Bias at the Ballot Box? Testing Whether Candidates' Gender Affects Their Vote*

Amy King, *Oxford University*

Andrew Leigh, *Australian National University*

Objective. This study examines whether women's electoral fortunes in Australia have improved in line with changing social norms over the past century. We use new strategies to explore whether female candidates face discrimination by the voting public, or by political parties' preselection systems. *Methods.* Using data from all elections to the House of Representatives between 1903 and 2004, we examine the relationship between candidates' gender and their share of the vote. We consider the electoral performances of female independent candidates, female incumbents, and female candidates from the Australian Labor Party (after 2001) in order to determine whether the bias against female candidates is driven by voters or preselectors. We also make use of gender pay gap and attitudinal data to examine how the ballot box penalty has shifted in line with changing social norms. *Results.* We find that the vote share of female candidates is 0.6 percentage points smaller than that of male candidates (for major parties, the gap widens to 1.5 percentage points), but find little evidence that the party preselection system is responsible for the voting bias against women. Over time, the gap between male and female candidates has shrunk considerably as a result of changes in social norms (as proxied by the gender pay gap and attitudinal data) and the share of female candidates running nationwide. *Conclusions.* A statistically significant gender penalty has been a consistent feature of Australian federal elections since 1903. The penalty against female candidates has narrowed since the 1980s, and this bias lies with the voting public rather than with the political parties themselves. We find little evidence that party-based affirmative action policies have reduced the gender penalty against female candidates.

In 1902, Australia became the first country to allow women both the right to vote in and stand for national elections.[1] Although there have been a

*Direct correspondence to Andrew Leigh, Research School of Social Sciences, Australian National University, ACT 0200, Australia ⟨andrew.leigh@anu.edu.au⟩. The second-named author will share all data and coding information with those wishing to replicate the study. The authors are grateful to Clive Bean, Rachel Gibson, David Gow, Sophie Holloway, and Ian McAllister for providing them with access to the restricted-use version of the Australian Candidate Study. Jennifer Chesters, Susanne Schmidt, and Elena Varganova provided outstanding research assistance. The authors also benefited from feedback from seminar participants at the Australian National University.

[1]Although New Zealand was the first country to grant women voting rights in 1893, New Zealand did not permit women to stand for election until 1919. In Australia (non-Indigenous) women's suffrage was granted through the Commonwealth Franchise Act 1902.

SOCIAL SCIENCE QUARTERLY, Volume 91, Number 2, June 2010
© 2010 by the Southwestern Social Science Association

number of significant electoral milestones since that time, Australian women are still significantly underrepresented in national parliament in the first decade of the 21st century. After the 2004 election, the share of women in the Australian House of Representatives was 25 percent. According to a recent ranking by the Inter-Parliamentary Union, this placed Australia 36th highest out of 189 countries in terms of its representation of women in national lower houses (IPU, 2007).

Given the apparent paucity of women elected to the Australian Parliament over the past century, it is important to determine precisely how women have fared in Australian elections to the federal House of Representatives since Federation. These elections are based on single-member electorates, which are effectively two-party contests between Australia's major parties in most cases.[2] The Australian test case is also an interesting one, as voting has been compulsory since 1924 and around 95 percent of its citizens vote in each federal election (AEC, 2006). We ask two main questions in this article. First, have women's electoral fortunes improved in line with other gender equality milestones such as equal pay and anti-discrimination laws over the past century? Second, do female candidates face discrimination by the voting public, or is gender voting bias merely a function of the way the parties' preselection process operates?

To answer these questions, we examine the relationship between a candidate's gender and his or her electoral success, using data from federal elections to the Australian House of Representatives between 1903 and 2004. We find that female candidates *do* fare worse than their male counterparts in Australian elections, although the penalty against female candidates has shrunk over time. One possible interpretation of our results is that they are merely driven by preselector bias. However, we conclude that this is not the case (this is not to say that preselector bias is nonexistent, but simply that it does not affect our estimates of voter bias). Furthermore, we find that over the past century, changes in voter bias against women seem to be partly explained by social norms (as proxied by the gender pay gap) and by the share of candidates who are female. We also find that across electorates, women tend to perform better when there are fewer other women on the ballot paper, and do not benefit from running in electorates with more female voters.

The remainder of this article is organized as follows. The next section presents a review of the key literature. We then turn to electoral data,

However, women in South Australia and Western Australia voted in the 1901 federal election. For a detailed history of the representation of women in Australian politics, see Sawer (2001).

[2]There are three major political parties in Australian politics. The main left-wing party is the Australian Labor Party, and the two right-wing parties are the city-based Liberal Party of Australia and the rural National Party. The two right-wing parties operate in coalition with one another, meaning that each agrees not to run candidates against a sitting member of the other party.

analyzing election results from all Australian federal elections. The following section addresses "the preselection problem" by attempting to separate voter bias from preselector bias. Next, we analyze several factors that might explain trends in voter bias against women over time and across electorates. The final section concludes.

Existing Research on Candidate Gender and Voting

Little analysis on the relationship between gender and electoral outcomes has taken place in Australia; however, several U.S. studies have sought to explain the lack of female elected representatives in terms of either discriminatory voting behavior by the public, or discriminatory preselection practices by political parties themselves.

