

Games and Discrimination: Lessons From *The Weakest Link*

Kate Antonovics*, Peter Arcidiacono†, Randall Walsh‡§

March 24, 2003

Abstract

The earnings of both women and blacks have consistently lagged behind those of white men. Empirically determining whether these differences arise because of discrimination is extremely difficult, and distinguishing between the various theories of discrimination is harder still. This paper builds on the emerging experimental literature on discrimination and explores the use of a high-stakes game environment to reveal patterns of discrimination. We use data from the television game show, *The Weakest Link*, to determine whether contestants discriminate on the basis of race and gender and, if so, which theory of discrimination best explains their behavior. Our results suggest no evidence of discriminatory voting patterns by males against females or by whites against blacks. In contrast, we find that in the early rounds of the game women appear to discriminate against men. We test three competing theories for the voting behavior of women: preference-based discrimination, statistical discrimination and strategic discrimination. We find only preference-based discrimination to be consistent with the observed voting patterns. Finally, we demonstrate that our results are consistent with much of the existing literature on discrimination.

Key Words: Discrimination, Experiments, Games.

JEL: J7, C9, C7

*Department of Economics, University of California at San Diego. Email: kantonov@ucsd.edu.

†Department of Economics, Duke University Email: psarcidi@econ.duke.edu.

‡Department of Economics, University of Colorado at Boulder. Email: walshr@colorado.edu.

§We would like to thank Vince Crawford, Nora Gordon, Roger Gordon, Tom Nechyba, Joel Watson and seminar participants at the University of Colorado, Johns Hopkins University and the University of Pennsylvania for their helpful comments and advice concerning this paper. We also thank Patrick Dickinson and Pamela Maine for research assistance.

1 Introduction

The earnings of both women and blacks have consistently lagged behind those of white men. Empirically determining whether these differences arise because of discrimination is extremely difficult, and distinguishing between the various theories of discrimination is harder still. As a result, researchers have begun to look outside of labor markets to develop a better understanding of whether and why individuals discriminate. This paper builds on the emerging experimental literature on discrimination and explores the use of a high-stakes game environment to reveal patterns of discrimination. To this end, we use data from the television game show *The Weakest Link* to determine whether contestants discriminate on the basis of race and gender and, if so, which theory of discrimination best explains their behavior.

A number of existing papers analyze contestant behavior on television game shows in order to draw inferences about behavior in more policy-relevant settings.¹ The benefit of this type of analysis is twofold. First, in many situations, the strategic environment of game shows generates data that can be used to answer questions that cannot be easily addressed with data from more traditional settings. This is particularly relevant in the context of labor market discrimination where economists have had only limited success at distinguishing between the various theories of discrimination using standard data. Second, the stakes are much larger on game shows than in laboratory experiments and these stakes may influence both the extent and nature of discrimination. For example, Ball and Eckel (2003) find that favoritism towards members of “high status” groups in ultimatum games is sensitive to the dollar value of the prize being split.

There are a number of unique features of *The Weakest Link* that make it a particularly apt setting in which to examine discriminatory behavior. Before outlining these features, it is useful to first describe the game more generally. In *The Weakest Link*, players act cooperatively to build a pot of money by correctly answering a series of “general knowledge” questions. The game proceeds in a series of rounds, and at the end of each round of questioning, the players vote to remove one player from the game. At the end of the game, the two remaining players compete head to head to determine the sole winner of the accumulated pot of money.

One of the primary advantages of data from *The Weakest Link* is that we have excellent controls for individual ability. This is crucial because the principle difficulty in establishing

¹For example, Bennett and Hickman (1993), Berk et al. (1996), Gaines (2003), Gertner (1993), Gilbert and Hatcher (1994), Metrick (1995), and Tenorio and Cason (2002).

the presence of discrimination is that it is almost impossible to determine whether unequal outcomes arise because of discrimination or because of unobservable race and gender differences in productivity. However, in *The Weakest Link* we observe the same explicit measure of individual ability—the percentage of questions answered correctly²—that is observed by the contestants. Our ability to observe a large fraction of the information that the contestants have about one another limits the extent of the omitted variables bias problem that plagues discrimination research. However, even if there remain important unobserved characteristics, the structure of the game allows us to control for these omitted factors by using information on the voting patterns of the other contestants. In particular, if we are trying to examine the factors that determine whether player A votes against player B, we can control for the proportion of other contestants who vote against B. In this way, we can account for other factors that are not directly included in our data but that are relevant to the players’ voting decisions. Additionally, because data are collected from two slightly different versions of the show, a daily show and a weekly show, the veracity of our results can be tested against two independent samples.

The structure of the game also allows us to examine why contestants discriminate. In particular, we are able to distinguish preference-based discrimination (in which discrimination arises because people simply do not like members of certain groups) from statistical discrimination (in which discrimination arises because people use group identity as a proxy for unobserved ability). In a strategic environment such as *The Weakest Link*, we also need to consider a third type of discrimination: what we term “strategic discrimination.” The notion here is that in a strategic setting race and gender may serve as focal points for collusive behavior. For example, Holm (2000) studies a one-shot coordination game similar to the Battle of the Sexes. He finds that players appear to use gender as a focal point for coordinating actions. While the game that Holm studies is very different from the game in *The Weakest Link*, the analysis highlights the importance of ruling out strategic discrimination as an explanation for disparate outcomes by race and gender. A major contribution of this paper is to highlight various methods of distinguishing between the three theories of racial discrimination discussed above.

To look for evidence of discriminatory behavior, we evaluate the voting decisions of individual players using a series of conditional logit models. A major advantage of the conditional logit model relative to analysis of the total votes cast against a particular group is that it

²As discussed below, the game also involves individual players’ decisions about “banking” money currently in the pot vs. accepting the risk of losing this money to facilitate a potentially larger reward. These decisions provide an additional dimension of information regarding player performance.

allows us to consider interactions between the characteristics of the individuals doing the voting and the characteristics of those receiving the votes. Thus, for example, one can examine whether men are more likely to vote against women *and* whether women are more likely to vote against men. The results reveal no evidence of discrimination against either women or blacks. However, we consistently find that women are substantially more likely to vote against men in the early rounds of the game, even after controlling for a broad set of performance measures. That is, based on these findings, it appears that women discriminate against men (and in favor of women) in their voting decisions. We are able to show that neither statistical discrimination nor strategic discrimination explains this pattern. Instead, the evidence seems broadly consistent with women simply preferring to play the game with other women.

How well do these results conform with the realities of the labor market? While the fact that we find no evidence of discrimination against women or blacks may be viewed as inconsistent with a growing number of empirical studies that find evidence of labor market discrimination, this inconsistency is actually quite informative. The data reveal no significant race and gender differences in contestants' ability to answer questions correctly. As a result, even if statistical discrimination is a feature of the labor market, it is unlikely that it will arise in this game.³ Thus, given that we find no evidence of preference-based discrimination or strategic discrimination against women or blacks, our results point towards statistical discrimination as an explanation for the discriminatory patterns found in other studies. Interestingly, the notion that women have discriminatory preferences in favor of other women is reflected in a number of recent studies. We discuss these studies in greater detail below.

The remainder of the paper is organized as follows. Section 2 discusses the relevant background literature and outlines the problems associated with identifying and distinguishing between different types of discrimination. Section 3 describes the details of the game and the data that we have collected. Section 4 shows how the game show data can be used to distinguish between preference-based, statistical, and strategic discrimination. Section 5 shows that males and whites do not discriminate, but that females discriminate against males in the early rounds. Section 6 tests which of the theories of discrimination best explains female voting behavior. Section 7 relates the game show results to the labor market and Section 8 concludes.

