Affirmative Action Policies Promote Women and Do Not Harm Efficiency in the Laboratory

Loukas Balafoutas, *et al.*

Science **335**, 579 (2012);
DOI: 10.1126/science.1211180

Affirmative Action Policies Promote Women and Do Not Harm Efficiency in the Laboratory

Loukas Balafoutas[1] and Matthias Sutter[1,2]*

Gender differences in choosing to enter competitions are one source of unequal labor market outcomes concerning wages and promotions. Given that studying the effects of policy interventions to support women is difficult with field data because of measurement problems and potential lack of control, we evaluated, in a set of controlled laboratory experiments, four interventions: quotas, where one of two winners of a competition must be female; two variants of preferential treatment, where a fixed increment is added to women's performance; and repetition of the competition, where a second competition takes place if no woman is among the winners. Compared with no intervention, all interventions encourage women to enter competitions more often, and performance is at least equally good, both during and after the competition.

Affirmative action programs try to promote an equal representation of women in upper-level positions, both in business and in politics, as well as among members of the academic and scientific community (*1*). The reason for the existence of such programs is that, despite improvements over the past decades, there are still substantial gender differences in labor markets, both in the private sector and in the public sector, including the underrepresentation of women in the sciences. Of particular concern with respect to gender equality in labor markets are wage differentials between men and women and fewer opportunities for career advancement of women (*2*–*4*). Such gender differences are often attributed to differences in preferences regarding the type of employment, rank or position, problems in combining family and career, or to discrimination against women (*5*, *6*). A recent line of research has highlighted the contribution of another important factor—namely, the weaker inclination of women to participate in competitions (*7*–*15*). These studies provide evidence that, in general, men increase their performance in competitive environments more than women, and women more often opt out of competition, even when they are equally qualified. As a consequence of gender differences in competitive behavior, women may get fewer promotion opportunities and subsequently receive lower wages than men.

Policy interventions to support the promotion of women often face the criticism that they are inefficient in assigning the best available candidates, irrespective of gender, to a particular job when several candidates compete for it (*16*). While this is difficult to measure in the field because it is hard to exactly identify a candidate's qualifications, laboratory-based economic experiments allow for an unambiguous assessment of the efficiency of affirmative action programs in promoting the best candidates, although this is measured in the artificial context of a well-controlled, quantitative task. We evaluate and compare four alternative types of policy interventions within a unified experimental framework, examining whether they actually induce more women to compete and whether the performance of the selected winners is harmed by policy interventions. We also look at competitive behavior and teamwork after the competition ends. Competition within firms often means that one member of a working group receives a promotion but that he or she still needs to work together with the other group members afterwards, so that efficient teamwork requires the successful coordination of activities of the losers and winners of the competition. It is open to investigation, however, whether policy interventions in the spirit of affirmative action programs might backfire by spoiling the willingness of losers to coordinate efficiently in subsequent tasks.

We examine the following types of policy interventions: (i) Quotas that guarantee a certain minimum fraction of winners to be female (*1*). For instance, many European parliaments have quotas on parliamentary seats that are reserved for women. (ii) Two variants of preferential treatment of women. Preferential treatment schemes are often encountered in practice as a means to increase the participation of women in leading positions. A weak variant of preferential treatment can be a tie-breaking rule that favors women in case of equal performance or qualifications. In a stronger variant, preferential treatment may imply discrimination against better-performing men. (iii) Repetition of the competition unless a critical number of female winners is reached in the first attempt. For instance, in competitions for academic jobs (and, more generally, for jobs in the public sector) in many continental European countries—for example, in Austria—it is possible for the process of filling a vacant position to be nullified and reset to the start if no woman is shortlisted for the position. This is then equivalent to repeating the competition.

The experiment was run with 360 undergraduate students from various academic backgrounds (N = 360). Subjects were randomly assigned into groups of three men and three women [see the notes on the experimental procedure in the supporting online material (SOM)]. All groups went through several stages. The experimental task in each of the stages 1 to 4 was to add as many sets of five two-digit numbers as possible within 3 minutes (*1*). The task in stage 5 was a simple coordination game.

