167)
Female	.503 (.500)
Old (age > 49)	.070 (.257)
Education:				
High school	.344 (.475)
College	.421 (.494)
Postcollege professional school	.084 (.277)
Doctorate	.036 (.187)
Still in school	.076 (.265)
Uncertain	.038 (.192)
On prime-time show	.197 (.398)
Round level ($N = 3,273$):				
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)
% Correct this round (deviation from episode average)	.00 (.28)	.00 (.30)	.00 (.27)	.00 (.26)
Cumulative % correct (deviation from episode average)	.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)
N	3,273	1,191	1,599	483

NOTE.—Standard deviations are in parentheses. Early rounds correspond to the first two rounds of the prime-time show (1 hour long) and the first round of the daytime show (30 minutes long). Middle rounds correspond to rounds 3–5 of the prime-time show and rounds 2 and 3 of the daytime show. Final round corresponds to round 6 of the prime-time show and round 4 of the daytime 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.

education (although many of them may nonetheless have higher levels of education). Over 40 percent are classified as college educated, with an additional 8 percent having professional degrees (such as law) and another 3.6 percent with doctorates (Ph.D. or M.D.).

The lower part 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. 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.

Also shown in Table 1 are statistics that report the deviation of that player

from other contestants on the same show in the percentage of questions that were answered correctly this round and cumulatively up to 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.

III. STRATEGIC CONSIDERATIONS OF CONTESTANTS

The primary strategic decision that a player faces is for whom to vote.²¹ The complexity of the situation precludes a formal model. Within each round there 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 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 daytime 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.²² Players who survive to the

²¹ As noted earlier, the question of whether to bank poses an interesting question of decision making under uncertainty addressed in Haan et al., *supra* note 11.

²² 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, *supra* note 11) and bank money optimally, a lower bound on the expected value of a correct answer can be arrived at under the assumption that 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 daytime 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 $1/6^8$; for the daytime show, the payoff to six straight correct answers is \$12,500, and 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 daytime 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 of winning 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 (for example, 70 percent correct answers versus an opponent with 50 percent correct) will win the final round about 80 percent of the time. On the daytime 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, performance in the early rounds of the show is a powerful empirical 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, toward the end of the game, the desire to eliminate strong competitors becomes paramount in importance.

The third consideration relates to the impact that a player's vote has on others' voting behavior toward him or her. 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.²³ Thus, in general there is incentive to shy away from voting for players who have done well in the current round, aside from the first two considerations that were 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

probability that any one answer is pivotal is $1/6^5$. If banking is done optimally, a risk-neutral player will bank only if it increases the expected value of the pool, making this a lower bound.

²³ 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 percentage of answers correct in that round (players who were asked more questions are penalized). When a tie remains, the allocation rule used is unclear.

him or her in 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 that was 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.²⁴

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 considerations (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 toward voting off weak players, whereas at the end (at least in relative terms), the incentives to vote off strong players are greater.

Self-reported data collected from a subset of the daytime shows confirm the basic sensibility of the arguments made above. After the votes have been cast, the host questions some of the contestants regarding their motives for voting for a particular player. While there are strong reasons to suspect that contestants will not always tell the truth (it seems highly unlikely, for instance, that race would be cited as a reason for casting a vote), the pattern of self-reports in Table 2 is nonetheless somewhat elucidating. Contestant explanations for why they cast their vote are categorized into five mutually exclusive groups: bad play by the opponent, good play by the opponent, personal characteristics (such as the opponent's job, looks, manner of dress, and so on), revenge for past votes, and all other reasons. Results are shown round

²⁴ See Fevrier & Linnemer, *supra* note 11.

TABLE 2
 SELF-REPORTED EXPLANATIONS OF WHY CONTESTANTS CAST THEIR VOTES, BY ROUND

Round	Bad Play by Opponent	Good Play by Opponent	Personal Characteristics of Opponent	Revenge	Other	Total
1 (<i>N</i> = 139)	59.7	2.9	28.1	.0	9.4	100.0
2 (<i>N</i> = 153)	67.3	2.0	15.0	6.5	9.2	100.0
3 (<i>N</i> = 132)	59.8	6.1	15.2	7.6	11.4	100.0
4 (<i>N</i> = 133)	36.8	32.3	14.3	7.5	9.0	100.0