³This holds as long as contestants have correct prior beliefs about one another's ability. We test for this possibility and find no evidence to support it.

2 Previous Work on the Nature and Extent of Racial Discrimination

2.1 Evidence From the Labor Market

It is not surprising that there exists an enormous literature on racial and gender discrimination in the labor market. Here we divide these studies into two groups: those that test for evidence of discrimination and those that attempt to distinguish between different theories of discrimination.

Although many studies look for evidence of discrimination by examining wage differences using basic regression techniques,⁴ for the sake of brevity, we focus attention on audit studies of hiring practices (i.e. Cross et al. (1990), Turner et. al (1991) and Neumark et al. (1996)). The majority of these studies find substantial evidence of discrimination against women and blacks in both the labor and housing markets. However, as Heckman and Siegelman (1992) and Heckman (1998) point out, these studies may fail to overcome the omitted variables bias that plagues standard regression techniques. In particular, the implicit assumption of audit studies is that the auditors are identical apart from their race and gender. However, the researcher who selects the auditors may not know or may not be able to observe all of the characteristics that are relevant to employers. Thus, if there is any relationship between race and gender and these unobserved components, then the differences in outcomes across these groups may not truly represent discrimination.

Two studies which circumvent this problem are Goldin and Rouse (1997) and Bertrand and Mullanaithan (2002). Goldin and Rouse examine whether the adoption of “screens” that hide the identity of auditioning musicians from judges has been responsible for the increase in the proportion of females hired at orchestras. Overall their evidence is consistent with discrimination against female musicians, and indicates that the use of the screen has been responsible for a large fraction of the increase in the number of women hired. Bertrand and Mullanaithan (2002) conduct a field experiment in which they send out a series of resumes to employers in the Boston and Chicago areas that are essentially identical apart from the name at the top of the page.⁵ They find that the response rates on resumes with white sounding names are 50 percent higher than those with African American names. While both of these provocative studies find evidence of discrimination, they do not fully reveal *why* the

⁴See Altonji and Blank (1999) for a complete discussion of these techniques.

⁵Neumark et al. also partially address this critique by emphasizing the use of resumes in the pre-interview stage.

discrimination occurs.

Within the field of economics, the two leading theories of discrimination are preference-based discrimination and statistical discrimination. In models of preference-based discrimination, employers act as if there is some cost associated with hiring workers from a particular group, and, in equilibrium, workers from these groups are paid lower wages than workers from other groups (Becker, 1957). In contrast, in models of statistical discrimination firms cannot perfectly observe worker productivity, and base their assessment of worker productivity on prior beliefs about the productivity of workers in different groups. Thus, discrimination can arise either because of group differences in average productivity or in the quality of information that firms have about workers (See, for example, Phelps(1972), Arrow (1973), Aigner and Cain (1977), Lundberg and Startz (1983), and Coate and Loury (1993)).

A handful of studies explicitly test for evidence of preference-based discrimination. Three of these, Ashenfelter and Hannan (1986), Hellerstein et al. (2002), and Kahn and Sherer (1988), all find evidence consistent with preference-based discrimination. A smaller number of studies test for evidence of statistical discrimination. For example, both Oettinger (1996) and Altonji and Pierret (2001) build models of statistical discrimination in which employers slowly learn about worker quality. Unfortunately, their models generate conflicting predictions about what differences in black-white age-earnings profiles imply about statistical discrimination.⁶

2.2 Evidence From Experiments

Complementing the empirical work on labor markets are a number of papers by both psychologists and economists that look for evidence of discrimination in laboratory settings. A common finding is that individuals display “in-group bias” in the sense that they favor members of their own group (for example, Turner (1978), Turner and Brown (1978) and Vaughn et al. (1981)). The evidence also suggests that group status plays a role in determining the outcome of bargaining games (Ball et al. (2001)). A number of papers also look for evidence of statistical discrimination in laboratory settings (for example, Anderson and Hauptert (1999), Davis (1987) and Fryer et al. (2001)). These papers consistently show that statistical discrimination tends to arise in the presence of incomplete information. Generally, the groups defined in these settings are artificial (for example, the groups might be “green” and “yellow”). One exception is Fershtman and Gneezy (2001) who analyze discrimination

⁶Oettinger interprets the earnings gap between blacks and whites increasing over the life-cycle as evidence for statistical discrimination while Altonji and Pierret interpret the same feature as evidence against statistical discrimination.

among Ashkenazic and Eastern Jews. They find discrimination against Eastern Jews in ‘trust’ games by both Ashkenazic and Eastern Jews. Conducting a series of experiments, Fershtman and Gneezy argue that this result is due to incorrect expectations regarding the ‘trustworthiness’ of Eastern Jews.

Simultaneous to our paper, Levitt (2003) also examines discrimination in *The Weakest Link*.⁷ Our approaches are very different. Levitt focuses on aggregate discrimination (total votes cast against a particular group) rather than whether discrimination against a given group depends on the demographics of those doing the voting. For instance, we are interested in determining whether women vote against men and whether men vote against women, effects which are not observable and may cancel each other out under Levitt’s approach.⁸ Levitt’s results also differ from our own because of differences in our interpretation of the strategic incentives in the game. In particular, Levitt’s methodology relies on two assumptions. First, in Levitt’s paper the extent of taste-based discrimination must be the same across all rounds of the game. However, since the implicit cost of taste-based discrimination rises as the game progresses (because one’s probability of winning the game is higher in later rounds), discriminatory outcomes due to taste discrimination should diminish over time. Second, Levitt assumes that voting incentives switch as the game progresses, so that contestants first want to vote off weak players and later want to vote off strong players. Given the structure of the game, it is not clear that the voting incentives will truly reverse. We find that although the incentives to vote off the weakest player diminish as the game progresses, at no point is the strongest player ever more likely to be voted off than the weakest player.⁹ An advantage of Levitt’s work is that he has a larger sample of daily shows and analyzes discrimination against the elderly and Hispanics.

⁷Indeed, in late December of 2002, we both found out the other was close to completing their project. The result was two January 2003 working papers: Levitt (2003) and Antonovics et al. (2003).

⁸Additionally, because we use a conditional logit, the characteristics of the other contestants affect the probability of receiving a vote and the total number of predicted votes will always equal the number of votes cast. By simply regressing the number of votes against on one’s characteristics, shows with weak contestants will predict more votes than are actually available.

⁹We also differ in the way that we treat the data. Levitt increases the precision of his estimates by pooling data across the weekly and the daily shows and across rounds. However, given that there are substantial differences in the amount of prize money available, the number of rounds per game, and the likely sample of contestants in the weekly and daily shows, we chose not to pool the data in this manner.

3 Data and Rules of the Game

We use data collected from recordings of the nationally televised game show *The Weakest Link*. There are two versions of the show, an hour long weekly show and a half-hour daily show, with both versions following the same general structure. After excluding celebrity episodes where the contestants play for charity, our data consist of 28 weekly shows and 75 daily shows.¹⁰ Each show is divided into a series of timed rounds, with the number of rounds corresponding to the number of players: eight rounds in the weekly show and six rounds in the daily show. Within each round, players are sequentially asked to answer general trivia questions where correct answers translate to an increase in the prize money. The first correct answer is worth \$1,000 in the weekly show and \$250 in the daily show. After a correct answer, a player can choose to ‘bank’ the money for the team. If the player banks, the next correct answer is again worth \$1,000 in the weekly show and \$250 in the daily show. Should the player decide not to ‘bank’, the amount of money added to the pot following a correct answer increases. However, failure to answer a question correctly leads to the loss of any unbanked money for that round. A successive chain of eight (six) correct answers with no intermittent ‘banks’ leads to an \$125,000 (\$12,500) *increase* in the pot. Money banked from each round is accumulated into a team bank.