In stage 1 (piece rate), each subject receives €0.50 for each correct calculation. In stage 2 (tournament), group members compete against each other. The two members who solve the most calculations correctly are paid €1.50 per calculation. The other four group members receive nothing. In stage 3 (choice), subjects choose whether they want to solve the calculations under a piece-rate scheme or a tournament scheme. If the tournament is chosen, a subject's performance in stage

Table 1. Payoff matrix in the coordination game (stage 5). Each group member plays the two-person coordination game illustrated in Table 1 with each of the other five group members. This game has seven Nash equilibria that are Pareto-ranked along the diagonal. Before picking a number from 1 to 7, a subject is informed about the gender of the other player and whether this player has won or lost in the tournament of stage 4. With this information, each subject has to choose five times a number for the interaction with each of the other group members. All decisions are made simultaneously. If a subject wishes to obstruct efficient coordination, she should pick a low number. This decreases total payoffs and at the same time makes it likely that the subject has a higher payoff than her paired member. For instance, a choice of 1 guarantees the subject's own payoff of €3.50, while it implies a payoff of at most €3.50 for the other player. On the contrary, choices of high numbers create the potential for a large payoff for both subjects, provided that both coordinate on choosing a high number.

		Other person's number						
		7	6	5	4	3	2	1
Your number	7	€6.50	€5.50	€4.50	€3.50	€2.50	€1.50	€0.50
	6	€6.00	€6.00	€5.00	€4.00	€3.00	€2.00	€1.00
	5	€5.50	€5.50	€5.50	€4.50	€3.50	€2.50	€1.50
	4	€5.00	€5.00	€5.00	€5.00	€4.00	€3.00	€2.00
	3	€4.50	€4.50	€4.50	€4.50	€4.50	€3.50	€2.50
	2	€4.00	€4.00	€4.00	€4.00	€4.00	€4.00	€3.00
	1	€3.50	€3.50	€3.50	€3.50	€3.50	€3.50	€3.50

[1]University of Innsbruck, Department of Public Finance, Innsbruck, Austria. [2]University of Gothenburg, Department of Economics, Gothenburg, Sweden.

*To whom correspondence should be addressed. E-mail: matthias.sutter@uibk.ac.at

3 is compared with the other group members' performance in stage 2 (*17*). In this stage, we vary the competition rules across five treatments to examine the effects of different types of policy interventions. The treatments are described below. In stage 4 (tournament with policy intervention), all subjects compete (against each other's performance in stage 4) and the competition rules vary according to the policy intervention that is applied in each treatment. At the end of stage 4, we inform subjects about the outcome of the competition, but not about the performance of competitors, before moving on to the coordination task in stage 5 (*17*). Each winner in stage 4 receives an additional €5 as an initial endowment in stage 5, and each loser receives only €2. The reason for this unequal payment is to introduce a clear distinction between winners and losers before starting with the postcompetition stage.

The rules for determining the winners in stages 3 and 4 differ across treatments as follows: In the control treatment (CTR), the winners are the two group members with the largest numbers of correct calculations, regardless of gender. In the minimum quota treatment (QUO), there has to be at least one woman among the two winners of the tournament, irrespective of the ordinal ranking of group members' performances. This implies that the best-performing woman is always a winner. In preferential treatment 1 (PT1), each woman's performance is automatically increased by one unit (i.e., one correct calculation), whereas in preferential treatment 2 (PT2), each woman receives two extra units as a head start. Finally, in repetition of the competition (REP), the competition is repeated once if there is not at least one woman among the two winners. In this case, the rules of the repeated competition are the same as in the control treatment.

Postcompetition stage 5 (coordination game) is identical in all treatments. Each group member plays a two-person coordination game with each of the other five group members. The outcome of this coordination game, which is described in detail in Table 1, depends on players' cooperativeness and on their expectations about the cooperative choices of other players (*18*–*20*).

Figure 1 presents the average performance of men and women across all treatments as the number of correctly solved calculations in stages 1 to 4. The main pattern emerging from Fig. 1 is similar to earlier experimental evidence on the performance of men and women (*1*, *13*). On average, we find that men perform better than women, but the difference is not statistically significant (except in the tournament of stage 2; Mann-Whitney test, $z = -2.16$, $P = 0.031$). The relatively strong increase in performance from stage 1 to stage 2 is probably due to competition in stage 2, although parts of this increase might also be driven by learning effects, because Fig. 1 indicates an upward trend in performance.

