IZA DP No. 6579

Tall or Taller, Pretty or Prettier: Is Discrimination Absolute or Relative?

Daniel S. Hamermesh

May 2012

Tall or Taller, Pretty or Prettier:
Is Discrimination Absolute or Relative?

Daniel S. Hamermesh
University of Texas at Austin,
Maastricht University, NBER and IZA

ABSTRACT

Tall or Taller, Pretty or Prettier:
Is Discrimination Absolute or Relative?

Using several microeconomic data sets from the United States and the Netherlands, and the examples of height and beauty, this study examines whether: 1) Absolute or relative differences in a characteristic are what affect labor-market and other outcomes; and 2) The effects of a characteristic change when all agents acquire more of it – become taller or better-looking. Confronted with a choice among individuals, decision-makers respond more to absolute than to relative differences among them. Also, an increase in the mean of a characteristic's distribution does not alter market responses to differences in it.

NON-TECHNICAL SUMMARY

Do people discriminate based on others' relative position in the distribution of a characteristic, or do they concentrate instead on absolute characteristics? Is it how tall you are or whether you are taller than someone else? Is it how good-looking you are, or whether you are better-looking than someone else? Using three different data sets, from the U.S. and the Netherlands, the evidence shows that absolute position matters more than relative position. Do people discriminate the same way if everybody else acquires more of a desirable characteristic, or is discrimination reduced? If everyone gets taller, do differences in height matter as much as before? If we all get better-looking, do differences in beauty matter as much as they had? Using some of the same data, plus several sets of data on heights of Dutch men, which showed striking increases between 1980 and 2010, the evidence suggests that increases in a characteristic do not alter people's ability to discriminate based on that characteristic.

Corresponding author:

Daniel S. Hamermesh
Department of Economics
University of Texas
Austin, TX 78712-1173
USA
E-mail: hamermes@eco.utexas.edu

1   Introduction

The literature on the economics of discrimination is immense, going back at least to Becker (1957).   While research in the area has mostly been empirical—concerned with measuring the *ceteris paribus* impact of an ascriptive characteristic on some economic outcome, often earnings or wages, a small theoretical literature has made additional fundamental contributions (see the summaries by Cain 1986; Altonji and Blank 1999).   With only one exception (Fryer and Jackson 2008), however, the theoretical literature appears to have been unconcerned about how agents form their views of the characteristic against which they discriminate—how they organize their impressions of the characteristic that in turn affect their treatment of members of other groups.   The lack of concern with this question in the empirical literature seems to have been complete.

That the question is generally important seems clear. How do wage differences respond to differences in height in the work force if new cohorts of workers are taller than their predecessors? How would earnings differentials that arise from differences in workers' beauty be altered if workers generally became better-looking?   How does the impact of looks on electoral success change if the distribution of candidates' looks changes? Persico et al (2004), Case and Paxson (2008), Hamermesh and Biddle (1994), Möbius and Rosenblat (2006), Benjamin and Shapiro (2009) and Berggren et al (2010) have studied the market responses of these outcomes to differences in the characteristics. None of these studies, nor any other, has considered the general question of how perceptions of the characteristic affect the outcome. With Americans, and especially northern Europeans, becoming taller, the treatment of height as an earnings-enhancing labor-market characteristic may change. To the extent that the distribution

of looks is changeable by an increasingly affluent and beauty-obsessed public, how those possible changes would affect the returns to beauty is also important.

In this study I examine these issues on a number of data sets covering several different characteristics and outcomes, with the data coming from the United States and the Netherlands. In several cases I run a "horse race" between models specifying the characteristic as absolute and those specifying it as relative—in percentiles. Where possible I estimate the kernel density of the characteristic and use kernel estimation to obtain a nonparametric representation of its impact on the outcome, thus obviating spurious results that might arise from the imposition of a particular functional form on the relationship. Whether there is a general answer—a consistent way in which agents form the perceptions that affect how a characteristic alters labor- and other market outcomes—is not clear. But by examining several characteristics in a variety of contexts I may be able to shed light on the two general questions posed above.

2 Modeling the Nature of Responses to Personal Characteristics

Write a general statistical relationship between some characteristic X~f(X) and an outcome Y as:

(1) $Y = g(X)$,

where I ignore the error term and any conditioning variables in a vector Z that might also affect Y. Throughout I assume that X is solely ascriptive and that the response of Y to it reflects discrimination. I do not inquire (and, indeed, the literature only very rarely considers) whether the response of Y to X reflects market discrimination or the productivity-enhancing effects of X. I follow the literature and assume the former.

The focus throughout is on $\partial Y/\partial X$ and how it changes in response to changes in f(X). In particular, I first examine:

(2) $\partial(\partial Y/\partial X)/\partial\sigma_X \mid \mu_X$,

that is, how the responsiveness of Y to changes in X is altered when there is a mean-preserving spread in X. The world does not appear to have generated any exogenous shocks on which we can obtain data that would allow examining this change directly. I thus need to formulate proxies for the shock in (2) that can capture the change.

One way to do this is to note that one could estimate a linear (or log-linear) specification of (1) to obtain ($\partial Y/\partial X$), as is standard practice. If a mean-preserving spread in X leaves $\partial Y/\partial X$ unchanged, a re-specification of (1) with X defined as C(X), the centiles of X, would be an inferior description of g(X) compared to the linear (or log-linear) specification. In other words, does the relationship g(X) describe the responses of Y to absolute changes in X or to changes in an agent's rank in the distribution of X? In the empirical sections, absent the desired exogenous shocks to f(X), I follow this approach to infer the shape of (2).

The second issue is the estimation of:

(3) $\partial(\partial Y/\partial X)/\partial\mu_X \mid \sigma_X$,

that is, how the responsiveness of Y to X changes when there is a variance-preserving increase in the mean of X. In this case there are three examples in which shocks unrelated to Y generated such changes. It is a simple matter either to compare estimates of (1) obtained from times when μ_X differs or to examine (1) in the presence of simultaneous exogenous shocks to μ_X across groups.

The difficulty with this entire approach is that g(X) may not be linear, not based on centiles or any other simple transformation of X, but may instead be some high-powered, non-monotonic and perhaps even discontinuous function of X. While I cannot solve this difficulty

3

generally, in some of the examples the distribution of X has sufficient support to allow kernel estimation of g(X) and thus to enable me to examine these potential problems.