NOTE.—Contestant explanations are based on a random subset of daytime shows. “Bad play by opponent” includes any comment related to the quality of a player’s answers, speed of answering, and so on. “Good play by opponent” refers to cases in which a player specifically mentions that an opponent poses a threat because of skill at the game. “Personal characteristics of opponent” include comments about gender, age, manner of dress, style of speaking, attitude, and ethnic origin. “Revenge” refers to a player specifically mentioning that his or her vote this period is retribution for a past vote by another player. All remaining answers are classified as “other.” The total number of self-reports that the percentages are based on is in parentheses in the first column.

by round. In all four rounds of the show, the most commonly stated rationale for voting a player off was bad play. Roughly 60 percent of comments related to bad play in the first three rounds of voting. In the fourth round, however, the prevalence of that answer fell (to 37 percent), as would be expected. In the first three rounds, few votes are cast because of a perceived threat posed by a strong player. In the final round, almost one-third of all votes are justified by a player being too strong. Given the moral suasion by the moderator of the show, this category of responses is likely to be underreported. Nonetheless, the direction of change is consistent with the theory above. Personal characteristics are most frequently cited in the first round (before there is much of a history on which to base explanations). Revenge, after the first round, receives a relatively stable share of the weight.

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 target group throughout all rounds. In an information-based model, on the other hand, the discrimination arises from beliefs regarding the talent level of the target group. The discriminator believes that the target group has lower ability, controlling for observable characteristics, and thus should be more likely vote against that group in early rounds but less likely to vote against that group toward the end, when the desire to face a weak final round opponent becomes paramount.

IV. EMPIRICAL RESULTS

Table 3 presents raw data on the number of votes received as a function of contestant characteristics in early, middle, and final rounds of the game.

TABLE 3
RAW DATA ON VOTES RECEIVED BY RACE, GENDER, AND AGE OF CONTESTANTS

	Male	Female	White	Black	Asian	Hispanic	Age 50+
Early rounds	1.05 (.06)	.95 (.06)	.98 (.05)	1.10 (.10)	.62 (.17)	1.29 (.29)	1.32 (.18)
<i>N</i>	590	601	897	229	37	28	82
Middle rounds	1.00 (.04)	1.00 (.04)	1.00 (.03)	.93 (.06)	1.24 (.18)	1.38 (.26)	1.27 (.13)
<i>N</i>	775	824	1,217	300	50	32	92
Final round	1.02 (.05)	.98 (.05)	1.02 (.04)	.97 (.08)	.83 (.25)	.71 (.29)	1.50 (.16)
<i>N</i>	230	253	370	93	13	7	24

NOTE.—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. Early rounds correspond to the first two rounds of the prime-time show (1 hour long) and the first round of the daytime show (30 minutes long). Middle rounds correspond to rounds 3–5 of the prime-time show and rounds 2 and 3 of the daytime show. Final round corresponds to round 6 of the prime-time show and round 4 of the daytime show.

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, we see that the numbers are similar across gender in all three rounds, with women slightly 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, we see that blacks have a slightly elevated rate of vote receipt in the early rounds and a 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 stands out in terms of proficiency in answering 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.²⁵ 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 is a small number of observations for Hispanics, which makes the estimates imprecise. Contestants aged 50 and over receive excess votes at all three stages of the game, which suggests the possibility of taste-based discrimination.

The simple statistics in Table 3 may be misleading if contestants syste-

²⁵ The percentage of correct answers is very similar for all other races. Men are 2 percent more likely to provide correct answers on average than are women. Contestants age 50 and over correctly answer 2 percent fewer of their questions than does the typical player. Another dimension of player skill is the speed with which the answers are given. Although I do not have a direct measure of this, I am able to analyze whether the composition of the players affects the total number of questions asked in a round. Controlling for the round of the game, marginally fewer questions are asked when more black, old, Asian, and female players are in the game, which implies that these types of players may answer more slowly than whites and Hispanics.

matically 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:

$$\text{Vote}_{cre} = X_{ce'}\Gamma + Z_{cre'}\theta + \lambda_r + \varepsilon_{cre}, \quad (1)$$

where “c,” “r,” and “e” index contestants, rounds, and episodes, respectively. The variable *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, and region of residence), as well as variables that change by round (such as 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 daytime rounds treated as different (that is, there is a dummy for the first round of a prime-time show and a separate dummy for the first round of a daytime show).²⁶