Figure 2 shows the relative frequency with which men and women choose to compete in stage 3 (instead of choosing the piece-rate scheme). We find that in the control treatment CTR men compete about twice as often as women, which is consistent with the existing literature (*13*). However, the four different policy interventions reduce, and even reverse, this gender gap. The relative frequency of women opting for competition increases from 30.6% in CTR to 38.9% in REP, 52.8% in QUO, 58.3% in PT1, and 69.4% in PT2. Bilateral comparisons with the control treatment CTR reveal that the frequency of competing women is significantly higher in PT1 [$\chi^2(1) = 5.62$, $P = 0.018$] and PT2 [$\chi^2(1) = 10.89$, $P = 0.001$], whereas the increase is not quite significant in QUO [$\chi^2(1) = 3.66$, $P = 0.056$]. Repetition of the competition (REP) is the only policy intervention that has no significant effect on women's entry choices.

Although the fraction of men choosing competition is slightly affected by interventions (see Fig. 2), there is no significant difference across treatments [$\chi^2(4) = 2.67$, $P = 0.614$]. Hence, the main impact of the different policy interventions is on the choices of women and not on those of men. For a more formal analysis of the determinants of tournament entry choices and the effects of policy interventions, we have run a number of probit regressions, which confirm all of our main findings (see SOM).

Interventions that promote the entry of women may have two opposing effects on the overall efficiency in selecting the best candidates as winners. On the one hand, any intervention that gives an advantage to women may yield efficiency losses by passing over better-performing men for the sake of promoting women. On the other hand, interventions may induce more high-performing women to choose competition instead of going for the piece rate, leading to efficiency gains.

Figure 3 shows for each treatment the average ability level (measured in terms of performance) of those subjects who have entered and won the competition in stage 3 (*21*). The winners' average ability is higher than in the control treatment in two out of the four treatments with a policy intervention (QUO and PT1). However, all differences compared with the control treatment are statistically insignificant (pairwise Mann-Whitney tests, $P > 0.4$). These findings suggest that the two opposing effects discussed above cancel out in the aggregate, so that interventions do not entail efficiency costs.

Next, we examined how the various interventions affect the profile of candidates who choose to compete. We found that our four different interventions increase the likelihood of weak and strong performers entering the competition—always

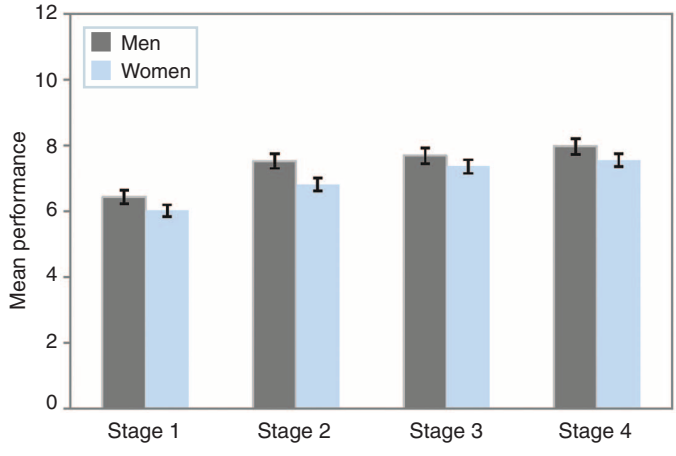

Fig. 1. Performance in the experimental task, by gender ($N = 360$ subjects). Each bar shows the average performance of participants (number of correctly solved calculations) in a particular stage, across all five treatments. Performance in treatments PT1 and PT2 excludes the extra units given to competing women. Error bars, mean ± SEM.