The first body of literature related to this study examines the possibility that voters display discriminatory voting behavior based on candidate gender. A range of U.S. studies (Darcy and Schramm, 1977; Ekstrand and Eckert, 1981; Huddy and Terkildsen, 1993a; Berch, 2004) have experimented with "changing" the gender of a particular candidate in order to determine the effect on voters' choice of candidate. Each found that voters display no systematic bias for or against female candidates. However, while Fox and Smith's (1998) experimental study also found that voters in California did not display evidence of gender bias, a simultaneous experiment in Wyoming found that voters gave an additional 10 percent of the vote to the male candidate. In seeking to explain some of the bias against female candidates in real-world elections, Kenworthy and Malami (1999), Inglehart and Norris (2003), Kittilson (2006), and Paxton and Hughes (2007) concluded that society's attitudes toward women had a pervasive effect on their numbers in Parliament or Congress.

A second body of literature finds that voters discriminate on the grounds of gender as they tend to view men and women as better suited to particular elected roles. That is, male candidates are more successful in elections to national or executive offices and those that require decision making around the economy or defense, while female candidates fare better when running for "nurturing" roles in local government or in "feminine" policy areas such as education and childcare (Huddy and Terkildsen, 1993b; Fox and Oxley, 2003; King and Matland, 2003; Herrnson, Lay, and Stokes, 2003).

However in contrast to the first two literatures, many studies blame political parties themselves for an internal bias against women in the candidate-selection process (Darcy and Schramm, 1977; Kelley and McAllister, 1984; Bean and McAllister, 1990; Caul, 2001; Conway, 2001; Kittilson, 2006). In Australia, the left-wing minor parties have been far more likely to nominate women than the major parties; in 1990, 44 percent of the Green candidates and 30 percent of the Democrat candidates were women, compared with just 15 percent of candidates nominated by the ALP or Liberal Party (Bean

and McAllister, 1990). Furthermore, the major parties have tended to nominate their female candidates for the most unwinnable seats (Mackerras, 1977). Bean and McAllister (1990) showed that in the 1990 federal election, 56 percent of female Liberal Party candidates were nominated for safe (i.e., "unwinnable") Labor seats, while 21 percent of female ALP candidates ran for safe Liberal-National Party seats.

Given this distinction in the literature between the effect of voter bias and party-based preselection bias on female electoral outcomes, this article seeks to separate these two effects in the Australian context. As far as we are aware, ours is the first article to study the effects of gender on electoral outcomes in all elections to the Australian House of Representatives over the past century.

How Do Female Candidates Fare in Australian Elections?

We now turn to the first question in our study—an analysis using data from all 40 federal elections to the Australian House of Representatives conducted between 1903 and 2004 (as we explain below, we use the 1901 election to calculate an expected vote share measure, so observations from 1901 appear in Figures 1 and 2, but not in the regression analysis and subsequent graphs). So far as we are aware, our article is the first to use data on every individual electoral race since Federation. This is most likely because the data have not previously been available in a comparable electronic form. In our case, we have benefited from electronic spreadsheets provided

FIGURE 1

Share of Candidates Who Were Women

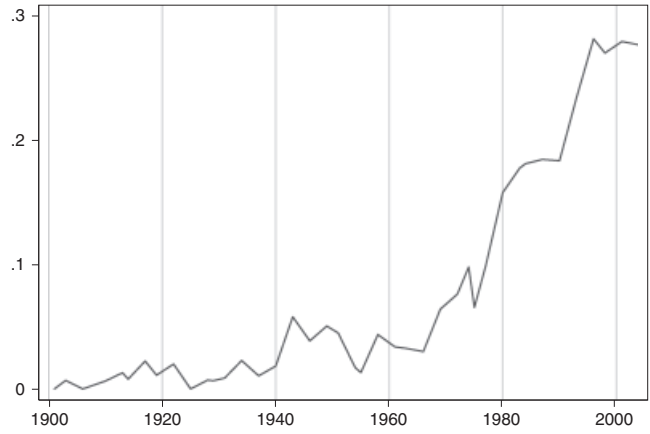

FIGURE 2

Share of Incumbents Who Were Women

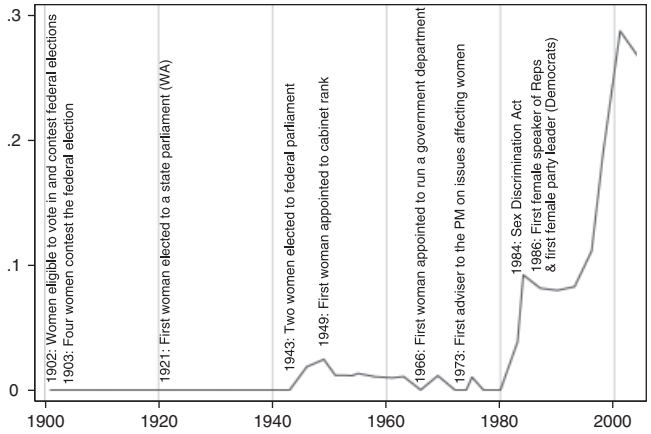

to us by Robert Pugh of the Australian Electoral Commission (AEC). Although these saved us the burden of manually entering the data, we nonetheless spent a considerable amount of time reformatting the data in a consistent manner. Additionally, since the AEC spreadsheets did not always include information on the incumbent candidate, we merged in data on incumbency using records available at Adam Carr's election website.