After each round, each player votes independently as to which player he would like to remove from the show, and the player who receives the most votes must leave the game. In the event of a tie, the ‘strongest link’ chooses which player to remove from the subset of players who received the most votes. The strongest link is the player who answers the highest percentage of his or her questions correctly.¹¹ Once the field of players is reduced to two (this occurs in round 7 of the weekly show and round 5 of the daily show), these two players first accumulate prize money in the same fashion as in the earlier rounds, after which the two players compete directly against each other with the winner taking all the money in the team bank.¹²

Some concern may be raised as to how representative a group of game show contestants are of the population at large. *The Weakest Link* draws contestants from a diverse set of groups. Besides variation in gender and race, contestants come from a broad range of ages, occupations and educational backgrounds.

¹⁰The data are incomplete for some rounds due to broadcast interruptions.

¹¹Should there be a tie for the ‘strongest link,’ amount of money banked is used as the tie-breaker.

¹²In the final round, five (three) questions are asked of each contestant in the weekly (daily) show with the winner being determined by who answers the most questions correctly. All other contestants leave with nothing.

4 Discrimination and Strategic Behavior in *The Weakest Link*

The Weakest Link provides a unique environment in which to distinguish between various theories of discrimination. Playing the game well not only involves answering questions correctly, but also making astute inferences about the other players' ability to play the game. Since we observe players' decisions about who to vote off, we can determine whether race and gender are relevant factors in voting decisions even after controlling for each player's performance in the game.

Analysis of contestant behavior is complicated by the game's evolving strategic environment. In the early stages of the game, when the contestants act cooperatively to build the pot of prize money, there is a clear incentive to vote off weak players. However, voting incentives may shift as the game progresses because of the structure of the final round. As discussed above, in the final round, the two remaining players first build the pot of prize money as they do in the earlier rounds of the game. However, they then face one another in a head-to-head competition to determine the winner of the entire game. Thus, we would expect that concern about the head-to-head competition would create an incentive to vote off strong contestants in later rounds. However, it is unclear whether the incentive to vote off strong players in later rounds will *outweigh* the incentive to vote off the weak players because even in the final round the two remaining players cooperatively build the prize money.¹³ Due to this ambiguity in voting incentives, we base the majority of our analysis on the early rounds of the game where incentives are more clear-cut.

The following section presents a brief discussion of the expected voting behavior under the assumptions of a) statistical discrimination, b) preference-based discrimination, and c) strategic discrimination.

4.1 Statistical Discrimination

In models of statistical discrimination uncertainty is the root cause of discrimination. In this situation, individuals may use race or gender as a proxy for unobserved ability. However, uncertainty alone is not sufficient to generate unequal outcomes. One of two additional factors must be present; either there must exist group differences in average ability or in the accuracy of the information available to others.

¹³Indeed, as mentioned earlier, we find that although the probability of voting off the weakest link diminishes as the game progresses, players are always more likely to vote off the weakest link than the strongest link.

Because contestants are uncertain of one another's ability to play the game, uncertainty is likely to play an important role in *The Weakest Link*. As a result, contestants may rationally use race or gender as a proxy for underlying ability when attempting to assess the skills of their fellow contestants. In this setting, discriminatory voting patterns can arise either because of real group-level differences in the contestant's ability or because contestants have better information about the abilities of some players relative to others. These two sources of statistical discrimination (differences in ability vs. differences in the quality of information) have different predictions about voting patterns as the game progresses. This provides us with a method for uncovering the mechanism underlying that discrimination. That is, we can tell whether players are statistically discriminating because of perceived differences in players' average abilities or because of asymmetries in the accuracy of information that players receive about one another.

In order to see how these two "sources" of statistical discrimination differentially affect the voting patterns of players, assume that there are two groups: group A and group B. In addition, suppose that on average contestants from group A play the game better than contestants from group B, but that all players are equally capable of assessing one another's ability, regardless of group membership. Then we would expect that contestants from *all* groups would be more likely to vote against members of group B in the early rounds of the game when contestants clearly have an incentive to vote off players who they believe to be relatively weak. One simple method for examining whether this type of statistical discrimination could account for observed voting patterns is to look for group-level differences in the two most relevant factors that determine player ability: the percentage of questions correctly and decisions about whether and when to bank.

Now, suppose instead that contestants from group A and group B are equally skilled at playing the game, but that contestants from group A have less accurate information about the ability levels of contestants from group B than they do about their fellow members from group A. If this is the case, members of group A will be more likely to vote off contestants from group A than they will contestants from group B. The reason is that, because members of group A have more accurate information about contestants from their own group than they do about contestants from group B, contestants from group A will place a higher weight on mean ability when assessing the ability level of contestants from group B than they do for contestants from group A. A contestant from group B with a relatively low signal will have a higher perceived ability than a contestant from group A who has that same signal. Thus, perhaps unexpectedly, in the early rounds of the game when the incentives are such that

contestants will want to vote off weak players, if group ability distributions are the same and informational asymmetries lie at the root of statistical discrimination, then we would expect that contestants would be more likely to cast votes *against* members of their own group.¹⁴

A final possibility is that the contestant's prior beliefs about ability are simply inaccurate. It is possible, for example, that members of group A begin the game with the erroneous belief that members of group B are inferior at playing the game. In this case, we would expect to see that members of group B are more likely to be voted off in the early rounds of the game, and that this pattern would disappear as the incorrect prior is updated. Two types of analysis allow us to examine the possibility that players have such erroneous expectations. First, we examine data from episodes of the daily show for which all participants have the same occupation. Presumably members of the same occupation have better information about how particular groups will perform. If bad information is the reason that particular groups are voted off, one would expect to see a weaker pattern of discriminatory voting in shows where everyone has the same occupation. Second, if incorrect prior beliefs are driven by a characteristic that is not common to all members of the discriminated group, then the test for explicit collusion discussed below can also be used to rule out incorrect prior beliefs as a source of discriminatory voting behavior.

4.2 Strategic Discrimination

In a strategic setting such as *The Weakest Link*, collusion may also play a role in discriminatory outcomes. Two types of collusion are possible here: explicit collusion (where agreements are made between groups of contestants on their voting behavior) and implicit collusion (where no agreements are discussed but individuals play strategies that yield focal points).

Contestants meet one another prior to the game and presumably there are opportunities to discuss with one another various (quite possibly collusive) strategies for playing the game. Whether and how players are able to commit to this strategy of collusion remains an open question, however, assuming that players are able to successfully act cooperatively, coalitions may be likely to form along racial or gender lines. To test for explicit collusion, we examine whether the votes of the individual players are correlated with the votes of the same group members after controlling for the votes of all the contestants. That is, there may be unobservable characteristics which make particular individuals more likely to be voted off.

¹⁴The exact opposite would occur in the labor market where high signals would be weighed more heavily from one's own group. The difference is that, in the labor market, the hired individual will be one with the highest signal whereas in *The Weakest Link* the one with the lowest signal will be voted against.

These will be captured by the total votes cast against a person by all the other players. If, however, group coalitions are formed ahead of time, then the colluding group will not only vote against members of the other group, but against the *same* member of the other group.