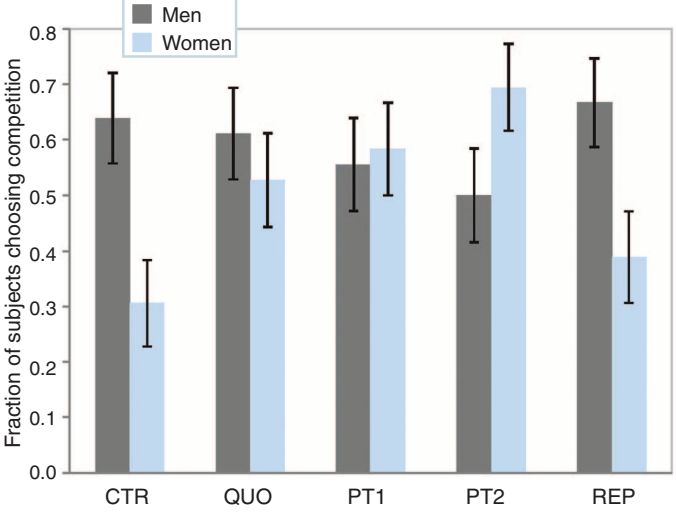

Fig. 2. Proportion of subjects choosing competition ($N = 360$ subjects; 72 per treatment). The bars show, for each treatment, the proportion of subjects (between 0 and 1) who chose the tournament compensation scheme in stage 3. Error bars, mean ± SEM.

compared with the control as a benchmark—whereas they have no effect on intermediate performers (see fig. S1). In particular, the large increase in competition entry by strong female performers shows the potential of policy interventions to improve the quality of participants. It is also encouraging to observe that strong male performers do not respond to policy interventions in a negative way.

The absence of an efficiency-decreasing selection effect is largely due to the fact that hardly any better-qualified men were passed over as a result of interventions. In treatment REP, there was no single instance in which the competition had to be repeated in stage 4, meaning that no better-qualified men could have been passed over. In treatment QUO, there were five cases where a man performed better but lost to a woman. Finally, in only one case in PT1 (PT2) a man performed better by one unit (two units) and lost.

We evaluated post-competition outcomes by means of a group's total payoff from the coordination game in stage 5, expressed as a fraction of the highest possible group payoff of €195 (Fig. 4). This measure can be considered an indicator of efficiency in a coordination task. Although efficiency differs slightly across treatments—fluctuating between 73 and 79%—the differences are not statistically significant [Kruskal-Wallis test comparing across treatments, $\chi^2(4) = 2.67$, $P = 0.614$]. Moreover, when comparing the individual choices of the winners and losers from stage 4, we do not find significant differences between the two groups in any of our treatments; this indicates that the losers of the competition do not react by obstructing efficient coordination after the competition. Nor do we detect significant gender differences in postcompetition behavior or differences between those who entered the competition in stage 3 and those who did not. We thus conclude that, in the aggregate, introducing any of our policy measures does not entail efficiency losses in a coordination task after the competition has been completed.

Studying the effects of different policy interventions on competitive behavior of men and women, both in the course of a competition and after it, is important for companies and politicians alike. Of course, companies and their human resource departments have a general interest in selecting the best candidates for a job, irrespective of gender. However, the gender composition of a company's workforce can have implications for a company's success as well (*22*). Therefore, companies may want to consider the gender of competitors in various ways, and also how different policy interventions affect postcompetition behavior. Likewise, politicians may want to provide an institutional and legal framework for a level playing field of men and women on labor markets. This requires comparing different alternative measures and their effects on behavior and efficiency, something that is difficult to examine in a controlled way with field data. Laboratory experiments are an attractive alternative methodology for the study of these issues (*23*). We provide controlled laboratory evidence on the behavioral consequences and the efficiency of different intervention schemes to promote women in competition and on how intervention schemes may affect postcompetition behavior.

The results of this study allow us to draw some lessons regarding the effects of certain affirmative action programs on individual incentives. In particular, we have looked at incentives to enter into and perform under competitive environments, as well as at incentives to coordinate efficiently in the aftermath of a competition. On the basis of our findings, we conclude that strong preferential treatment is the most successful intervention in terms of encouraging women—and, in particular, high-performing women—to enter competition. Weak preferential treatment and quotas have quantitatively similar effects in this respect, whereas repetition of the competition appears to generate the weakest incentives. Moreover, our policy interventions of minimum quotas, weak or strong preferential treatment, and repetition of the tournament do not entail any efficiency losses, neither in terms of the selection of winners in a competition nor in the postcompetition stage.