3 The Impact of Changing Variance of a Characteristic

In this Section I study the responses in (2) using the example of beauty. I examine how changes in the distribution of beauty affect the impacts of looks on the success of fund-raisers for a charity, on retention as a participant in a television game show and on electoral success in a professional organization.

3.1. The Beauty of Charitable Solicitors

Landry et al (2006) conducted a field experiment in which solicitors for a charity went door-to-door seeking funds, with different treatments applied randomly to potential target households. As part of the experiment the impacts of solicitors' characteristics, including their physical attractiveness, BMI and personality traits, were also assessed (separately, not by those being solicited). In their published study the authors estimated linear regressions over all the households surveyed, of which over two-thirds contributed nothing. Here I estimate probits describing whether or not a household contributed, then tobits describing that and the amount donated. In the first set of estimated equations I use exactly the same variables as Landry et al, a vector that includes solicitors' beauty as evaluated by ten raters whose assessments were normalized and then averaged. (Thus the mean beauty was 0.06, the standard deviation 0.61.[1]) In the second set of estimates I replace each male solicitor's beauty by his percentile in the distribution of male solicitors' looks, and similarly for the female solicitors.

Table 1 presents the estimates of the two sets of equations, with Columns (1) and (3) containing the estimates that follow Landry et al by including absolute beauty measures,

---

[1]The mean differs slightly from zero because better-looking solicitors appear to have contacted more households, and each observation is a household.

Columns (2) and (4) reporting the results with the beauty variables re-specified as percentiles in the distribution of own-sex beauty. The probit estimates present derivatives showing the impacts of one-unit increases in the independent variables. The estimated standard errors are clustered on the solicitors' identification numbers. As in the original study, male beauty has negative, but quite insignificant effects, while female beauty has significant and quite large positive impacts.[2] (E.g., a one-standard deviation increase in female beauty in this sample raises the probability that a household contributes to the charity from 0.30 to 0.38.) The explanatory power of the equations containing the absolute measures of beauty exceeds that of the equations containing percentiles, although the differences are small. In this example using market-wide data there is some weak evidence that the absolute effects of the relevant characteristic dominate its relative impacts.

3.2 Beauty in a Dutch Game Show, 2002

Belot et al (2012) describe a television game show in which groups of five people answer questions posed by the quizmaster. At the end of the first round of questions the contestant who earned the most points selects one of the other four group members for expulsion from the group (and from further participation in the show). Belot et al demonstrate that, holding "productivity" (questions answered) constant, those contestants who were rated (on a 7 down to 1 scale, averaged over ten raters) as being worse-looking were more likely to be expelled. Using these data we can analyze whether the likelihood of expulsion in this round of five players was greater if being relatively worse-looking had the same effect on the probability of expulsion regardless of absolute differences in looks among the players.

---

[2]There is no effect of solicitors' beauty on the size of the contribution conditional on it being positive. The impact of female solicitors' beauty in these data works through its inducements to make some contribution.

The first column in Table 2 shows the average absolute rating of the beauty of contestants who did not "win" in the first round of the game and were thus eligible for expulsion. The average beauty ratings of these players ranged from 5.73 down to 1.70 on the seven-point scale. Since productivity and/or the "winner's" preferences for expelling other players may depend on characteristics other than beauty, I control for each player's score in the round of questions, a quadratic in the player's age, and an indicator of gender. Column (2) of this table lists conditional logit estimates of the impact of a one-unit increase in the absolute beauty rating on the probability of expulsion. Column (3) shows the estimated impact of moving up one in the ranking of beauty among the four "losers" in the round.

There is a substantial difference in the explanatory power of the absolute as opposed to the relative differences in the "losing" contestants' looks on their likelihood of expulsion. The pseudo-R^2 is noticeably higher in the conditional logits based on the specification of differences in beauty as absolute. Column (4) presents estimates with both measures (and the controls) included. A likelihood-ratio test comparing this to the conditional logit in Column (2) yields $\chi^2(1)=0.62$ (p=0.43); the same test compared to the estimates in Column (3) yields $\chi^2(1)=2.66$ (p=0.10). Although not highly significant statistically, the estimates clearly show that the "losers'" absolute beauty, not their standing in a ranking of beauty, determines how the "winner" of a round in this game treats them.

Figure 1 shows the kernel density of the average beauty ratings of the "losing" contestants in the first round of the game, and Figure 2 presents the kernel estimates of its effects on the probability of expulsion. The density is slightly right-skewed, due entirely, as the points in Figure 2 show, to the presence of one outlier whose beauty was one standard deviation above that of the second best-looking among the 276 "losing" contestants. That this outlier contestant

was expelled from the show explains the strange upturn in the expulsion-beauty kernel estimate in Figure 2. Ignoring this individual, the kernel estimate implies an especially large penalty to being very bad-looking, so that it is unsurprising that absolute differences in looks describe the relationship better than relative differences.

3.3 Economists' Beauty and AEA Elections, 1966-2004

Another example that allows examining the roles of absolute position and rank in the distribution of a characteristic is provided by the data collected and analyzed by Hamermesh (2006). The looks of each candidate in the 78 elections for office (vice-presidents and members of the Executive Committee) in the American Economic Association that were held from 1966 to 2004 were rated by a panel of four incoming economics graduate students, with the average for each rater normalized and then averaged across the four raters. In each four-person election the two candidates obtaining the most votes from the Association's membership won the election. Since candidates' pictures were mailed out with the ballots, the voters at least had the opportunity to choose (discriminate?) on the basis of the candidates' looks. The question is whether they did, and whether any impact of looks worked through the candidate being among the better-looking of the four in his/her election, or whether the absolute extent of differences among the candidates' looks is what mattered for voters' electoral choices.

Column (1) in Table 3 shows the mean, standard deviation and range of the average ratings of each candidate. That the range of the average beauty ratings is large shows that the raters were able to make fairly sharp distinctions among the candidates' looks. Figure 3 presents the kernel density of the distribution of the average ratings and suggests that the distribution of the averages (of the four standardized ratings) has an extended right tail.