The results of the regressions are reported in Table 4, 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. 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 that capture nonlinearities 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 1 standard deviation above or below the mean aids in the interpretation. In the early rounds, a player 1 standard deviation below the mean in correct answers receives approximately one vote more than a contestant who is 1 standard deviation above the mean—a large effect 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 1 standard deviation below the mean on cumulative performance garners an extra .25 votes relative to a player 1 standard deviation above the mean. In the final round, however, the picture changes. The punishment for poor performance is greatly reduced to about .20, consistent with the incentive not to vote for the strongest link in the final round because of that player’s tie-breaking power. The tendency to punish bad cumulative performance, however, completely disappears. There is no difference in the number of votes received for a player 1 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. The players succumb to the host’s

²⁶ The results reported are from ordinary least squares regressions. Negative binomial specifications that incorporate the fact that the dependent variable is a count variable yield virtually identical results.

TABLE 4
REGRESSION ANALYSIS OF VOTES RECEIVED

Variable	EARLY ROUNDS (<i>N</i> = 1,191)		MIDDLE ROUNDS (<i>N</i> = 1,599)		FINAL ROUND (<i>N</i> = 483)	
	(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 alive12 (.29)27 (.04)09 (.04)
<i>R</i> ²	.013	.387	.010	.257	.043	.089
Education and region dummies included?	No	Yes	No	Yes	No	Yes

NOTE.—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 by 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 (1 hour long) and the first round of the daytime show (30 minutes long). Middle rounds correspond to rounds 3–5 of the -time show and rounds 2 and 3 of the daytime show. Final round corresponds to round 6 of the prime-time show and round 4 of the daytime show.

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 that 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 who a player unsuccessfully tried to vote off in the past translates into between .09 and .27 extra votes against that player 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 who perceive that these highly educated contestants are strong players, even after controlling for observed performance up to this time.²⁷ 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 inconsistent with either model of discrimination. In the early and middle rounds, both Hispanics and older contestants have 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 elderly in the final round. While caution is warranted given the imprecision of the estimates, it is nonetheless intriguing that Hispanics have 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, which suggests a reversal in voting behavior toward Hispanics over the course of the game. These parameter estimates are consistent with statistical discrimination

²⁷ This pattern is consistent with statistical discrimination in favor of those with doctorates.

TABLE 5
SENSITIVITY ANALYSIS OF THE KEY COEFFICIENTS

VARIABLE	BASELINE (1)	DAYTIME SHOWS ONLY (2)	PRIME-TIME SHOWS ONLY (3)	DROP IF SUSPECT EXPLICIT COLLUSION (4)	SEPARATE COEFFICIENTS BY GENDER	
					Male (5)	Female (6)
Early rounds:						
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:						
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:						
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)

NOTE.—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 daytime 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 that allow 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, which are clustered by episode and round, are in parentheses.

toward 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.²⁸

Table 5 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 5. Results are again shown separately for early, middle, and final rounds. The first column simply reports the baseline results presented in columns 2, 4, and 6 of the preceding table. Columns 2 and 3 of Table 5 divide the sample into daytime and prime-time shows. The same general

²⁸ 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 older players in the final round. Hispanics win three of their five head-to-head opportunities, whereas older contestants are just one for seven in winning the final round.

TABLE 6
DO OBSERVABLE PLAYER CHARACTERISTICS PREDICT FUTURE
SUCCESS IN ANSWERING QUESTIONS?

VARIABLE	PERCENT OF ANSWERS CORRECT THIS ROUND			WINNER IN
	Early Rounds	Middle Rounds	Final Round	HEAD-TO-HEAD FINALE
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)
(Cumulative % correct) ² (deviation from other players)	-.506 (.242)	-.241 (.138)	.506 (.645)	-.370 (2.292)
R ²	.024	.016	.058	.172
N	1,191	1,599	483	322

NOTE.—The dependent variables are given as column heads. 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, which is conditional on making it to the final round. Education and region dummies are included for all regressions. Standard errors, which are clustered by episode and round, are in parentheses.