Discriminatory outcomes may also result from implicit collusion where individuals naturally play particular equilibria. For example, discriminating against members of group B may be a best response for members of group A given that the other members of group A are playing a discriminatory equilibrium. These focal point equilibria should be reinforcing. Hence, if group A succeeds in removing a member of group B in round 1, then this should make it more likely that they will vote off a member of group B in round 2. We can therefore test for explicit collusion by examining how who is voted off affects voting behavior in the next round.

4.3 Preference-Based Discrimination

In the prototypical model of preference-based discrimination, members of the majority group simply dislike working with members of the minority group and act as if there is some non-monetary cost associated with hiring them. As a result, in equilibrium, workers from minority groups may earn lower wages on average.

In *The Weakest Link*, contestants from one group may dislike playing the game with contestants from some other group. Thus, if we believe that members of group A dislike members of group B, we would expect that members of group A would be consistently more likely to vote off members of group B in every round. As the game progresses, however, the probability that any single player will win the game increases. As a result, the implicit cost of preference-based discrimination *also* increases. For this reason alone, even if the contestants have discriminatory preferences, their propensity to vote off blacks will fade as the game progresses. Therefore, if preference based discrimination exists, we would expect it to diminish with each round of the show. These predictions are similar to what we would expect to see with statistical discrimination and may make distinguishing between the two theories difficult.¹⁵

¹⁵As it turns out, we are able to distinguish between the two theories as there is no difference in observed abilities across racial and gender lines, eliminating statistical discrimination from consideration.

5 Which Contestants Discriminate?

5.1 Descriptive Evidence

In order to understand the broad patterns in the data, we start by analyzing voting patterns by gender and race. Table 1 summarizes the voting behavior by round for both the daily and weekly shows. The table evaluates the voting behavior of three demographic groups: men, women, and whites¹⁶ to determine if members of these demographic groups discriminate against players who are not group members. For example, do men discriminate against women in their voting patterns? The task of discerning discrimination is complicated because even if both men and women voted randomly, men would cast more votes against women and women would cast more votes against men simply because a contestant will never vote against him or her self. To account for this problem, and the fact that the distribution of demographic types varies across episodes and rounds, we describe the voting behavior in terms of the mean “group bias statistic”. For individual i voting in round r , the “group bias statistic” is given by:

$$\text{Group Bias Statistic} = \frac{1_{ir}}{\left[\frac{G_{ir}-1}{N_r-1} \right]}. \quad (1)$$

where 1_{ir} is an indicator variable that takes on a value of one if individual i votes for a contestant from his or her group in round r . G_{ir} is the number of contestant i 's type in round r of his or her episode and N_r is the total number of contestants in round r of contestant i 's episode.¹⁷ A mean value of one for the group bias statistic implies no discrimination, a value less than one implies discrimination against the other group and a value greater than one implies discrimination against one's own group.

The descriptive evidence show some surprising results. First, there is virtually no evidence of discrimination by whites against blacks; all of the values of our discrimination statistic are indistinguishable from one. Second, there is no discrimination against women by men in the early rounds, though in round 3 of both shows men are more likely to vote against women. The most surprising result, however, is that in both the weekly and the daily show women are substantially more likely to vote against men in round 1. This pattern of women voting off men continues in rounds 2 and 3 of the daily show, though at a diminishing rate.

Obviously, these patterns may reflect race and gender differences in the average ability to play the game. In particular, if men are not as successful at playing the game well, then

¹⁶The voting behavior of blacks is not considered due to small sample size.

¹⁷Note that this statistic is only calculated for rounds in which contestant i has the option of voting for both group members and non-group members.

Table 1: Own-Group Bias in Voting Behavior

Show	Round	Females				Males				Whites			
		Sample	Group Bias Statistic [†]		Sample	Group Bias Statistic [†]		Sample	Group Bias Statistic [†]		Sample	Group Bias Statistic [†]	
			Mean	Std. Dev.		Mean	Std. Dev.		Mean	Std. Dev.		Mean	Std. Dev.
Daily	1	222	0.739*	1.141	222	0.976	1.220	331	0.949	0.634			
	2	189	0.783*	1.212	186	1.336	0.336	235	0.982	0.703			
	3	150	0.910	1.177	123	0.756*	1.205	119	0.966	0.856			
	4	66	0.909	1.003	62	0.935	1.006	68	1.000	1.007			
Weekly	1	111	0.783*	1.086	113	1.000	1.185	133	1.004	0.479			
	2	98	0.907	1.119	98	0.978	1.344	103	1.019	0.526			
	3	84	0.933	1.160	83	0.689*	1.172	82	1.037	0.500			
	4	64	1.115	1.307	72	0.806	1.162	62	0.871	0.689			
	5	44	0.784	0.252	53	1.301	1.349	32	0.703	0.851			
	6	26	1.077	0.287	22	1.091	1.019	12	0.833	1.030			

[†] The ‘group bias statistic’ is equal to $1_{ir} / \left[\frac{G_{ir}-1}{N_r-1} \right]$. Where 1_{ir} is an indicator variable that takes on a value of one if individual i votes for a contestant from his/her group in round r . G_{ir} is the number of contestant i 's type in round r of his/her episode and N_r is the total number of contestants in round r of contestant i 's episode. A mean value of one implies no discriminatory outcomes, a value less than one implies discrimination against the other group and a value greater than one implies discrimination against one's own group.

* Statistically different from 1 at the 95% level.

this may explain why they are more likely than women to receive votes in the early stages of the game. The analysis of the following section controls for these differences.

5.2 Evidence from Conditional Logits

In order to control for any possible race or gender differences in the contestants' abilities to play the game, we estimate conditional logits by show and by round of the probability contestants cast votes against other players as a function of those players' characteristics. We model the utility of player i voting for contestant j in round r as:

$$U_{ijr} = X_{ijr}\beta_r + \epsilon_{ijr}, \quad (2)$$

where β is a vector of coefficients to be estimated and ϵ_{ijr} is the unobserved preference individual i has for voting for contestant j . Included in X_{ijr} are controls for the percentage of the questions the player answered correctly in that round, whether the player was statistically the weakest link (answered the fewest questions correctly) in that round and also whether the player was the strongest link in that round. In addition, we control for the gender and races of the other contestants and whether the individual voting was of the same gender or race. Due to small sample sizes, we restrict our analysis of race to the daily show.¹⁸ We assume that the ϵ_{ijr} 's are distributed i.i.d. extreme value, implying that the probability of voting for contestant j in round r is given by:

$$P_{ijr} = \frac{\exp(X_{ijr}\beta_r)}{\sum_{k=1}^{N_r} \exp(X_{ikr}\beta_r)} \quad (3)$$

where N_r is the number of contestants in round r .

The advantage of using a conditional logit here is threefold. First, consider the probability that player i votes against player j . As the abilities of the other contestants increase, so too will the probability of voting against player j . That is, the characteristics of the all the contestants influence the probability of voting against any one contestant. Second, the conditional logit allows us to examine interactions between the voting characteristics of the individual and the other contestants. Hence, we allow men to treat women differently *and* women to treat men differently. Third, the predicted number of votes cast by each contestant is constrained to be one and therefore the total number of votes cast in a round equals the number of contestants in that round.

Results for each round are shown in the appendix, with the results for round 1 of both the daily and the weekly show given in Table 2. We have the most confidence in the round

¹⁸None of the qualitative results change if race is included in the weekly show.

1 results for three reasons. First, sample sizes are much larger here than in the later rounds. Second, strategies may change as the game progresses and may depend upon the history of play. Finally, the pool of contestants in all but the initial round is endogenously determined by the voting behavior.