Since the AEC data do not contain information on the gender of the candidate, we coded candidate gender using their first names. Across the period 1901–2004, 65 candidates did not spell out their first name on the ballot paper; providing only their initials. We exclude these individuals, since we were unable to determine their gender. In all other cases, we coded candidates' names as male or female. Where we were unsure of the gender of a particular candidate, we attempted to consult official sources. This was not always possible, so it is conceivable that we have mis-coded candidate gender in some instances. However, given that there is empirically very little overlap between men's and women's names, this is unlikely to be a significant problem in practice.[3] To the extent that this measurement error is uncorrelated with candidate vote share, this should not bias our results.

[3]To check this, we contacted the Registry of Births Deaths & Marriages in New South Wales (Australia's largest state) and were supplied with a file containing the most popular names in that state (annual birth counts for approximately 400 names per gender for 1952–2008). Within this group, 96 percent of women have names that are extremely feminine (i.e., more than 99 percent of those with the name are female), and 99 percent of women have names that are highly feminine (i.e., more than 90 percent of those with the name are female). The statistics for male names are very similar. This suggests that over this period, names provided an extremely clear signal of an Australian's gender.

Figure 1 shows the share of candidates in each election who were female. This figure first exceeded 5 percent in 1943, and only exceeded 10 percent in 1980. During the 1980s and 1990s, the share of female candidates steadily rose. In the 2004 election, 28 percent of candidates were female, a slightly higher share than in federal parliament.

Figure 2 graphs the share of incumbent candidates who are female (we graph incumbents rather than winners, since it directly pertains to the regression analysis that follows). We also note on the graph some important milestones for women in Australian politics.

A surprising aspect of these data is that prior to 1970, only three women were elected to federal parliament: Enid Lyons, Doris Blackburn, and Kay Brownbill. Perhaps more striking is the fact that despite social changes that took place during the 1970s, only one woman (Joan Child) was elected to parliament in that decade (AEC, 2007). This is consistent with data from the upper echelons of the federal public service (the Senior Executive Service), where the share of women rose more rapidly in the 1980s than in the 1970s (ABS, 1997), but it is not consistent with the pattern in state and territory parliaments, where female representation increased at about the same rate in the 1970s as it did in the 1980s (Sawer, 2001). However, this pattern may simply reflect that barriers to electoral success in state and territory parliaments are lower than at the federal level.

Our formal analysis involves estimating regressions that take the form:

$$Voteshare_{ijkt} = \beta FemaleCandidate_i + \gamma' Z_{ijkt} + \varepsilon_{ijkt}. \quad (1)$$

In this regression, the dependent variable is the share of the first-preference vote received by candidate i, representing party j, in electorate k, and election t. This variable, which we term *Voteshare*, ranges from 0 to 1. Throughout this article, our dependent variable is a candidate's share of the primary vote.[4] The variable *FemaleCandidate* is an indicator variable, taking the value 0 for male candidates, and 1 for female candidates.

Z is a vector of candidate-specific, party-specific, and election-specific characteristics, which we include because they may be correlated with both *Voteshare* and *FemaleCandidate*. The more candidates who run in a given race, the fewer votes each will be likely to garner. To account for this effect, all our regressions control for the reciprocal of the number of candidates standing in that particular electorate in a given year. Because incumbency may confer electoral benefits, all regressions also control for whether the candidate was the incumbent (i.e., whether he or she won that seat at the previous election, or in a by-election). In some instances, politicians win in

[4]Australian elections use a preferential voting system (also termed "instant runoff"). Although we have data on the full preference distribution for most electorates, we focus on the primary vote, since using the two-party preferred vote shares would by definition involve excluding all but the top two candidates in each race. Using the primary vote share as the dependent variable allows us to include minor party candidates in our sample, and significantly increases our sample size.

one electorate and run in a different electorate at the next election. We code these as nonincumbents.

As noted above, a number of studies of women in politics have referred to the "sacrificial lamb hypothesis"; the notion that political parties choose women to run for unwinnable seats, but put forth men as candidates when the party has a reasonable chance of winning the seat (see, e.g., Berch, 2004). To account for the possibility that women may stand in electorates where the ex ante probability of their party winning is lower (or higher), we include a control that we term the "expected vote share." Ideally, this variable should capture a candidate's probability of winning given his or her political party. However, we do not want the expected vote share measure to capture anything about the particular candidate, or we run the risk that it will attenuate the *FemaleCandidate* coefficient. Therefore, the expected vote share takes the value of the vote share received by a different candidate of the same party running in the most recent election. For example, suppose that Candidate A received 10 percent and 15 percent in Elections 1 and 2, and was then replaced by Successor B from the same party. For each election in which B stands, the expected vote share measure would be set to 15 percent. Only elections that are within five elections of one another are coded in this manner. In the 2 percent of cases where it was not possible to create the expected vote share measure in this way, we used one of three alternative techniques (for details, see King and Leigh, 2009b).