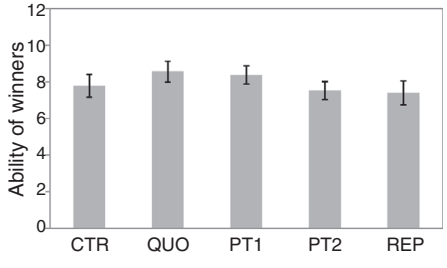

Fig. 3. Average ability of the winners in stage 3, by treatment ($N = 360$ subjects; 72 per treatment). The bars show, for each treatment, the average stage 1 performance (number of correct calculations) of those subjects who entered and won the tournament in stage 3. Error bars, mean ± SEM.

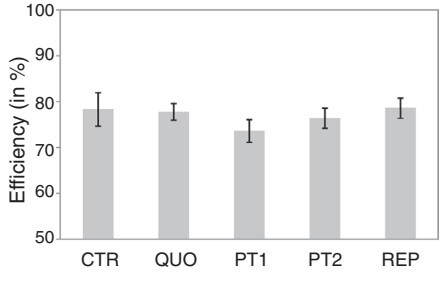

Fig. 4. Efficiency in the coordination game of stage 5 ($N = 360$ subjects; 72 per treatment). The bars show, for each treatment, the efficiency achieved on average by all groups. Efficiency is defined as the ratio of the average total group payoff in stage 5 to the maximum possible total group payoff of €195, which would emerge if all group members always chose an effort level of 7. Error bars, mean ± SEM.

Of course, many other aspects of affirmative action schemes need to be considered in policy discussions. Our study focuses on output (in stage 1 to 4) and total payoffs (in stage 5) as a measure of individual and group performance, something that companies will put great emphasis on, but it remains largely agnostic about the welfare implications of affirmative action policies because these are difficult to assess without imposing any welfare functions on individuals. Furthermore, our study has not been able to address the proper scope and extent of affirmative action policies. These aspects refer to issues such as whether affirmative action policies should apply to the public sector and the private sector to the same extent or whether governments might go too far by imposing them on private sector companies.

References and Notes

1. M. Niederle, C. Segal, L. Vesterlund, How costly is diversity? Affirmative action in light of gender differences in competitiveness (Stanford University Working Paper, 2011); available at www.stanford.edu/~niederle/AAPaper.pdf.
2. F. Blau, M. Ferber, A. Winkler, *The Economics of Women, Men and Work* (Prentice Hall, Englewood Cliffs, NJ, 2010).
3. M. Bertrand, K. F. Hallock, *Ind. Labor Relat. Rev.* **55**, 3 (2001).
4. D. Weichselbaumer, R. Winter-Ebmer, *Econ. Policy* **22**, 235 (2007).
5. J. Altonji, R. Blank, in *Handbook of Labor Economics*, O. Ashenfelter, D. E. Card, Eds. (Elsevier, 1999), vol. 3, pp. 3143–3259.
6. C. Goldin, C. Rouse, *Am. Econ. Rev.* **90**, 715 (2000).
7. B. Bartling, E. Fehr, M. Maréchal, D. Schunk, *Am. Econ. Rev.* **99** (no. 2), 93 (2009).
8. R. Croson, U. Gneezy, *J. Econ. Lit.* **47**, 448 (2009).
9. N. Datta Gupta, A. Poulsen, M. C. Villeval, *Econ. Inq.* (2012); available at http://onlinelibrary.wiley.com/doi/10.1111/j.1465-7295.2011.00378.x/abstract.
10. U. Gneezy, K. Leonard, J. List, *Econometrica* **77**, 1637 (2009).
11. U. Gneezy, M. Niederle, A. Rustichini, *Q. J. Econ.* **118**, 1049 (2003).
12. U. Gneezy, A. Rustichini, *Am. Econ. Rev.* **94** (no. 2), 377 (2004).
13. M. Niederle, L. Vesterlund, *Q. J. Econ.* **122**, 1067 (2007).
14. M. Niederle, L. Vesterlund, *J. Econ. Perspect.* **24**, 129 (2010).
15. D. Wozniak, W. T. Harbaugh, U. Mayr, The menstrual cycle and performance feedback alter gender differences in competitive choices (University of Oregon Working Paper, 2010); available at http://harbaugh.uoregon.edu/Papers/Menstruation_Wozniak_201010.pdf.
16. R. G. J. Fryer Jr., G. C. Loury, *J. Econ. Perspect.* **19** (no. 3), 147 (2005).
17. See SOM for details on this procedure.
18. J. Brandts, D. J. Cooper, *Am. Econ. Rev.* **96**, 669 (2006).
19. G. Charness, M. Jackson, *J. Econ. Theory* **136**, 417 (2007).
20. In stage 5, we have kept the hierarchy flat in the sense that the winners of stage 4 do not have superior power over losers. This is motivated by the observation that teamwork, even if there is a formal hierarchy, depends crucially on the successful coordination of activities of team members, be they former losers or winners in the competition stage. We were particularly interested in coordination in groups after a tournament has been concluded.
21. We have used the winners' performances in stage 1 as an appropriate measure of a subject's ability, because this performance is unaffected by competition and any policy intervention.
22. A. Weber, C. Zulehner, *Am. Econ. Rev.* **100** (no. 2), 358 (2010).
23. A. Falk, J. J. Heckman, *Science* **326**, 535 (2009).