As in the previous sub-section, Column (2) of the table shows the conditional logit estimates of the impact of a one standard-deviation increase in absolute beauty, in this case on the probability of winning the election; Column (3) shows the estimated impact of a one-unit increase in the rank of the beauty distribution in the election; and Column (4) includes both of these variables. Also contained in each equation is a set of controls including, most importantly, the candidate's rank in scholarly productivity among the candidates (measured by lifetime citations in the *Social Science Citation Index* up to the election year), an indicator of gender and whether the candidate had previously held or currently holds high public office.[3]

The results are qualitatively remarkably similar to those in the previous sub-section. Again the specification of beauty as absolute describes the outcome better than does the candidate's position in the ranking of beauty. The differences are not, however, as large as in the previous example. Likelihood-ratio tests comparing the conditional logit in Column (4) to the estimates in Column (2) yields $\chi^2(1)=0.21$ (p=0.64); the same test compared to the estimates in Column (3) yields $\chi^2(1)=1.64$ (p=0.20).

That the results indicate that voters respond to absolute differences in the candidates' looks is also suggested by the kernel estimates shown in Figure 4. The response is monotonically increasing over the entire range of the average beauty ratings. It is especially strong, however, as a response to increases in beauty as one approaches the upper tail of looks. This result suggests that here too, and more generally than is possible in the conditional logits, absolute beauty rather than relative position in the distribution of looks determined the outcome.

---

[3]Estimates of these equations that did not include these controls, and other estimates that include the share of citations instead of the rank in this measure, yield the same conclusions about the relative importance of absolute and relative differences in beauty.

4   The Impact of a Variance-Preserving Increase in a Characteristic's Mean

In this Section I examine three natural experiments that allow inferring the impact of an increase in the average of some characteristic that occurs without any change in its variance. I use these to examine whether and how the agents' treatment changes in order to estimate the cross-partial derivative in (3). The first experiment—a change in height in a population—occurred naturally over a period of several decades. The other two—differences in average beauty across small groups of people—were generated by apparently random mixing of the individuals to form those groups.

4.1 The Increasing Height of Dutch Men, 1981-2010

Two striking facts stand out about international differences in the distributions of human heights over time: 1) During the first half of the 20[th] century American males were the tallest in the world; 2) In the early 21[st] century Dutch males are the world's tallest; and American men have fallen (actually, gained only very slightly in height) far behind their counterparts in most northern European countries.[4] To examine how this striking change altered the relation between earnings and height in the labor market we use two Dutch data sets that allow us to study this question. The focus on the Netherlands is due to the availability of data and to the unusually large and rapid shift in the distribution of heights that occurred there.[5]

In this sub-section I use Dutch data for 1981-82, 1995 and 1996, and 2006-2010 from two sources: 1) The POLS (*Permanent Onderzoek Leefsituatie*), a household survey from which data are available beginning in 1981, which continued to collect earnings data through the mid-

---

[4]See Komlos and Lauderdale (2009) for evidence, and http://www.thedailyshow.com/watch/thu-june-21-2007/stature-of-liberty for a humorous popular presentation of this phenomenon.

[5]We restrict the analysis here to men to avoid concerns about the changing labor force participation of Dutch women. Suffice it to note, however, that Dutch women's heights in these samples also rose significantly and sharply over these periods.

1990s and which also included anthropometric data. Here I use the first two years of the POLS, 1981-82, which like later years only presented data on height in five-centimeter categories. I also merge the 1995 and 1996 POLS data to enlarge the sample for the middle of this nearly thirty-year time period; and 2) The DNB Household Survey, a panel study begun in 1993 and continuing through today. I use data from the 1995 wave, to match the POLS data for that period, and independent observations from the most recent waves of the DNB panel, 2006-10. The samples are restricted to men who are at least 150 cm tall and who are between the ages of 25 and 59 inclusive. [6]

Some evidence on the startling change in Dutch men's heights is provided in Table 4, which presents summary statistics for men 25-59 from the POLS and the DNB for each of the three periods, 1981-82, 1995-96 and 2006-10. Over this period the average heights of men in this age group increased by a highly statistically significant 4.3 cm (1.7 inches), an increase in the median height over this period to what would have been about the 75[th] percentile of Dutch men's heights in 1981-82. It is also worth noting that the means for the two different samples in the mid-1990s are very similar (although the DNB estimate is statistically significantly above the POLS estimate).

The variances are nearly the same in both POLS samples; they are also identical, but lower, in both DNB samples. It is thus fairly likely that that the variance in men's heights did not increase over this period. Assuming that all the information is correct (and the fact that average heights in both 1990s samples are quite similar is encouraging), comparing the 1981-82

---

[6] In the four years of the POLS that are used here this restriction excludes seven men. It excludes seven men from the DNB in 1995, and one man from the DNB in 2010. I supplement the 2010 sample going backward through 2006 with men who did not appear in the 2010 wave. No individual appears more than once in the 2006-10 sub-sample that we use here. Of the 872 men used in that sub-sample, 449 are observed in 2010, 136 in 2009, 83 in 2008, 109 in 2007 and 95 in 2006.

POLS to the 2006-10 DNB we can conclude that the assumption of a constant variance is reasonable.[7]

To examine the earnings-height relationship I estimate log-(annual) earnings regressions, controlling for vectors of indicators of educational attainment, a quadratic in age, and marital status, with the results presented in Table 5. The estimates for both samples from the mid-1990s are nearly identical. Also, they are remarkably similar to those for 1981-82—there was very little change in the earnings-height relationship in the POLS data between the early 1980s and the mid-1990s. As a comparison of the two set of results using the DNB data shows, however, the strong relationship observed in the mid-1990s had become smaller by 2010 (although not significantly so due to the larger standard error on the estimate for 2006-10). The earnings-height elasticity dropped from 0.85 to 0.30 in these data over this period. With a linear specification the evidence suggests that the variance-preserving increase in men's heights in the Netherlands sharply reduced the impact of height on earnings.