Our preferred specification includes a party $\times$ election fixed effect. This effectively allows us to control for party-specific swings in each election. In this specification, the coefficient on *FemaleCandidate* is effectively the difference in the electoral performance of male and female candidates running for the same party in the same election (accounting for incumbency and expected vote share). For completeness, we also show results without party $\times$ election fixed effects, or with separate party and election fixed effects. All specifications are estimated using ordinary least squares, with standard errors clustered at the election $\times$ electorate level, to account for the fact that, within a given race, vote shares must sum to 100 percent.

Our sample covers 40 elections held between 1903 and 2004. It includes 4,213 separate contests, 10,528 candidates, and 16,767 candidate-election observations. Summary statistics and further details on our data are provided in King and Leigh (2009b). Female candidates comprise 14 percent of the sample.

Table 1 presents the results of these specifications. In the absence of election and party fixed effects, the *FemaleCandidate* coefficient is very small (0.001) and statistically insignificant. However, once we add election and party fixed effects (Column 2) or election $\times$ party fixed effects (Column 3), the coefficient rises to around two-thirds of a percentage point, and is significant at the 1 percent level. The fact that the results without fixed effects are insignificant is itself curious, and suggests election-specific and party-specific factors that are correlated with the success of female candidates

TABLE 1

Basic Gender Effects

	1	2	3 Preferred Specification
Female candidate	0.001	−0.006***	−0.006***
	[0.002]	[0.001]	[0.001]
Incumbent	0.128***	0.132***	0.131***
	[0.003]	[0.003]	[0.003]
Expected vote share	0.741***	0.401***	0.405***
	[0.008]	[0.012]	[0.012]
1/Total candidates	0.309***	0.380***	0.363***
	[0.009]	[0.012]	[0.012]
Observations	16767	16767	16767
R^2	0.81	0.86	0.89
Election and party FE	No	Yes	No
Election × Party FE	No	No	Yes

***,**,and* denote statistical significance at the 1 percent, 5 percent, and 10 percent levels, respectively.

NOTE: Standard errors in brackets.

(indeed, the results in Column 2 are largely unchanged if we include only election fixed effects, or only party fixed effects).

In our preferred specification (Column 3 of Table 1), we estimate that being female reduces a candidate's primary vote share by 0.6 percentage points. One can glean some sense of the magnitude of this effect from the fact that the benefit of being male is about two-thirds as large as the benefit of drawing the top spot on the ballot paper (King and Leigh, 2009a). Another potentially useful comparison is to note that in recent Australian elections, around 4 percent of seats have been decided by a margin of 0.6 percent or less.[5] The other covariates have the expected sign and magnitude, with the incumbency coefficient around 0.13, the expected vote share coefficient ranging from 0.74 (without party and election fixed effects) to 0.4 (with fixed effects), and the coefficient on the reciprocal of the number of candidates ranging between 0.3 and 0.4.[6]

The Preselection Problem

Although we find evidence of a bias against female candidates in Australia, it is possible that part (or all) of this bias could be due to preselectors, rather

[5]This comparison uses two-party preferred data from the 1996–2004 elections, in which 24 out of 595 races were decided by a margin of 0.6 percent or less.

[6]Some might wonder why the coefficient on 1/*TotalCandidates* is not 1. The answer is that in a bivariate regression of *Voteshare* on 1/*TotalCandidates*, the coefficient is precisely 1, but adding other controls (incumbency, expected vote share, election effects, and party effects) attenuates the coefficient.

than to voters. As Kelley and McAllister (1983) have suggested, the preselection process used by the major political parties in Australia may be a source of disadvantage for female candidates.[7] Our concern here is not with estimating the relative size of preselector bias and voter bias, but with a more subtle question: To what extent are our estimates of voter bias merely a reflection of preselector bias?

Here, we adopt two strategies that help us to determine whether or not party-based discrimination is driving our results. First, we restrict the sample to incumbents, and compare female incumbents with male incumbents. Incumbents, unlike their challengers, do not typically face a hotly contested preselection process in order to become a party's candidate. In effect, we hypothesize that if the voter bias differs between incumbents and challengers, then party bias is likely to affect estimates of the gender voting bias. This strategy follows Milyo and Schosberg (2000), who found that when the sample was restricted to U.S. incumbents, female incumbents outperformed male incumbents, controlling for challenger quality (see also Berch, 2004). Second, we restrict the sample to independent candidates, and compare female and male candidates who run as independents. By definition, independent candidates do not have to face preselection, so this strategy provides an alternative way of estimating voter bias in a way that is uncontaminated by preselector behavior.

Table 2 shows the results from these two approaches. In both cases, we use our preferred specification (Table 1, Column 3), which includes election $\times$ party fixed effects. We find that female nonincumbents are penalized by 1.8 percentage points at the ballot box (Column 1), while incumbent females receive 11.0 percentage point fewer votes (Column 2).[8] These effects are both larger than those shown in Table 1, since the sample in Columns 1 and 2 is restricted to major party candidates. A formal test cannot reject the hypothesis that the two *FemaleCandidate* coefficients are the same. The results in Columns 3 and 4 bear out this conclusion. Separating the sample into independent candidates and those who ran with the support of a political party, we again find no difference between the two. Since independents do not have to face a preselection, this suggests that our estimates of voter bias are not affected to any large extent by preselector behavior.