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 blocs of players voting for exactly the same progression of competitors in the initial rounds of the program.²⁹ The similarity of voting may, however, simply have arisen by chance, which makes this a crude indicator of explicit collusion. There are few systematic differences relative to the baseline estimates. The final two columns of Table 5 allow the effect of race and age to vary by gender. The omitted category is white males. The coefficient that is 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 evidence of discrimination against Hispanics and older players.

An important consideration in whether one interprets the voting pattern data above as discriminatory depends on whether these traits (such as Hispanic or elderly) are indeed predictive of future performance, conditional on other observable characteristics. Table 6 reports results that relate player attributes to success in answering questions. The first, second, and third

²⁹ While explicit collusion against minorities is not at odds with either model of discrimination, the mechanism that underlies explicit collusion is very different from that of individual-level discrimination, so it is useful to try to distinguish the two.

TABLE 7
 PROPENSITY TO VOTE FOR CONTESTANTS OF A PARTICULAR RACE, GENDER,
 OR AGE (by Contestant's Own Race, Gender, and Age)

VOTING CONTESTANT	BLACK		FEMALE		OLD	
	Actual	Predicted If Vote Randomly	Actual	Predicted If Vote Randomly	Actual	Predicted If Vote Randomly
Black	.208 (.029)	.242	.547 (.021)	.522	.288 (.034)	.240
<i>N</i>	207		510		153	
Not black	.285 (.011)	.279	.486 (.010)	.506	.327 (.018)	.248
<i>N</i>	1,713		2,209		553	
Female	.278 (.014)	.275	.372 (.013)	.430	.324 (.022)	.249
<i>N</i>	1,016		1,396		370	
Male	.275 (.014)	.275	.630 (.015)	.592	.312 (.023)	.249
<i>N</i>	904		1,323		336	
Old	.250 (.039)	.286	.479 (.037)	.523	.412 (.073)	.255
<i>N</i>	128		169		34	
Not old	.278 (.010)	.274	.499 (.010)	.508	.314 (.016)	.246
<i>N</i>	1,792		2,550		672	

NOTE.—The 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 the 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 or none of the remaining competitors are in the target group are excluded from the calculations. Similar results are obtained when the predicted fraction of votes controls for past performance, education, and region.

columns show results for early, middle, and final rounds, respectively. In the last column, the dependent variable is an indicator that corresponds to whether a contestant wins the head-to-head finale, conditional on 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 6, 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, older players do extremely poorly in the head-to-head finale. That result makes it even more striking that older 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 subgroups toward other subgroups. For instance, discrimination by whites against blacks might be disguised if blacks rarely vote for one another. To examine that hypothesis further, Table 7 breaks down

players into three dichotomous sets of groups by race, gender, and age.³⁰ The voting patterns within these groups are reported in the table. Rows correspond to the group casting votes; columns are the group receiving the votes. 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.

The results with respect to race once again show little evidence of discrimination. Nonblacks are slightly more likely to vote for blacks than would be predicted by chance, but this difference is not statistically significant. Blacks vote for other blacks somewhat less than would be expected (.208 versus .242), but again the difference is substantively small and not statistically significant. There is little evidence of discrimination against blacks by males, 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 (for instance, 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 elderly, both young and old contestants are more likely to vote against older competitors. Indeed, older players show an even stronger propensity to vote against other older contestants, although this coefficient is imprecisely estimated because of the small number of cases in which older players have the opportunity to face one another.

V. CONCLUSIONS

Using the unique institutional setup of *Weakest Link*, this paper tests for the presence and type of discrimination. Perhaps surprisingly, no evidence of discrimination toward blacks or women is found, whereas there are substantively large magnitudes of observed discrimination toward Hispanics and the elderly. The data are consistent with statistical discrimination toward Hispanics and taste-based discrimination toward the elderly. There is also evidence that women tend to vote more frequently for men and vice versa. It is important to emphasize, however, that strategic concerns, rather than characteristics such as race, gender, and age, appear to be the primary determinants of voting behavior.

³⁰ 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.

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 toward blacks in this artificial context might arise precisely because of the presence of real-world discrimination toward 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.