Is the linear specification correct, however? Because information on height was provided in five-centimeter ranges in the POLS data, there are insufficient support points in those data to present meaningful kernel densities and estimates. There are sufficient support points in the DNB data, and Figure 5 shows the kernel density of the height measure in the 1995 DNB sample. Comparing it to Figure 6, which presents the kernel density for these data for 2006-10, the densities are shaped fairly similarly, with the 2006-10 density shifted rightward, as is also suggested by a comparison of the means and standard deviations in Table 4. The kernel estimates of the log-earnings-height relationship for the mid-1990s, shown in Figure 7, do not look at all like those shown in Figure 8 for 2006-10. In the recent period the relationship rises up

---

[7]We can reject the hypothesis that the means of the two distributions are the same (t=23.54). Given the sample sizes, we cannot reject the hypothesis that the variance of the distribution was unchanged over this period.

to just short of 180 cm—below the mean height, falling thereafter over most of the density in the rest of the range. In the mid-1990s the relationship was increasing much further up the distribution of heights.

Without any further considerations the results suggest that an increase in the mean of the distribution of this characteristic reduced the cross-partial derivative in (3). It is, however, worth noting that the earnings-height relationship weakened in the most recent sample because it turned flat or even slightly negative at a height slightly above a level that in the early 1980s would have been just above the mean. One explanation for the change might be that in 2006-10 we are observing a market that had not yet adjusted to the long-run equilibrium. It seems likely that employers, presumably those making decisions that led to an earnings-height relationship in 2006-10, were older (and thus shorter) than the average Dutch male ages 25-59. If they did not perceive, or at least did not react to differences in height among potential workers who are substantially taller than they, we would observe an earnings-height relationship shaped exactly like that shown in the kernel estimation shown in Figure 8.

To examine this possibility I separated the 2006-10 DNB sample into sub-samples of workers age 25-42 and ages 43-59 and re-estimated the equation for which results were shown in the last column of Table 5. Of course, the men in the sub-sample of older workers are shorter than those in the younger sub-sample (181.05 cm vs. 184.28 cm), and, indeed, the variance is higher in the younger sub-sample. The crucial thing to note is that the parameter estimate of the impact of height on log-earnings in the older sub-sample is 0.0049, nearly identical to the estimates shown in Table 5 for the mid-1990s for all men ages 25-59.[8] Among workers in the younger sub-sample the estimated impact of height on log-earnings is -0.0039 in 2006-10.

---

[8]Re-estimating the model for the DNB 1995 sample only on workers ages 29 through 45 (presumably a random sample of those who would reach 43 through 59 in 2009, the average year in which members of the 2006-10 sample were observed), the estimated impact of height is even larger than it was for this cohort in 2009.

Apparently the market only rewarded the height of older workers, those whose distribution of heights matched more closely that of the people who were likely to have been their employers than did the right-shifted distribution of the heights of younger workers.

4.2 Varying Average Beauty in the Dutch Game Show and the AEA Elections

The winner's decision about whom to "send away" in the Dutch game show and the results of the four-person elections for office in the American Economic Association analyzed in Sections 3.2 and 3.3 provide natural experiments for inferring how a one-unit increase in a characteristic, in this case beauty, alters an outcome when the characteristic's mean changes. Columns (1) and (3) of Table 6 indicate that there is substantial variation in the average looks of "losing" contestants across games and of candidates across elections. In the game show the range of average looks is nearly five times the standard deviation of the averages, and the range of average looks in the AEA elections is also roughly that relatively large.

I assume that the average looks of participants in an individual game are independent of the outcomes. Indeed, one might imagine that, if anything, the television producers who considered the issue would attempt to keep the averages as close as possible in order to maintain audience interest in the game. That consideration suggests that any inter-game differences in the mean of looks are unplanned. The average looks of candidates in a particular AEA election are almost certainly exogenous: It is difficult to believe that the AEA's Nominating Committee, which selects the candidates, chooses sets of especially good- or bad-looking candidates to match up in a particular election in order to affect the outcome or generate additional voter interest in the process.

I expand the equations whose estimation underlay the results shown in Tables 2 and 3 to include interactions of the average looks of the candidates in the game or election with the

13

individual participant's or candidate's looks. The focus is on the interaction term—does the impact of a given difference in the looks of candidates change when the average looks in the game or election change?[9] The same controls as before are included, so that only the addition of the interaction terms distinguishes the results of estimating these equations from the results reported in Columns (2) of Tables 2 and 3.

As the estimates in Columns (2) and (4) of Table 6 show, the impacts of differences in beauty on the probability of being "sent away" are somewhat smaller when the average beauty in a game increases. The interaction term is, however, statistically insignificant (t=0.92). The impact of individual beauty on the probability of election in the AEA is slightly larger when the average candidate is better looking. The change is, however, also quite insignificant statistically (t=0.51) and substantively very minor. The appropriate inference from these results is that, if all the participants' or candidates' beauty increases by the same amount, the impacts of differences in their looks remain unaltered. A change in the mean beauty among the choices facing agents does not change the impact of differences in beauty among those choices.

5 Inferences and Conclusions

We have examined two distinct variations on the usual specifications of models measuring the impacts of personal characteristics that can be viewed as eliciting discriminatory responses by other agents in markets: 1) Whether the responses stem from agents comparing absolute or relative differences in others' characteristics; and 2) Whether changes in the average of a characteristic shared by agents who may be discriminated against affect the responses of the discriminating agents to the remaining differences. The differences between the ability of absolute and relative differences in the characteristic to characterize behavior were not large, but

---

[9]Since the average probability of being sent away is always 0.25, and that of election is always 0.5, the main effect of average looks in a game or an election cannot be included in the specification of the conditional logits.

absolute differences in the characteristic consistently described agents' discriminatory responses better than did relative differences. On the second question the results are somewhat ambiguous, although the evidence does indicate that responses to remaining differences are not changed when the average of a characteristic increases.

I have shown that, in the context of choices that discriminating agents make between well-defined small sets of individuals among whom they make simultaneous distinctions, i.e., in the studies of the impacts of beauty, the conclusions are unambiguous. The results suggest that the nature of responses to differences in ascriptive characteristics is discernible when we can examine the explicit comparisons that discriminating agents make among those against whom they discriminate. So too, the effects are fairly well determined when the distribution of a characteristic shifts with no change in its variance.