Unfortunately, both of the strategies employed in Table 2 have their limitations. While incumbents are often reselected without a contested preselection, their initial selection is often hotly contested. As a result, it is

[7]In King and Leigh (2009b), we explore the issue using the 1987 and 1990 Australian Candidate Studies, and find some suggestive evidence of preselector bias in those elections (though it is quite possible that this estimate does not generalize to the other elections in our sample).

[8]By contrast, Milyo and Schosberg (2000) found that female incumbents in U.S. House elections over the period 1984–1992 received a 6 percentage point higher vote share than male incumbents.

TABLE 2

Two Strategies for Addressing the Preselection Problem

	1	2	3	4
	Nonincumbent Major Party	Incumbent Major Party	Nonindependents	Independent Candidates
Female candidate	−0.018***	−0.010**	−0.006***	−0.006***
	[0.004]	[0.005]	[0.001]	[0.002]
Incumbent			0.129***	0.241***
			[0.003]	[0.023]
Expected vote share	0.513***	0.237***	0.419***	0.137***
	[0.017]	[0.014]	[0.012]	[0.039]
1/Total candidates	0.415***	0.342***	0.357***	0.471***
	[0.021]	[0.022]	[0.013]	[0.040]
Election × Party FE	Yes	Yes	Yes	Yes
Observations	5905	3099	14853	1914
R^2	0.5	0.43	0.88	0.5
P value on test of equality of female coefficients	0.27		0.99	

***, **,and *denote statistical significance at the 1 percent, 5 percent, and 10 percent levels, respectively.

NOTE: Standard errors in brackets.

plausible that male and female major party incumbents may differ on some important dimension. In the case of independent candidates, it is plausible that the typical voter with a propensity to vote independent has a different gender bias than other voters in the electorate. Nonetheless, our results would be unlikely to look as they do in Table 2 if gender differences for candidates were driven solely by preselector bias.

One natural experiment that may serve to shed light on this question is the 1994 decision by the Australian Labor Party to enact an affirmative action policy. This rule required that in the first election following the year 2001, 35 percent of Labor's candidates in winnable seats must be women (Whip, 2003).[9] The change required was quite substantial, given that in 1994, only 10 of Labor's 80 House of Representatives members were women. If we regard this as an exogenous shock to the preselection system, it can help shed light on the extent to which preselector bias affects our estimates of voter bias.

[9]Since 2001, the ruling holds that at least 40 percent of preselected ALP candidates must be women so that, by 2012, at least 40 percent of the seats held by Labor will be filled by women (Australian Labor Party, 2004:7–8).

To test the effect of the rule change, we compare the performance of female candidates standing for the Labor Party before and after the 35 percent rule took effect. In order that our results are identified only from the rule change, we include three sets of fixed effects: election × party fixed effects, which absorb party-specific swings from election to election; female × party fixed effects, which allow voters to provide a differential level of support to female and male candidates from each party; and female × election fixed effects, which allow the average degree of support for female candidates to change from one election to the next. Our results are therefore identified from the triple-difference interaction *FemaleCandidate* × *ALP* × *Post*, which denotes female candidates who ran for the Labor Party after the rule took effect. We present two specifications: one in which the postperiod denotes the elections after 1994 (when the rule change began to affect preselections), and one in which the postperiod denotes the election after 2001 (when the 35 percent rule change was operative).

Table 3 presents results from these specifications. In both cases, the coefficient on the triple-difference term is close to zero and statistically insignificant. This suggests that the new policy—which substantially increased the share of female candidates—did not affect the bias against women for ALP candidates. This provides further support for the notion that the preselection process is not driving our estimates of voter bias. (Again, it is important to stress that we are not measuring preselector bias; instead, we are checking whether preselector bias affects our estimates of voter bias.)

TABLE 3

Does an Exogenous Change in Preselection Affect the Electability of Female Candidates?

	1	2
Female × ALP × Post1994	−0.004	
	[0.008]	
Female × ALP × Post2001		0.000
		[0.011]
Incumbent	0.131***	0.131***
	[0.003]	[0.003]
Expected vote share	0.404***	0.404***
	[0.012]	[0.012]
1/Total candidates	0.362***	0.362***
	[0.012]	[0.012]
Observations	16767	16767
R^2	0.89	0.89
Election × Party FE	Yes	Yes
Female × Party FE	Yes	Yes
Female × Election FE	Yes	Yes

***,**, and *denote statistical significance at the 1 percent, 5 percent, and 10 percent levels, respectively

NOTE: Standard errors in brackets.

What Explains Gender Voting Bias?

In the previous section, we found evidence of a persistent bias against female candidates in Australia over the past century that is not a reflection of discrimination within the party preselection process. Given that the bias appears to lie with the voters themselves, in this section we therefore analyze the factors that might explain differences in voter bias against women between electorates and over time, seeking to tease apart the effect of several different theories. These include changing social norms about the role of women, the share of women sitting in parliament, the share of women standing for office in a given election, and the share of women competing in a given constituency.