The first set of rows examines the voting behavior of women. Consistent with the descriptive evidence, women appear to be more likely to cast votes against men than against other women. As shown in the appendix, this effect disappears immediately after round 1 in the weekly show while slowly diminishing in the daily show. This effect is quite large. Consider a daily show with six white contestants, three men, and three women, all with identical performance. The probability of a given female voting for a particular male contestant is 23.3% while the probability that she votes for a particular female contestant is 15.1%.

Also consistent with the descriptive evidence, the next two sets of rows show that there is no indication of discrimination by men against women or by whites against blacks. This holds true for all rounds of the game except for round 3 where in both the weekly and the daily show men are more likely to vote against women, though the effect is not statistically significant. However, this discriminatory behavior by men may be due to the discrimination against men and the corresponding selection of men who make it to round 3. We discuss these results in greater detail and relate them to other studies of labor market discrimination in Section 7.

As one would expect, for all demographic groups in the early stages of the game the higher the percentage of the questions that the player answers correctly the less likely other contestants are to cast votes against that player. However, as the game progresses, the percent correct becomes less and less important. This confirms the basic logic that players have an incentive to vote against weak players in the early rounds of the game but in the later rounds, this incentive is partially offset by the incentive to vote against stronger players.

6 Why Women Vote Against Men

This section attempts to distinguish between the three possible hypotheses for why women are more likely to vote against men: statistical discrimination, strategic discrimination, and preference-based discrimination.

Table 2: Conditional Logit Estimates of Round 1 Voting Behavior[†]

Characteristics of Other Contestants	Females Only			
	Daily Show		Weekly Show	
Male	0.433*	(0.164)	0.538*	(0.224)
Percent Correct	-2.667*	(0.470)	-2.913*	(0.506)
Weakest Link	0.305	(0.242)	0.429	(0.283)
Strongest Link	-0.016	(0.342)	0.275	(0.510)
Black	-0.145	(0.292)		
Same Race	-0.189	(0.291)		
Observations	222		111	
Characteristics of Other Contestants	Males Only			
	Daily Show		Weekly Show	
Female	-0.099	(0.150)	0.084	(0.210)
Percent Correct	-2.018*	(0.410)	-2.754*	(0.488)
Weakest Link	0.085	(0.234)	0.354	(0.263)
Strongest Link	-0.596	(0.352)	-0.728	(0.746)
Black	0.199	(0.308)		
Same Race	-0.201	(0.309)		
Observations	222		113	
Characteristics of Other Contestants	Whites Only			
	Daily Show		Weekly Show	
Female	-0.170	(0.124)		
Same Gender	-0.271*	(0.123)		
Percent Correct	-2.508*	(0.352)		
Weakest Link	0.057	(0.192)		
Strongest Link	-0.360	(0.284)		
Black	0.000	(0.144)		
Observations	349			

[†] Conditional logit estimates of the probability of an individual voting for a particular contestant. Standard errors in parenthesis.

* Statistically different from zero at the 95% confidence interval.

6.1 The Case Against Statistical Discrimination

Recall that statistical discrimination with correct priors can take two forms. In the first, the mean performance level is different across groups. For statistical discrimination of this type to explain females discriminating against males, the performance by males must be on average worse than the performance by females. Table 3 documents the average percent correct for males and females by round and by show type. There is virtually no difference in performance levels for males and females in either show in the early rounds, while in the later rounds males actually perform better than their female counterparts.¹⁹ In addition, if men are truly worse at playing the game, then both women *and men* should be more likely to vote against men in the early rounds of the game. We find no evidence of this. Hence, there is no evidence that statistical discrimination based upon mean group performance is driving females to vote against males.

The second type of statistical discrimination, where the signals on ability are more informative for one group than another, is ruled out by the fact that women are discriminating against men rather than against women. Recall that information-based statistical discrimination implies that within-group performance is more informative than out-of-group performance. Hence, poor performance is weighted more heavily when it is by a member of one's own group. Should this type of discrimination occur, women would be more likely to vote against women than against men. We find the exact opposite result.

Next, we turn to the case of incorrect prior beliefs regarding group performance. It is possible that in the early rounds of the game, women have not yet learned that men perform as well as women in the game. Our data set contains 13 episodes of the daily show in which all of the contestants have the same occupation. Should erroneous prior beliefs exist, it seems reasonable to expect that workers in the same occupation would have more informative priors on the abilities of their opposite-sex contestants than do contestants of differing occupations. In Table 4 we interact a dummy variable for the shows in which all of the contestants have the same occupation with the male indicator. Under the incorrect priors hypothesis, one would expect a negative coefficient on this interaction term. However, the coefficient on this term is positive, although insignificant—providing weak evidence that women in the same occupation are more likely to vote against men than are women in different occupations. Taken by itself, this result does not conclusively rule out the possibility of incorrect prior beliefs. However, given additional evidence against this hypothesis provided below in the

¹⁹That males perform better in later rounds may be a result of relatively low ability males being voted off earlier than relatively low ability females due to the discrimination.

Table 3: Performance by Race and Gender

Show	Round	Females			Males			Blacks			Whites		
		Sample	Percent Correct	Standard Deviation	Sample	Percent Correct	Standard Deviation	Sample	Percent Correct	Standard Deviation	Sample	Percent Correct	Standard Deviation
Daily	1	225	0.668	0.306	225	0.667	0.340	97	0.639	0.346	353	0.675	0.317
	2	186	0.630	0.341	184	0.632	0.336	79	0.657	0.325	291	0.624	0.342
	3	159	0.600	0.308	141	0.614	0.337	63	0.583	0.366	237	0.613	0.309
	4	115	0.580	0.353	110	0.659	0.292	48	0.647	0.334	177	0.611	0.325
Weekly	1	111	0.608	0.367	113	0.631	0.328	27	0.629	0.391	196	0.618	0.342
	2	98	0.613	0.311	98	0.611	0.344	23	0.563	0.343	173	0.618	0.325
	3	85	0.579	0.303	83	0.631	0.302	20	0.698	0.319	148	0.592	0.299
	4	67	0.569	0.274	73	0.655	0.308	18	0.580	0.323	122	0.619	0.292
	5	49	0.511	0.252	63	0.639*	0.246	12	0.636	0.277	100	0.577	0.254
	6	40	0.539	0.287	44	0.548	0.236	8	0.413	0.289	76	0.558	0.254

* Statistically different from female performance in the round at the 95% confidence level.

Table 4: Do Females Have Wrong Priors on Male Ability? Evidence From Same Occupation Shows [†]

Characteristics of Other Contestants	Daily Show- Females Only	
Male	0.398*	(0.178)
Male \times Same Occupation Show	0.238	(0.472)
Percent Correct	-2.649*	(0.471)
Weakest Link	0.309	(0.243)
Strongest Link	-0.026	(0.342)
Black	-0.139	(0.292)
Same Race	-0.192	(0.292)
Observations	222	

[†] Conditional logit estimates of the probability of a female voting for a particular contestant. Standard errors in parenthesis.

* Statistically different from zero at the 95% confidence level.

discussion of explicit collusion, we feel comfortable focussing our attention on other potential explanations for the observed voting patterns.