Is there any way to distinguish between the alternatives in these two questions in the more interesting context of labor markets generally rather than in the narrower contexts of charitable solicitations, a game show, elections in a professional organization or the unusually large and rapid increases in height that occurred in one country? Studies based on secondary data describing large national random samples of workers cannot make the required distinctions, as a number of attempts not reported here demonstrate.[10] One possibility would be to construct audit studies in which the characteristics of various lists of job applicants are manipulated to allow inferences about these two issues.[11] Another alternative would be to create laboratory

---

[10]I examined the impacts of absolute and relative rankings of heights on earnings in the Dutch data and the American Time Use Survey, with nearly identical results for specifications of height as absolute or relative. Similarly ambiguous results were produced for the impact of beauty on earnings in recent German data, and for estimates of the impact of immigrants' skin color on earnings (using data on skin color previously used by Hersch 2008).

[11]See, however, Heckman (1998) for a discussion of audit studies and some of their difficulties. Despite concerns about whether asking businesspeople to spend time evaluating resumés of phony applicants is even ethical, there

experiments in which groups of agents with appropriately manipulated different characteristics confront other "buying" agents (although the generalizability of any results would be questionable). Yet a third possibility would be econometric case studies of promotion choices (tournaments) in which small numbers of candidates who differ along one of the dimensions analyzed here are included. Overall, given the importance of such characteristics in determining labor-market outcomes, it would be worthwhile to understand more about how the structure of differences and changes in these characteristics alter individuals' labor-market success.

References

Altonji, J, Blank, R (1999) Race and gender in the labor market. In: Ashenfelter, O, Card, D (eds) Handbook of labor economics, volume 3C. Elsevier, Amsterdam, pp

Becker, G (1957) The economics of discrimination. University of Chicago Press, Chicago

Belot, M, Bhaskar, V, van de Ven, J (2012) Beauty and the sources of discrimination. J Hum Resources, forthcoming

Benjamin, D, Shapiro, J (2009) Thin-slice forecasts of gubernatorial elections. Rev Econ Stats 91:523-536

Berggren, N, Jordahl, H, Poutvaara, P (2010) The looks of a winner: Beauty and electoral success. J Public Econs 94: 8-15

Cain, G (1986) The economic analysis of labor market discrimination: A survey. In: Ashenfelter, O, Layard, PRG (eds) Handbook of labor economics, volume 2. North-Holland, Amsterdam, pp

Case, A, Paxson, C (2008) Stature and status: Height, ability and labor-market outcomes. J Pol Econ 116:499-532

Fryer, R, Jackson, M (2008) A categorical model of cognition and biased decision making. Berkeley Electronic J Theoretical Econs, Contributions 8:Article 6

Hamermesh, D (2011) Beauty pays. Princeton University Press, Princeton, NJ

------------------------- (2006) Changing looks and changing 'discrimination': The beauty of economists. Econ Lttrs 93:405-412

-------------------------, Biddle, J (1994) Beauty and the labor market. Am Econ Rev 84:1174-1194

Heckman, J (1998) Detecting discrimination. J Econ Perspectives 12:101-116

Hersch, J (2008) Profiling the new immigrant worker: The effects of skin color and height. J Labor Econ 26:345-386

Komlos, J, Lauderdale, B (2007) Underperformance in affluence: The remarkable relative decline in U.S. heights in the second half of the 20th century. Soc Sci Qtrly 88:283-305

Landry, C, Lange, A, List, J, Price, M, Rupp, N (2006) Toward an understanding of the economics of charity: Evidence from a field experiment. Qtrly J Econ, 121:747-782

Möbius, M, Rosenblat, T (2006) Why beauty matters. Am Econ Rev 96:222-235

Persico, N, Postlewaite, A, Silverman, D (2004) The effect of adolescent experience on labor market outcomes: The case of height. J Pol Econ 112:1019-1053

Table 1 Charitable donations, field experiment, probit and tobit estimates, N = 1754[a]

	Contribute?		Contributed Amount	
Beauty:				
Absolute:				
Male Beauty	-0.0263		-0.6005	
	(0.042)		(0.936)	
Female Beauty	0.1353		2.6563	
	(0.034)		(0.789)	
Rank:				
Male Beauty/100		-0.0739		-1.6554
		(0.076)		(1.838)
Female Beauty/100		0.2748		5.3858
		(0.070)		(1.599)
Pseudo-R^2	0.0765	0.0760	0.0227	0.0225

[a]The data are from Landry *et al* (2006). All equations include as controls a vector of attidudinal variables, indicators of the arm of the experiment and indicators for whether the solicitor was overweight or obese. Standard errors in parentheses below the parameter estimates are clustered on solicitor identification numbers.

Table 2 Descriptive statistics and conditional logit estimates, Dutch game show,
N = 276 (dependent variable is "sent away")[a]

	Mean Std. Dev. Range	Parameter Estimates	
Beauty:			
Absolute	3.52 (0.69) [1.70, 5.73]	-0.4911 (0.2551)	-0.8525 (0.5343)
Rank in round	2.50 (1.42)	-0.1346 (0.1002)	0.1627 (0.2082)
Pseudo R^2		0.1678 0.1571	0.1710

[a]The conditional logits include as controls a vector showing the rank of the person's score in the round, a quadratic in age and an indicator for gender. Standard errors in parentheses below the parameter estimates here and in Tables 3, 5 and 6.

Table 3 Descriptive statistics and conditional logit estimates, AEA elections, N=312 (dependent variable is "elected")[a]

	Mean Std. Dev. Range	Parameter Estimates		
Beauty:				
Absolute	0.00 (0.71) [-1.80,2.71]	0.3623 (0.2043)		0.5300 (0.4184)
Rank in election	2.50 (1.12)		0.1385 (0.1038)	-0.0981 (0.2123)
Pseudo R^2		0.1846	0.1795	0.1853

[a]The conditional logits include as controls a quadratic in the candidate's lifetime citations up through the year before the election, and indicators for gender, whether the person had previously held a high-level government position, was in a Top 5 economics department, was not an academic, was African-American or had won a Nobel Prize.

Table 4 Descriptive statistics, Dutch household data, men's height (in centimeters), 1981-82 -- 2006-2010[a]

	POLS 1981-82	POLS 1995-96	DNB 1995	DNB 2006-10
Mean	177.92	180.72	181.40	182.30
Std. Dev.	(7.61)	(7.62)	(7.17)	(7.17)
Range	[150, 205]	[155, 200]	[160, 206]	[159, 206]
N	2017	1926	1339	872

[a]Includes only men who are at least 150cm tall and ages 25-59.