To give some sense of the patterns that we are trying to account for, we estimated gender voting bias separately for each election by interacting it with a dummy for that election. Note that such an interaction can either be done by including the main effect (*FemaleCandidate*) and omitting one of the election interactions, or by omitting the main effect and including all election interactions. Here, we opt for the latter. Since our specification includes election $\times$ party fixed effects, it is unnecessary to include election fixed effects as well.

$$\begin{aligned} Voteshare_{ijkt} = {} & \beta_{1901} FemaleCandidate_i \times I_t^{1901} \\ & + \beta_{1903} FemaleCandidate_i \times I_t^{1903} + \cdots \\ & + \beta_{2004} FemaleCandidate_i \times I_t^{2004} + \gamma' Z_{ijkt} + \varepsilon_{ijkt} \end{aligned} \quad (2)$$

Figure 3 plots the beta coefficients from Equation (2), along with their associated standard errors. As can be seen, the degree of gender bias against female candidates has fallen substantially since the early part of the 20th century. We estimate the gender bias against female candidates (of the same party, running in the same election, and controlling for incumbency and expected vote share) to be over 10 percent until the 1920s, and between 5 and 10 percent until the 1940s. In the postwar decades, the penalty fell below 5 percent, where it has remained since. As the bars in Figure 3 show, most female candidates stood in elections since 1980. Between 1980 and 2004, the average penalty against female candidates was one-third of a percentage point, with female candidates outperforming male candidates in two elections (by one-third of a percentage point in 1996, and by one-tenth of a percentage point in 1998).

The shrinking gender penalty in Australia is consistent with trends in other postindustrialist societies around the globe (Paxton and Hughes, 2007). But why has the gender penalty narrowed? Here we test two possible explanations.[10] The first might be termed the "trail-blazing effect," in which

[10]A third theory about the gender penalty—that it is explained by differential voting patterns of male and female voters—is discussed in the working paper version of this article (King and Leigh, 2009b).

FIGURE 3

Ballot Box Penalty Against Women (1903–2004)

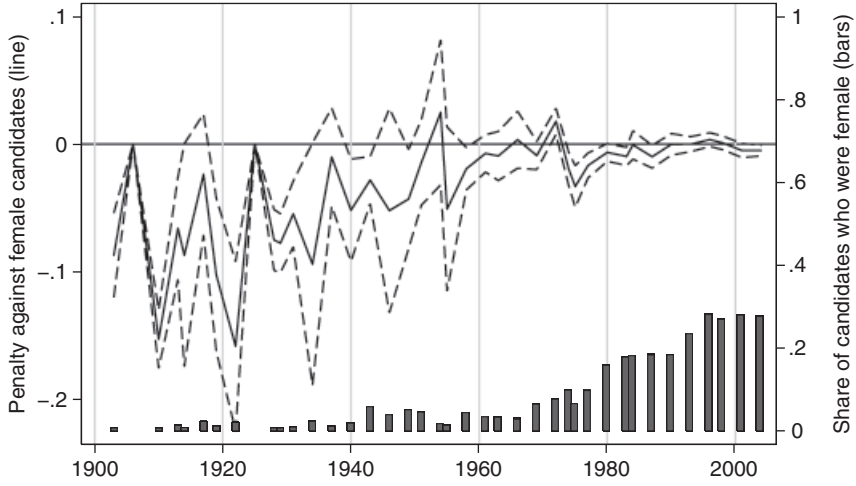

NOTE: Solid line denotes penalty, dashed lines denote 90 percent confidence interval. Bars denote share of candidates who were women.

women perform better in elections when more women stand for election or hold office. This would suggest that at a national level, the gender penalty has been driven by the trends shown in Figures 1 and 2. At a local level, it would also suggest that women candidates benefit if a larger number of their opponents are women.

A second plausible explanation is that the gender penalty is driven by changing social norms (Inglehart and Norris, 2003). As noted above, a series of studies in Western Europe and the United States have found that women are more likely to be elected to political positions in countries or states that hold more egalitarian attitudes toward gender equality. One way of gauging social norms over the entire past century is by looking at labor market differences between men and women. Although earlier studies (Inglehart and Norris, 2003; Paxton and Hughes, 2007) would characterize labor market differences as a structural factor rather than a social or cultural one, we suggest that the gender pay gap can be used as an effective proxy for social attitudes toward gender equality. To test this, we obtain annual data on the gender pay gap from the Australian Bureau of Statistics' average weekly earnings survey (covering the period since 1970) (ABS, various years), and splice this to data from Snooks (1994) (covering 1901–1969). Note that both series relate to full-time nonmanagerial employees, but that the gender gap is not adjusted to account for age, experience, education, or industry. Over the last two decades of our sample, we also have comparable