6.2 The Case Against Strategic Discrimination

We now test whether women are acting cooperatively with one another. There are two basic types of collusive behavior in which women might engage: implicit collusion and explicit collusion. Under explicit collusion, the presumption is that the women are following some agreed upon pattern of voting. Under implicit collusion, the presumption is that women have a tacit agreement to vote against men. To test for explicit collusion, we include in the female round 1 conditional logits the total votes cast for each contestant by the other contestants as well as the total votes cast by the other female contestants.²⁰ The first variable captures the fact that there are unobservable characteristics that may lead individuals to vote against particular contestants. In order for there to be explicit collusion among females, the

²⁰We also performed this analysis for the male sub-sample and for round 2 of the daily show. Gender again had no effect on the voting behavior of men while the results for round 2 are very similar to the results from round 1; female votes are not correlated except through total votes.

Table 5: Are Female Votes Correlated? Testing for Explicit Collusion[†]

Characteristics of Other Contestants	Females Only			
	Daily Show		Weekly Show	
Male	0.470*	(0.200)	0.480*	(0.242)
Percent Correct	-1.920*	(0.504)	-2.312*	(0.551)
Weakest Link	0.156	(0.262)	0.243	(0.300)
Strongest Link	0.028	(0.354)	0.198	(0.512)
Black	-0.161	(0.321)		
Same Race	-0.237	(0.321)		
Total Votes Cast By:				
Other Contestants	0.541*	(0.102)	0.223	(0.120)
Other Female Contestants	-0.054	(0.179)	-0.028	(0.206)
Observations	222		111	

[†]Conditional logit estimates of the probability of a female voting for a particular contestant. Standard errors in parenthesis.

* Statistically different from zero at the 95% confidence level.

coefficient on the latter variable, total votes cast by other women, must be positive. Results for this specification are shown in Table 5. Note that the total votes cast by other women has no more predictive power than the total votes cast by men as the coefficient on the total votes cast by other women is small and insignificant. Hence, there is no evidence of explicit collusion.

That the coefficient on female votes is small and insignificant allows us to reject many other explanations for the discriminatory behavior. In particular, it suggests that women are not discriminating on some unobserved characteristic correlated with being male as this too should have led to a positive coefficient on votes cast by women. If particular men are arrogant or make poor (in the eyes of women) banking decisions, then the votes of the other women should reflect this. Instead, the results suggest that whatever characteristics are unappealing to women are equally unappealing to men. Note that this also helps to rule out discrimination based upon bad expectations. In particular, the only bad expectations that can exist involve women expecting all men to perform equally poorly—that is, there can be

no correlation with any observable (to the voter, not to the researcher) characteristics.

We next test whether or not women are implicitly colluding to vote off men. The first evidence that this is not the case comes from the diminishing coefficient on same sex as we move to the later rounds. If implicit collusion was working, then there is no reason to stop colluding against men in the later rounds. The second indication that implicit collusion is not driving the results is that implicit collusion does not appear to be reinforcing. That is, if implicit collusion were occurring successfully (a man was voted off), then implicit collusion should be more likely to happen in the next round than if the implicit collusion did not work. Table 6 reports conditional logits for rounds 2 and 3 of both the daily and weekly shows with an indicator variable for whether a man (two men) were voted off prior to round 2 (3). Although not significant, three of the four interactions are negative, implying that, if anything, removing males makes it less likely that women will vote off men in future rounds. Hence, we find no evidence of strategic discrimination through either implicit or explicit collusion.

6.3 The Case for Preference-Based Discrimination

The only remaining explanation is that women simply prefer playing with women. Consistent with this explanation, the coefficient on male diminishes in later rounds as the price of discriminating based upon preferences increases. Further, this coefficient falls faster in the weekly show, disappearing after one round. This is consistent with the theory of preference-based discrimination since the cost of discriminating is higher in the weekly show where the total prize money is substantially larger.

We have attempted to further test this explanation by including controls for the amount of money banked (the size of the pot) at the voting stage. If preference-based discrimination exists, higher amounts of money banked should lead to less discrimination. Estimates of models of this type were mixed, with evidence that at the lowest quartile of money banked discrimination increases (consistent with the theory) but discrimination also increases at the highest quartile of money banked (inconsistent with the theory). One possible explanation for the latter is that, in order to bank a large sum of money, every contestant needs to perform well, and if every player is similarly talented at playing the game, then there is no extra cost associated with voting against men, regardless of the amount of money banked. Unfortunately, the small sample sizes make it difficult to further test this hypothesis.

Table 6: Do Females Use Males as Focal Points? Testing for Implicit Collusion[†]

Characteristics of Other Contestants	Females Only- Round Two			
	Daily Show		Weekly Show	
Male	0.538	(0.332)	-0.423	(0.413)
Male × Male Voted Off in Round 1	-0.314	(0.389)	-0.759	(0.505)
Percent Correct	-1.568*	(0.441)	-3.387*	(0.623)
Weakest Link	0.240	(0.247)	0.057	(0.308)
Strongest Link	0.053	(0.283)	-0.381	(0.774)
Black	-0.048	(0.298)		
Same Race	-0.079	(0.300)		
Observations	179		98	
	Females Only- Round Three			
	Daily Show		Weekly Show	
Male	0.242	(0.240)	0.302	(0.289)
Male × Male Voted Off in Rounds 1 & 2	-0.163	(0.394)	0.443	(0.624)
Percent Correct	-1.750*	(0.487)	-2.678*	(0.687)
Weakest Link	-0.028	(0.136)	0.342	(0.333)
Strongest Link	-0.140	(0.290)	-0.359	(0.588)
Black	0.665	(0.556)		
Same Race	-1.081	(0.553)		
Observations	156		85	

[†] Conditional logit estimates of the probability of a female voting for a particular contestant. Standard errors in parenthesis.

* Statistically different from zero at the 95% confidence level.

7 Discussion

Given that we are interested in learning about labor market discrimination, it is important to confirm that our findings are consistent with what we observe in other contexts. A number of recent audit studies (for example, Goldin and Rouse (1997) and Bertrand and Mullanaithan (2003)) have documented evidence of discrimination against women and blacks. We find no evidence of discrimination against these groups. This discrepancy, however, is revealing. As Table 3 shows, there are no significant performance differences between men and women or between blacks and whites in either the first or the second moment of the ability distribution. As a result, even if statistical discrimination is an important feature of the labor market, it would be unlikely appear in our analysis. Thus, one interpretation of the fact that we find no evidence of discrimination against either blacks or women in our data is as support for the role of statistical discrimination in the discriminatory patterns found in recent studies of the labor market.

Somewhat surprisingly, we also find evidence that that women have discriminatory preferences against men. There are a number of possible explanations for this type of behavior. First, women may dislike certain aspects of how men play the game.²¹ Women may also feel more compassionate towards women than towards men, and women may not like playing with men because they fear that they will not compete as well against men in the later rounds of the game. Gneezy, Niederle and Rustichini (2003), for example, find evidence in an experimental setting that women perform worse when competing against men than against women.²²

Evidence from other settings supports the notion that women might give preferential treatment to other women. For example, a recent analysis of gender equity in the workplace finds that while female business owners hire roughly half men and half women (52% vs 48%), male business owners hire on average 38% women and 62% men.²³ These differences persist when controlling for industry differences. There is also evidence that, all else equal, women are more likely to vote for women in political contests. For example, in her analysis of voting

²¹Though for this to be true, it would have to be how males in general play the game and not correlated with any other feature (such as how certain males speak or how certain males look). If the discrimination result was driven by females disliking the way males with particular features played the game, the test we performed for explicit collusion would have shown that female votes were correlated beyond the correlation with male votes.

²²If this were the case, we would expect discrimination to persist and even increase as the game progresses which we do not see in the data. The selection of contestants into the later rounds, however, may be why we do not see this in the conditional logit estimates.