Table 5 Regression estimates, Dutch household data, men's height, 1981-82 – 2006-10, (dependent variable is ln(annual earnings))[a]

	POLS 1981-82	**POLS 1995-96**	**DNB 1995**	**DNB 2006-10**
Height:	0.00269	0.00480	0.00467	0.00179
	(0.00105)	(0.00141)	(0.00272)	(0.00339)
Adj. R^2	0.3084	0.2334	0.1949	0.1078
N	2017	1926	1339	872

[a]The regressions use sample weights. Covariates include a vector of four indicators of educational attainment, a quadratic in age, and an indicator of marital status. Only men who are at least 150cm tall and ages 25-59 are included.

Table 6. Descriptive Statistics and Conditional Logit Estimates, Dutch Game Show and AEA Elections (Dependent Variable is "Sent Away" or "Elected")[*]

	Dutch Game Show		AEA Elections	
	Mean Std. Dev. Range	Parameter Estimates	Mean Std. Dev. Range	Parameter Estimates
Beauty:				
Own	0.00	-2.8587	0.00	0.3044
	(0.71)	(2.6098)	(0.71)	(0.2097)
	[1.50,5.73]		[-1.79,2.71]	
Average in election	3.50		0.00	
	(0.36)		(0.41)	
	[2.52,4.20]		[-1.01,0.94]	
Interaction:		0.6639		0.2522
Own x Average		(0.7242)		(0.4960)
Pseudo R^2		0.1723		0.1994
N	276		312	

[*]The conditional logits for the Dutch game show include as controls a vector showing the rank of the person's score in the round, a quadratic in age and an indicator for gender. Those for the AEA elections include as controls the person's share of total citations to the four candidates in the round, and indicators for gender, whether the person had previously held a high-level government position, was in a Top 5 economics department, was not an academic, was African-American or had won a Nobel Prize.
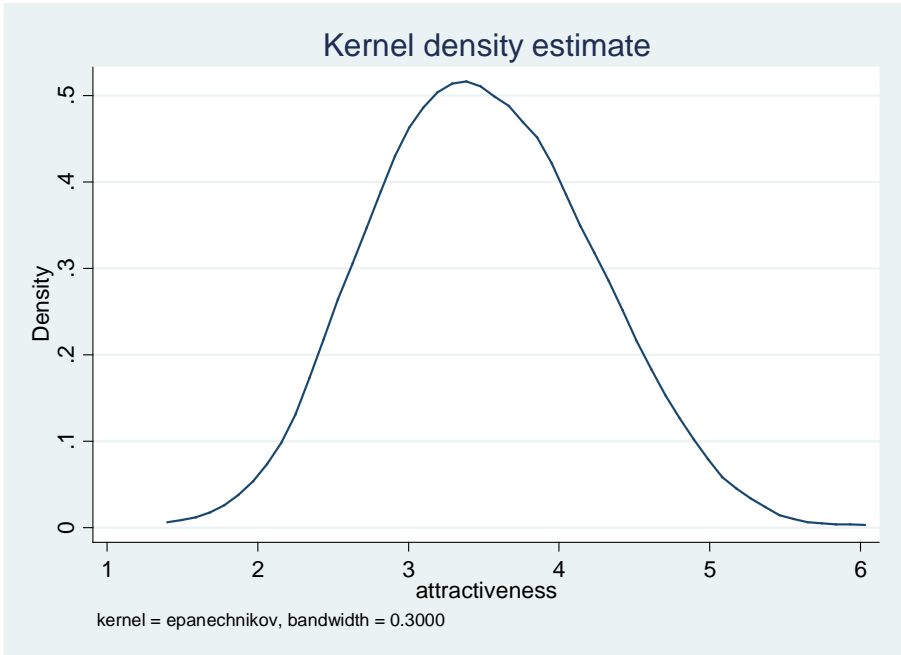
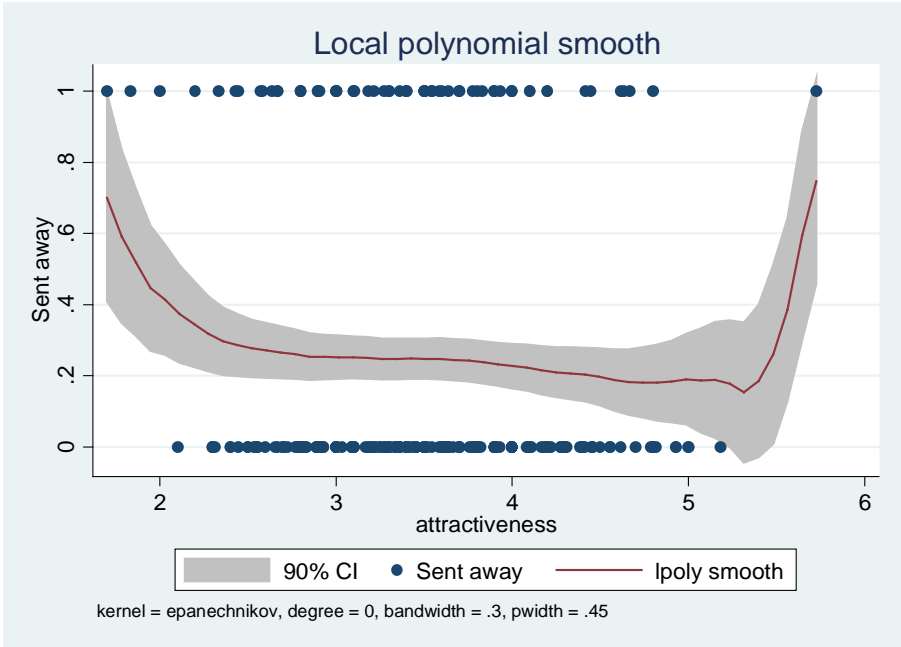**Fig. 1** Kernel of beauty density, Dutch game show, 2002