BIBLIOGRAPHY

- Aigner, Dennis, and Cain, Glen. "Statistical Theories of Discrimination in the Labor Market." *Industrial and Labor Relations Review* 30 (1977): 175–87.
- Altonji, Joseph, and Blank, Rebecca. "Race and Gender in the Labor Market." Pages 3C:3143–59 in *Handbook of Labor Economics*, edited by David Card and Orley Ashenfelter. Amsterdam: Elsevier Science, 1999.
- Altonji, Joseph, and Pierret, Charles. "Employer Learning and Statistical Discrimination." *Quarterly Journal of Economics* 116 (2001): 293–312.
- Antonovics, Kate; Arcidiancono, Peter; and Walsh, Randall. "Race and Gender Discrimination in the Weakest Link." Unpublished manuscript. Durham, N.C.: Duke University: Department of Economics, 2004.
- Arrow, Kenneth. "The Theory of Discrimination." Pages 91–100 in *Discrimination in Labor Markets*, edited by Orley Ashenfelter and Albert Rees. Princeton, N.J.: Princeton University Press, 1973.
- Ayres, Ian, and Siegelman, Peter. "Race and Gender Discrimination in Bargaining for a New Car." *American Economic Review* 85 (1995): 304–21.
- Becker, Gary S. *The Economics of Discrimination*. Chicago: University of Chicago Press, 1957.
- Bertrand, Marianne, and Mullanaithan, Sendhil. "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. Chicago: University of Chicago, Graduate School of Business, 2002.
- Coate, Stephen, and Loury, Glenn. "Will Affirmative-Action Policies Eliminate Negative Stereotypes?" *American Economic Review* 83 (1993): 1220–40.
- Fevrier, Philippe, and Linnemer, Laurent. "Equilibrium Selection: Payoff or Risk Dominance? The Case of the 'Weakest Link.'" Unpublished manuscript. Chicago: University of Chicago, Department of Economics, 2002.

- Fryer, Roland. "Economists' Models of Discrimination: An Analytic Survey." Unpublished manuscript. Chicago: University of Chicago, Department of Economics, 2001.
- Fryer, Roland, and Jackson, Matt. "Categorical Cognition: A Psychological Model of Categories and Identification in Decision Making." Unpublished manuscript. Cambridge, Mass.: Harvard University, Department of Economics, 2002.
- Gertner, Robert. "Game Shows and Economic Behavior: Risk Taking on 'Card Sharks.'" *Quarterly Journal of Economics* 108 (1993): 507–22.
- Haan, Marco; Los, Bart; Riyanto, Yohanes E.; and van Geest, Martin. "The Weakest Link: A Field Experiment in Rational Decision Making." Unpublished manuscript. Groningen: University of Groningen, 2002.
- Heckman, James, and Siegelman, Peter. "The Urban Institute Audit Studies: Their Methods and Findings." Pages 187–258 in *Clear and Convincing Evidence: Measurement of Discrimination in America*, edited by Michael Fix and Raymond Struyk. Washington, D.C.: Urban Institute, 1992.
- Howell, Susan, and Sims, Robert. "Survey Research and Racially Charged Elections: The Case of David Duke in Louisiana." *Political Behavior* 16 (1994): 219–36.
- Kahn, Lawrence. "Discrimination in Professional Sports: A Survey of the Literature." *Industrial and Labor Relations Review* 44 (1991): 395–418.
- Knowles, John; Persico, Nicola; and Todd, Petra. "Racial Bias in Motor Vehicle Searches: Theory and Evidence." *Journal of Political Economy* 109 (2001): 203–29.
- Kuklinski, James H.; Sniderman, Paul M.; Knight, Kathleen; Piazza, Thomas; Tetlock, Philip E.; Lawrence, Gordon R.; and Mellers, Barbara. "Racial Prejudice and Attitudes towards Affirmative Action." *American Journal of Political Science* 41 (1997): 402–19.
- List, John. "Friend or Foe: A Natural Experiment of the Prisoner's Dilemma." Unpublished manuscript. College Park: University of Maryland: Department of Economics, 2002.
- Lundberg, Shelly, and Startz, Richard. "Private Discrimination and Social Intervention in Competitive Markets." *American Economic Review* 73 (1983): 340–47.
- Metrick, Andrew. "A Natural Experiment in 'Jeopardy!'" *American Economic Review* 85 (1995): 240–53.
- Phelps, Edmund. "A Statistical Theory of Racism and Sexism." *American Economic Review* 62 (1972): 659–61.