FIGURE 4

The Ballot Box and Labor Market Penalty Against Women

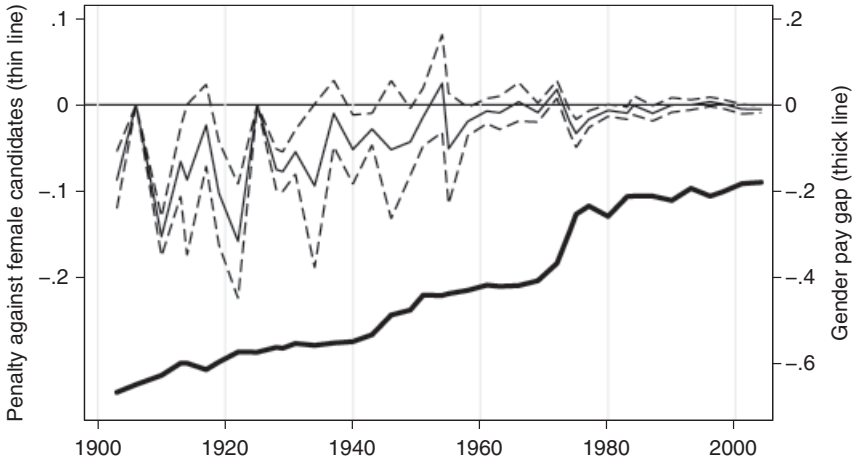

NOTE: Thin line denotes penalty, dashed lines denote 90 percent confidence interval, thick line denotes gender pay gap.

data on attitudes to gender equality, and we find that this closely tracks the gender pay gap.[11] Moreover, when we use our attitudinal variable in place of the gender pay gap in the regression analysis, we obtain similar results.

In Figure 4, we chart the gender pay gap against our estimate of the ballot box penalty against women. In 1903, the gender pay gap was 67 percent (meaning that the typical woman earned 33 percent as much as the typical man). By 1960, the gap had narrowed to around 40 percent. Following the equal pay decisions of 1969, 1972, and 1974, the gap shrunk to around 30 percent, and was 18 percent in 2004. In general, the shrinking of the gender pay gap appears to track the ballot box penalty against women quite well, with both gaps shrinking markedly in the first part of the 20th century, and (at a much slower rate) during the 1980s and 1990s. The only stage at which the two series do not seem to track closely is the late 1960s and early 1970s. This may reflect the fact that the substantial reduction in the gender pay penalty in this period was driven by legislative changes (the equal pay decisions) rather than social attitudes. Since we are using the gender pay gap as a proxy for social attitudes toward women, it may be a less accurate proxy of attitudes in these years.

We now turn to a more formal analysis of these two gender voting hypotheses (the "trail-blazing effect" and the social norms theory), in which

[11]This variable is the share of people in the Australian Election Study who agree with the statement "equal opportunity for women has gone too far." The figure fell steadily from 26 percent in 1987 to 9 percemt in 2004.

TABLE 4

What Affects the Gender Penalty Over Time?

	1	2	3
Female candidate	−0.011***	−0.025***	0.017***
	[0.002]	[0.004]	[0.005]
Female candidate × Share of all incumbents who are female	0.037***		
	[0.012]		
Female candidate × Share of all candidates who are female		0.092***	
		[0.018]	
Female candidate × Gender pay gap			0.098***
			[0.021]
Incumbent	0.131***	0.131***	0.131***
	[0.003]	[0.003]	[0.003]
Expected vote share	0.405***	0.404***	0.404***
	[0.012]	[0.012]	[0.012]
1/Total candidates	0.363***	0.363***	0.362***
	[0.012]	[0.012]	[0.012]
Observations	16767	16767	16767
R^2	0.89	0.89	0.89
Election × Party FE	Yes	Yes	Yes

***,**, and *denote statistical significance at the 1 percent, 5 percent, and 10 percent levels, respectively. Note that the share of incumbents who are female, the share of candidates who are female, and the gender pay gap only vary at the election level, so the main effects of these variables are captured by the election × party fixed effects.

NOTE: Standard errors, clustered at the election × electorate level, in brackets.

we augment Equation (1) by adding an interaction between the possible explanator and an indicator variable denoting whether the candidate is female. We begin by focusing on the three variables that change over time: the share of incumbents who are female, the share of candidates who are female, and the gender earnings penalty.[12] Since the specifications already include election × party fixed effects, these capture the main effects of the three variables, so we need only include the interactions.

Table 4 shows the results from this specification. In general, the coefficients accord with expectations. In line with the shifting social norms theory (Inglehart and Norris, 2003), female candidates tend to do better if more women are in parliament, if more women are running in that election, and if the gender pay gap is smaller.[13] The interactions suggest that a 10 percent-

[12]Since these variables are measured at the national level, we do not exclude the individual from the calculation. Doing so makes no substantive difference to the results, but does slightly complicate the presentation of results, since we need to show the main effects as well.

[13]However, if we include these three factors together, the female incumbent interaction flips sign. This suggests that—controlling for the gender pay gap and the share of candidates who are female—having more women in parliament has a negative effect on the electoral performance of other women.

age point increase in the share of female incumbents would shrink the gender penalty by 0.4 percentage points, while a 10 percentage point increase in the share of candidates who are female would shrink the gender penalty by 0.9 percentage points, and a 10 percentage point reduction in the gender pay gap would shrink the gender penalty by nearly a full percentage point.