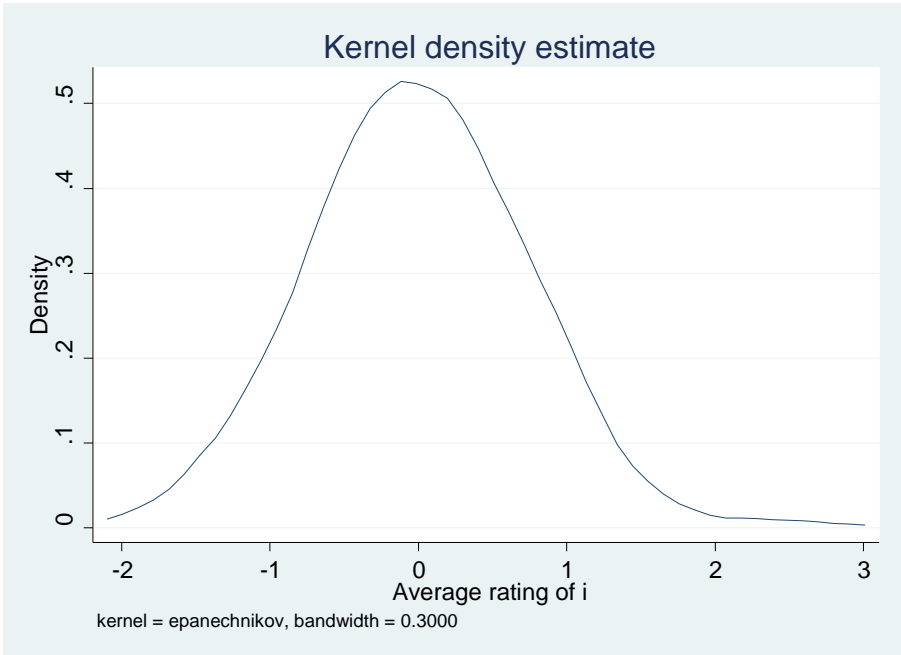**Fig. 2** Kernel estimation of the effect of beauty on expulsion probability, Dutch game show, 2002

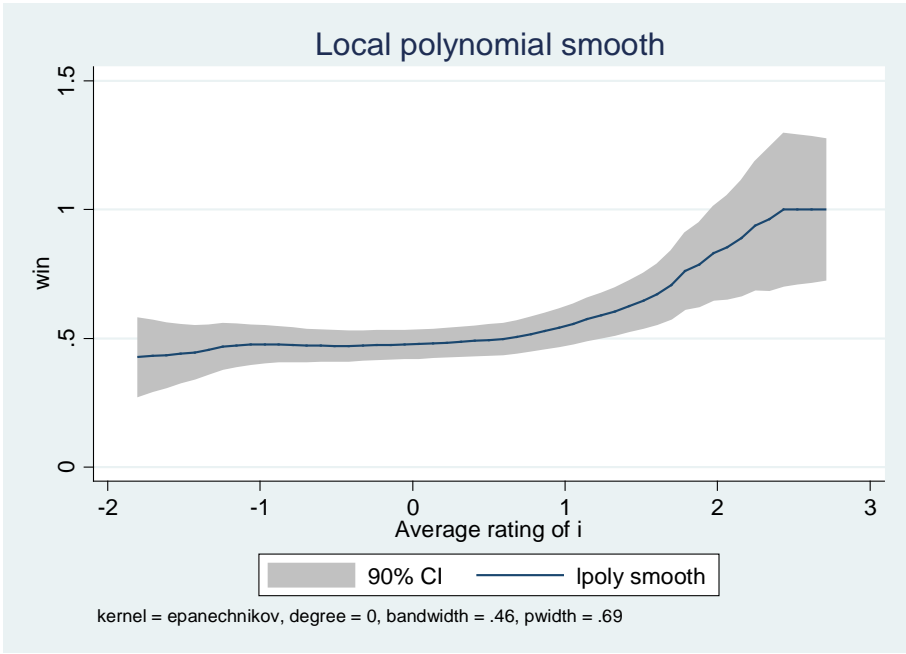**Fig. 3** Kernel of beauty density, AEA elections, 1966-2004

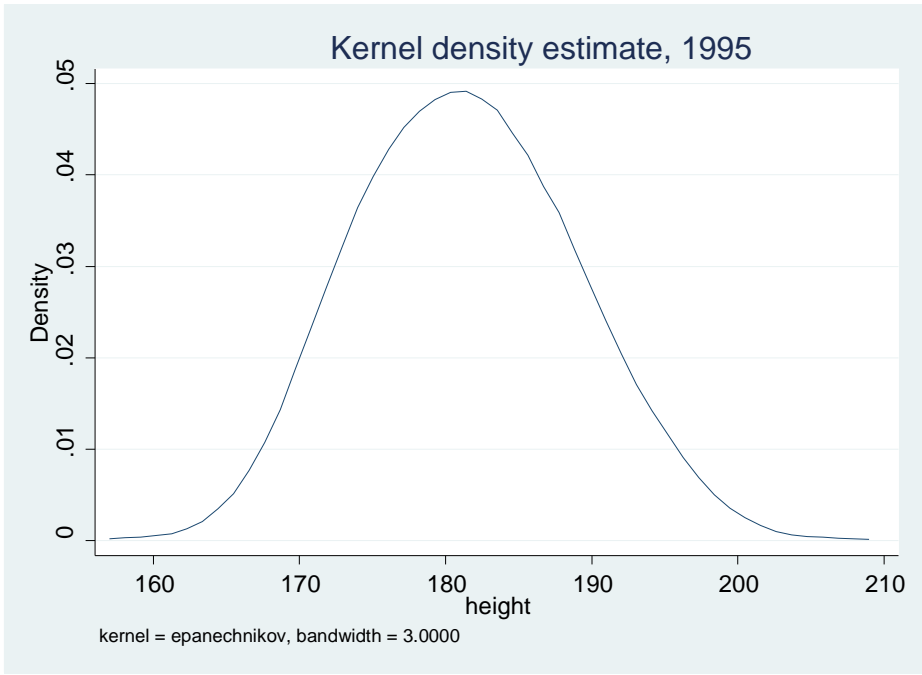**Fig. 4** Kernel estimation of the effect of beauty on win probability, AEA elections, 1966-2004