²³National Foundation for Women Business Owners (2000).

behavior during the 1992 U.S. election, Dolan (1998) finds that women were more likely than men to vote for women candidates, even after controlling for a number of ideological, issue, and party concerns. Similar results are reported by Smith and Fox (2001) for open seat house races between 1988 and 1992. Further, these authors find candidate sex does not matter to male voters once controls for other factors are included. Finally, Derose et al. (2001) examine the relationship between patient satisfaction and the gender of emergency department physicians. They find that even after controlling for the patient's age, health status, literacy level and a number of other covariates, women gave significantly higher performance ratings to female physicians. Men's satisfaction, on the other hand, was not associated with physician gender. While beyond the scope of this paper, understanding more about the sources of women's preferences towards women is clearly an intriguing area for future research.

8 Conclusion

Understanding the nature of discrimination and its contribution to both racial and gender earnings inequality involves tackling two questions. First, we would like to know whether individuals discriminate. Second, we would like to know why individuals discriminate. In this paper we attempt to address both of these questions by examining the voting behavior of contestants in *The Weakest Link*. While the game show environment is clearly different from that of the labor market, difficulties associated with evaluating discrimination directly within the labor market motivate an analysis of behavior in more stylized settings. This research builds on the emerging experimental literature on discrimination and provides an opportunity to consider individual behavior in a high-stakes environment.

Interestingly, we find no evidence of either preference-based or strategic discrimination against either blacks or women. Statistical discrimination was ruled out in *The Weakest Link* in part because of the lack of performance differences across racial and gender lines. However, this may not be the case in the labor market and points towards statistical discrimination as a possible explanation for the discriminatory patterns against these groups found in a number of recent studies. In addition, we find that women discriminate against men in the early rounds of both the daily and weekly shows. The one theory consistent with the observed voting trends is preference-based discrimination. In other words, it appears that women simply prefer playing with women rather than with men.

In addition to revealing a number of provocative patterns in the data, our paper sheds light on the types of experimental designs that could generate the data needed to distin-

guish between statistical discrimination and preference-based discrimination. In particular, although the complex structure of *The Weakest Link* precludes us exploiting the fact that voting incentives shift as the game progresses, this feature of the game suggests that one could use changing incentives to distinguish between preference-based discrimination and statistical discrimination in a laboratory setting. Because statistical discrimination is sensitive to switches in incentives whereas preference-based discrimination is not, it would be possible to distinguish between the two hypotheses by exploiting their different implications for choice behavior.

References

- [1] AIGNER, D. and CAIN, G. (1977), "Statistical Theories of Discrimination in the Labor Market," *Industrial and Labor Relations Review*, **30**, 175-87.
- [2] ALTONJI, J. and BLANK R. (1999), "Race and Gender in the Labor Market," in Orley Ashenfelter and David Card (eds.), *Handbook of Labor Economics*, **3c** (New York, NY: Elsevier).
- [3] ALTONJI, J. and PIERRET C. (2001), "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics*, **116**, 313-350.
- [4] ANDERSON, D. and HAUPERT, M. (1999) "Employment and Statistical Discrimination: A Hands-On Experiment," *The Journal of Economics*, **25**, 85-102.
- [5] ARROW, K. (1973), "The Theory of Discrimination" in Orley Ashenfelter and Albert Rees (eds.), *Discrimination in Labor Markets* (Princeton, NJ: Princeton University Press).
- [6] ASHENFELTER, O. and HANNAN, T. (1986), "Sex Discrimination and Market Competition: The Case of the Banking Industry." *Quarterly Journal of Economics*, 149-173.
- [7] ANTONOVICS, K., ARCIDIACONO, P. and WALSH, R. (2003) "Games and Discrimination: Lessons from The Weakest Link." *UCSD Working Paper 2003-03*.
- [8] BALL, S., ECKEL, C., GROSSMAN, P. and ZAME, W. (2001), "Status in Markets," *Quarterly Journal of Economics*, **116**, 161-188.
- [9] BALL, S. and ECKEL, C. (2003), "Stars Upon Thars," Working Paper.
- [10] BECKER, G. (1957), *The Economics of Discrimination* (Chicago, IL: The University of Chicago Press).
- [11] BENNETT, R. and HICKMAN, K. (1993), "Rationality and the Price is Right," *Journal of Economic Behavior and Organization*, **21**, 99-105.
- [12] BERK, J., HUGHSON, E., and VANDEZANDE, K. (1996), "The Price is Right, but Are the Bids? An Investigation of Rational Decision Theory," *American Economic Review*, **86**, 954-70.

- [13] BERTRAND, M. and MULLAINATHAN, S. (2002), "Are Emily and Brendan More Employable than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," University of Chicago, Graduate School of Business Working Paper.
- [14] COATE, S. and LOURY, G. (1993), "Will Affirmative Action Eliminate Negative Stereotypes?" *The American Economic Review*, **83**, 1220-1240.
- [15] CROSS, H., KENNY, G. and ZIMMERMAN, W. (1990), *Employer Hiring Practices: Differential Treatment of Anglo and Hispanic Job Seekers* (Washington, DC: Urban Institute Press).
- [16] DAVIS, D. (1987), "Maximal Quality Selection and Discrimination in Employment," *Journal of Economic Behavior and Organization*, **8**, 97-102.
- [17] DEROSE, K., HAYS, R., MCCAFFREY, D. and BAKER, D. (2001), "Does Physician Gender Affect Satisfaction of Men and Women Visiting the Emergency Department?" *Journal of General Internal Medicine*, **16**, 218-226.
- [18] DOLAN, K. (1998), "Voting for Women in the Year of the Woman," *American Journal of Political Science*, **42**, 272-293.
- [19] FERSHTMAN, C. and GNEEZY, U. (2001), "Discrimination in a Segmented Society: An Experimental Approach." *The Quarterly Journal of Economics*, **116**, 351-377.
- [20] FRYER, R., GOEREE, J., and HOLT, C. (2001) "Classroom Games: Experienced-Based Discrimination," *Discussion Paper, University of Virginia*.
- [21] GAINES, B. (2003) "Trueling for Dollars: Theory and Empirics of Final-Round Voting on *The Weakest Link*", University of Illinois, mimeo.
- [22] GERTNER, R. (1993), "Game Shows and Economic Behavior: Risk Taking on 'Card Sharks'," *Quarterly Journal of Economics*, **108**, 507-521.
- [23] GILBERT, G. and HATCHER, R. (1994) "Wagering in Final *Jeopardy!*," *Mathematics Magazine*, **67**, 268-277.
- [24] GNEEZY, U., NIEDERLE, M. and RUSTICHINI, A. (2003), "Performance in Competitive Environments: Gender Differences," forthcoming *Quarterly Journal of Economics*.
- [25] GOLDIN, C. and ROUSE, C. (1997), "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians," *NBER Working Paper No. 5903*.