Acknowledgments: Financial support from the Austrian Science Foundation (FWF Project P22772) and the D. Swarovski & Co. Förderfonds of the University of Innsbruck is gratefully

acknowledged. We thank G. Charness, N. Feltovich, U. Gneezy, M. Kocher, P. Kuhn, and P. Rey-Biel, as well as seminar participants at the University of Arizona at Tucson, the University of California at Santa Barbara, and the University of Athens and participants at the Economic Science Association meetings in Innsbruck and Melbourne, the Markets, Organization and Votes in Economics meeting in Barcelona, and the Behavioral and Experimental Economics Symposium in Maastricht, for valuable comments.

Supporting Online Material
www.sciencemag.org/cgi/content/full/335/6068/579/DC1
Materials and Methods
SOM Text
Fig. S1
Tables S1 to S3

14 July 2011; accepted 15 December 2011
10.1126/science.1211180

Female Leadership Raises Aspirations and Educational Attainment for Girls: A Policy Experiment in India

Lori Beaman,[1]* Esther Duflo,[2]* Rohini Pande,[3]* Petia Topalova[4]*

Exploiting a randomized natural experiment in India, we show that female leadership influences adolescent girls' career aspirations and educational attainment. A 1993 law reserved leadership positions for women in randomly selected village councils. Using 8453 surveys of adolescents aged 11 to 15 and their parents in 495 villages, we found that, relative to villages in which such positions were never reserved, the gender gap in aspirations closed by 20% in parents and 32% in adolescents in villages assigned a female leader for two election cycles. The gender gap in adolescent educational attainment was erased, and girls spent less time on household chores. We found no evidence of changes in young women's labor market opportunities, which suggests that the impact of women leaders primarily reflects a role model effect.

Quotas are an increasingly common response to the enduring underrepresentation of women in various domains, from science and politics to the corporate boardroom. Around 100 countries have adopted gender quotas in politics—mostly since the United Nations' Fourth World Conference on Women in 1995 (*1*)—and in 2004 Norway became the first country to mandate the presence of women on corporate boards, with a 40% quota. Policy-makers hope that quotas will have long-term effects on women's labor market outcomes over and above the immediate impact on leaders' gender balance, because the first women who become leaders may shape both parents' and children's beliefs about what women can achieve, through their policies and/or through a direct role model effect. In turn, this may raise their aspirations and shape educational and career choices. We present experimental evidence from a field setting in India that supports this view.