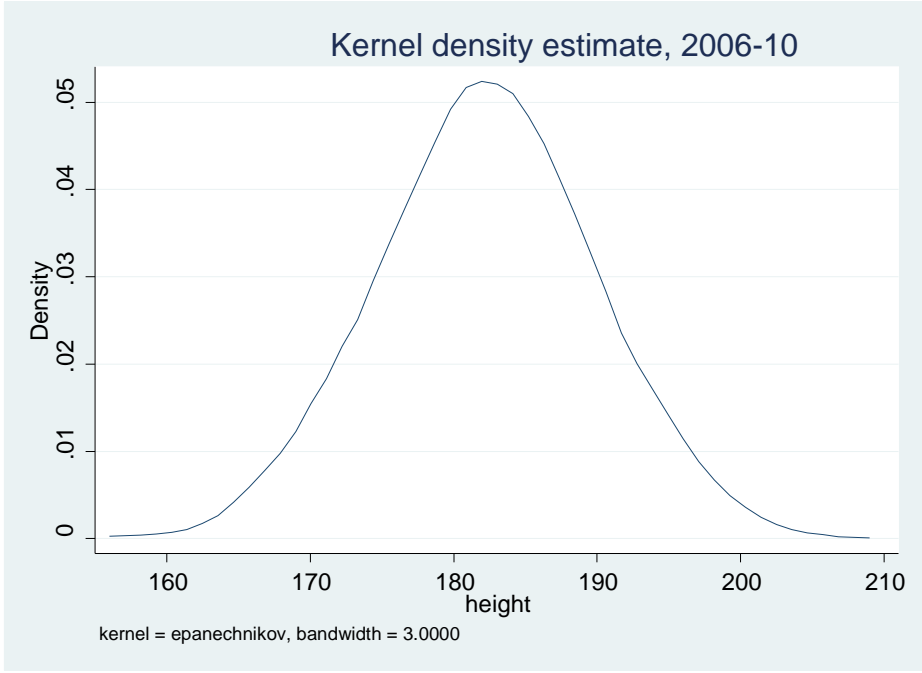**Fig. 5** Kernel of height density, DNB 1995

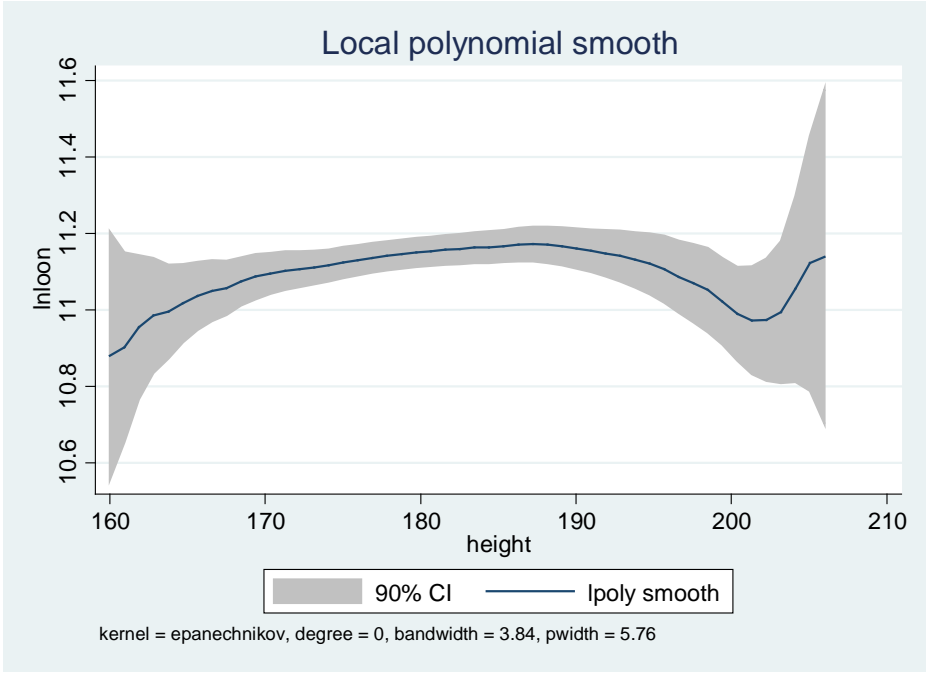**Fig. 6** Kernel of height density, DNB 2006-2010

27

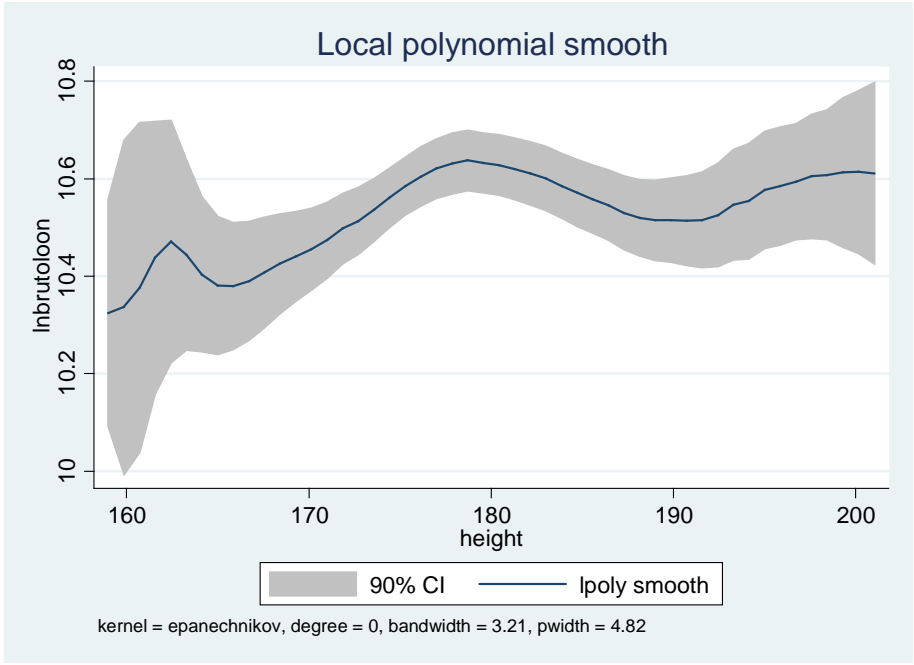**Fig. 7** Kernel estimation of the effect of height on ln(earnings), DNB 1995

Fig. 8 Kernel estimation of the effect of height on ln(earnings), DNB 2006-10