- [26] HECKMAN, J. (1998), "Detecting Discrimination", *The Journal of Economic Perspectives*, **12**, 101-116.
- [27] HECKMAN, J. and SIEGELMAN, P. (1992), "The Urban Institute Audit Studies: Their Methods and Findings," in Michael Fix and Raymond Struyk, *Clear and Convincing Evidence: Measurement of Discrimination in America* (Lanham, MD: Urban Institute Press).
- [28] HELLERSTEIN, J., NEUMARK, D. and TROSKE, K. (2002), "Market Forces and Sex Discrimination," *Journal of Human Resources*, **37**, 353-80.
- [29] HOLM, H. (2000), "Gender-Based Focal Points," *Games and Economic Behavior*, **32**, 292-314.
- [30] KAHN, L. and SHERER, P. (1988), "Racial Differences in Professional Basketball Players' Compensation," *Journal of Labor Economics*, **6**, 40-61.
- [31] LEVITT, S. (2003), "Testing Theories of Discrimination: Evidence from the Weakest Link," NBER Working Paper 9449.
- [32] LUNDBERG, S. and STARTZ, R. (1983), "Private Discrimination and Social Intervention in a Competitive Labor Market", *The American Economic Review*, **73**, 340-347.
- [33] MASON, C. PHILLIPS O. and REDDINGTON, D. (1991), "The Role of Gender in Non-Cooperative Games," *Journal of Economic Behavior and Organization*, **15**, 215-335.
- [34] METRICK, A. (1995) "A Natural Experiment in *Jeopardy!*," *American Economic Review*, **85**, 240-253.
- [35] NATIONAL FOUNDATION OF WOMEN BUSINESS OWNERS (2000), <http://www.nfwbo.org/Research/9-19-2000/9-19-2000.htm>.
- [36] NATIONAL FOUNDATION OF WOMEN BUSINESS OWNERS (2002), <http://www.nfwbo.org/USStateFacts/US.pdf>.
- [37] NEUMARK, D., BANK, R. and VAN NORT, K. (1995), "Sex Discrimination in Restaurant Hiring: An Audit Study," *NBER Working Paper No. 5024*.
- [38] OETTINGER, G. (1996), "Statistical Discrimination and the Early Career Evolution of the Black-White Wage Gap," *Journal of Labor Economics*, **14**, 52-78.

- [39] PHELPS, E. (1972), "The Statistical Theory of Racism and Sexism," *American Economic Review*, **62**, 659-61.
- [40] SMITH, E. and FOX, R. (2001), "A Research Note: The Electoral Fortunes of Women Candidates for Congress," *Political Research Quarterly*, **54**, 205-21.
- [41] TENORIO, R. and CASON, T. (2002), "To Spin or Not to Spin? Natural and Laboratory Experiments from The Price Is Right," *Economic Journal*, **112**, 170-95.
- [42] TURNER, J. (1978) "Social Categorization and Social Discrimination in the Minimal Group Paradigm," in Henri Tajfel (ed.), *Differentiation Between Social Groups* (New York, NY: Academic Press).
- [43] TURNER, J. and BROWN, R. (1978) "Social Status, Cognitive Alternatives and Intergroup Relations," in Henri Tajfel (ed.), *Differentiation Between Social Groups* (New York, NY: Academic Press).
- [44] TURNER, M., FIX, M. and STRUYK, R. (1991), *Opportunities Denied, Opportunities Diminished: Racial Discrimination in Hiring* (Washington, DC: Urban Institute Press).
- [45] VAUGHAN, G., TAJFEL, H. and WILLIAMS, J. (1981), "Bias in Reward Allocation in an Intergroup and an Interpersonal Context," *Social Psychology Quarterly*, **44**, 37-42.

Table 7: Conditional Logits of Voting Behavior in the Daily Show[†]

Characteristics of Other Contestants		Round 1	Round 2	Round 3	Round 4
Females Only	Male	0.433* (0.164)	0.314 (0.172)	0.183 (0.190)	0.088 (0.271)
	Percent Correct	-2.667* (0.470)	-1.573* (0.440)	-1.742* (0.490)	-0.832 (0.744)
	Weakest Link	0.305 (0.242)	0.239 (0.247)	-0.024 (0.135)	0.080 (0.354)
	Strongest Link	-0.016 (0.342)	0.049 (0.283)	-0.145 (0.290)	0.431 (0.318)
	Black	0.145 (0.292)	0.042 (0.298)	0.685 (0.555)	0.723 (0.575)
	Same Race	-0.189 (0.291)	-0.083 (0.299)	-1.102 (0.552)	-1.110 (0.579)
	Observations	222	179	156	113
Males Only	Female	0.099 (0.150)	-0.047 (0.167)	0.412 (0.226)	0.077 (0.265)
	Percent Correct	-2.018* (0.410)	-1.658* (0.428)	-2.234* (0.609)	-0.692 (0.756)
	Weakest Link	0.085 (0.234)	-0.001 (0.245)	0.037 (0.170)	0.123 (0.346)
	Strongest Link	-0.596 (0.352)	-0.452 (0.312)	-0.120 (0.353)	0.232 (0.346)
	Black	0.199 (0.308)	0.030 (0.311)	0.328 (0.434)	0.586 (0.473)
	Same Race	-0.201 (0.309)	0.034 (0.313)	0.204 (0.434)	-0.273 (0.480)
	Observations	222	179	138	110
Whites Only	Male	0.170 (0.124)	0.171 (0.135)	-0.044 (0.161)	0.162 (0.212)
	Same Gender	-0.271* (0.123)	-0.159 (0.134)	-0.295 (0.159)	0.089 (0.207)
	Percent Correct	-2.508* (0.352)	-1.818* (0.350)	-1.920* (0.455)	-0.638 (0.578)
	Weakest Link	0.057 (0.192)	-0.045 (0.197)	-0.086 (0.155)	0.092 (0.284)
	Strongest Link	-0.360 (0.284)	0.024 (0.228)	-0.203 (0.246)	0.324 (0.256)
	Black	0.000 (0.144)	0.002 (0.161)	0.000 (0.189)	0.020 (0.253)
	Observations	349	280	232	175

[†] Conditional logit estimates of the probability of an individual voting for a particular contestant. Standard errors in parenthesis.

* Statistically different from zero at the 95% confidence interval.

Table 8: Conditional Logits of Voting Behavior in the Weekly Show[†]

Characteristics of Other Contestants		Round 1	Round 2	Round 3	Round 4	Round 5	Round 6
Females Only	Male	0.538* (0.224)	-0.067 (0.235)	0.395 (0.259)	-0.170 (0.282)	0.547 (0.366)	0.018 (0.424)
	Percent Correct	-2.913* (0.506)	-3.472* (0.623)	-2.726* (0.686)	0.171 (0.770)	-0.177 (1.037)	0.333 (1.591)
	Weakest Link	0.429 (0.283)	-0.002 (0.301)	0.340 (0.334)	0.866 (0.383)	0.193 (0.433)	0.822 (0.614)
	Strongest Link	0.275 (0.510)	-0.377 (0.774)	-0.321 (0.585)	-0.788 (0.499)	-0.561 (0.470)	0.025 (0.562)
	Observations	111	98	85	66	49	39
Males Only	Female	-0.084 (0.210)	0.220 (0.230)	0.420 (0.263)	0.250 (0.287)	-0.487 (0.297)	-0.519 (0.495)
	Percent Correct	-2.754* (0.488)	-3.058* (0.602)	-2.247* (0.657)	-2.396* (0.851)	0.136 (1.095)	1.587 (1.363)
	Weakest Link	0.354 (0.263)	-0.107 (0.313)	0.346 (0.317)	0.170 (0.328)	0.379 (0.393)	1.266* (0.608)
	Strongest Link	-0.728 (0.746)	-34.060 [‡]	-0.452 (0.648)	-0.451 (0.585)	-0.039 (0.412)	0.304 (0.540)
	Observations	113	98	83	71	63	42

[†] Conditional logit estimates of the probability of an individual voting for a particular contestant. Standard errors in parenthesis.

[‡] No males voted for the strongest link in round 2 of the weekly show.

* Statistically different from zero at the 95% confidence interval.