The idea that gender identity is at least partially a social construction is widely acknowledged in sociology (*2*) and social cognitive theory (*3*). These literatures identify belief in one's own ability (self-efficacy) as a key mechanism for personal agency and show that this belief is highly correlated with educational aspirations and subsequent occupational choices (*4*). Interventions affecting these beliefs have been shown to influence long-term behavior—for example, effort and performance in schools (*5*). Gender disparities in efficacy beliefs, in turn, are cited as an important factor behind the difference in male and female aspirations (*6*), especially in leadership (*7*, *8*). Role incongruity is often emphasized as the source of the gap in beliefs (*9*, *10*), raising the possibility that role models can challenge prevalent stereotypes and help to reduce this gap. Studies show that girls may be less likely to aspire to become scientists because there are few female scientists (*11*–*14*). Exposure to own-gender experts can provide such role models, break stereotypes regarding gender roles (*15*), and improve individual women's aspirations and propensity to enter traditionally male-dominated areas (*15*–*17*).

Several open questions remain. These include whether role model effects can counteract potential backlash (*18*) and whether other social and economic constraints prevent aspirational improvements from enabling achievement gains. Further, much of the existing evidence comes from observational studies, where beliefs or actions of individuals exposed to different role models are compared (*19*–*22*), or from laboratory or short-term school-based experiments, where individuals are exposed to different role models either in person or on paper (*11*, *16*, *23*). In observational studies, people exposed to alternative role models may themselves have different preferences or opportunities. In the laboratory or school setting, data on beliefs and aspirations are often gathered after a brief exposure to a role model (*24*); this may not reflect real life, where exposure is more enduring. Similarly, the consequences of psychological processes on behavior are generally followed over a brief period [see (*5*) for an exception].

In this study, we took advantage of a large-scale randomized natural experiment in India, which quickly increased the number of women in leadership positions at the village level. Exploiting a rule that requires the random selection of villages where only women could compete for this position, we compared the aspirations of parents for their sons and daughters aged 11 to 15, as well as the aspirations of adolescents for themselves across villages with female and male leaders. We next examined whether changes in girls' aspirations were accompanied by changes in educational outcomes and time spent on domestic chores. Finally, we investigated the potential channels for this effect.

Since 1993, when India adopted gender quotas for elected positions on village councils, the gender balance of village leadership has markedly altered: Across India, the fraction of elected local leaders who are female has risen from less than 5% in 1992 to more than 40% by 2000 (*25*). In West Bengal, our area of study, as in most Indian states, one-third of village councils are randomly selected to be reserved for a woman chief councilor (or "pradhan") in every election. Random selection of reserved councils enables us to identify the causal relationship between the election of female leaders and villagers' aspirations: There should be no other difference between villagers living in councils reserved once or twice versus unreserved councils. Moreover, because of concurrent reservations for historically disadvantaged groups (scheduled caste and scheduled tribes), seats are not set aside to women by pure rotation: Between 1998 when the system was implemented in West Bengal and 2007 when we collected our data, a village council could have been reserved for a female pradhan once (in 1998 or 2003), twice (in 1998 and 2003), or never (*26*), giving us the opportunity to study the "dose response" in exposure to a female leader. Previous research has shown that women leaders invest in goods preferred by women, such as drinking water (*27*), and improve men's perception of women's leadership abilities (*28*).

Our survey team collected data in 2006 and 2007 in one West Bengal district, Birbhum, a rural and poor district located about 200 km from Kolkata. Our survey covered a random sample of 495 villages in Birbhum's 165 village councils. Tables S1 and S2 compare 1991 census data (before reservation) and household survey data according to reservation status of the council. Because of the randomization, councils have very similar characteristics. We interviewed 15 randomly selected households per village, administering separate questionnaires to the youngest married couple in each household and to all adolescents (aged 11 to 15) residing in the household. The final sample of respondents with adolescent children consisted of 2335 male and 2438 female respon-

[1]Department of Economics, Northwestern University, Evanston, IL 60208, USA. [2]Department of Economics, Massachusetts Institute of Technology (MIT), Cambridge, MA 02142, USA. [3]Harvard Kennedy School, Harvard University, Cambridge, MA 02138, USA. [4]International Monetary Fund, Washington, DC 20434, USA.

*To whom correspondence should be addressed. E-mail: l-beaman@northwestern.edu (L.B.); eduflo@mit.edu (E.D.); rohini_pande@ksg.harvard.edu (R.P.); ptopalova@imf.org (P.T.)