The three variables in Table 4 only vary from election to election; however, it is also possible to exploit variation across electorates *in the same election*. Two of our explanatory factors—the share of candidates in a particular electorate who are female (excluding the individual), and the share of voters in an electorate who are female—vary within elections. This allows us to estimate models in which we include both election $\times$ party fixed effects, and election $\times$ female fixed effects. We can then ask the question: Controlling for the average gender penalty in a particular election, why are some electorates more favorable toward women than others?

When we interact the female coefficient with the share of other candidates in the same race who are female, the coefficient is -0.018, which is significant at the 1 percent level (for results, see King and Leigh, 2009b). Since the typical race has four candidates, and the share variables exclude the individual, this suggests that a woman running against three men receives 0.6 percent more of the vote (0.018×0.33) than a woman running against two men and another woman.

We also explore whether female candidates do better in electorates with a larger share of female voters (restricting the sample to 1903–1966). Since we can observe only the aggregate share, this specification potentially suffers from the ecological fallacy problem. If the gender composition of an electorate has an independent effect on voting patterns, or if it is correlated with something else about the electorate, then the relationship between the female share of voters and the performance of female candidates may not reflect how female voters cast their ballots.

With this caveat in mind, we find only a weak relationship between the share of voters who are female and the performance of female candidates. The coefficient is substantively large (suggesting that a 10 percent increase in the share of female voters leads to a 3 percentage point drop in the vote share of female candidates), but is significant only at the 10 percent level. Including both the candidate composition interaction and the voter composition interaction has little effect on the point estimates (for results, see King and Leigh, 2009b), but neither is statistically significant. We therefore conclude that the gender composition of other candidates matters (using the full sample), but that the gender composition of voters does not have a significant effect.

Conclusion

Despite Australia's history as the first country to grant women the right to vote in and stand for national elections, and the fact that women make up a

majority of the Australian population (50.3 percent), women are still substantially underrepresented in the Australian parliament. Using data on elections since 1903, we show that, in contrast to previous U.S. studies, this is partly due to a systematic penalty against female candidates at the ballot box. Our regression analysis takes account of party affiliation, incumbency, expected vote share, and the number of candidates running in that election. We find that in Australia, female candidates faced a penalty at the ballot box of at least 5 percentage points until World War II, a couple of percentage points in the immediate postwar decades, and less than 1 percentage point in the 1990s and early 2000s. On average, female candidates received 0.6 percentage points fewer votes than male candidates. The effect was larger for female candidates representing major parties, who received 1–2 percentage points less, but is still smaller than previous Australian estimates (Kelley and McAllister, 1983; Bean and McAllister, 1990). We find no evidence of any consistent benefit to female candidates at the Australian ballot box.

In theory, differences in the electoral performance of male and female candidates could be explained solely by biases in the preselection system. However, unlike previous studies (Darcy and Schramm, 1977; Kelley and McAllister, 1984), we do not find evidence that the preselection system has had much of an effect on the ballot box penalty against female candidates. Three pieces of evidence support this. First, the gender penalty against women from major parties is similar among incumbents and challengers. Second, independents (who do not face preselection) receive a gender penalty similar to that received by nonindependents. And third, a substantial increase in the share of women preselected by the Australian Labor Party did not appear to affect the electoral performance of female Labor candidates. The implications of this finding are important. Although policies such as the ALP's gender quotas are a meaningful way of improving female candidates' chances within certain political parties, some have argued that affirmative action policies inevitably lead to a backlash against female candidates by voters. Empirically, we find no evidence of such a backlash.

Given that much of the bias against women candidates lies with the voters themselves, what explains changes in the gender penalty over time? Of the three explanations tested here, we find evidence to support arguments by Inglehart and Norris (2003) and Paxton and Hughes (2007) that changing social norms (as proxied by the gender pay gap) have an effect on the electoral gender penalty. Australia can expect to see more women in Parliament as the gender pay gap narrows. We also find that a higher share of female candidates running nationwide in a given election boosts the chances of a given female candidate winning. Conversely, however, we find that female candidates are harmed, not helped, by having more women on the *same ballot paper in the same election*. Female candidates fare worse when they have to compete against other female candidates. Finally, we show that

the share of voters who are women does not appear to have a positive impact on the performance of female candidates. Although our data are at an aggregate level, they are consistent with the theory that female voters are not more likely to support female candidates.

One should always be cautious about generalizing results across countries, but there are several reasons to think that our findings may have relevance to researchers and policymakers in other settings. In terms of gender equality, the gender pay gap in Australia is on par with the gap in other industrialized nations (Blau and Kahn, 1996, 2003). From a political perspective, Australian elections have some unusual features (e.g., compulsory voting and instant runoff vote counting), but its House of Representatives elections (based on single-member electorates, which are effectively two-party contests in most cases) have much in common with elections in other parts of the world. Nonetheless, it would be interesting to see the extent to which our findings on the voting bias against female candidates across the 20th century are common to other countries.

REFERENCES

Australian Bureau of Statistics (ABS). 1997. *Australian Women's Year Book 1997*. Cat. No. 4124.0. Canberra: ABS. Available at ⟨www.abs.gov.au⟩.

———. Various years. *Average Weekly Earnings, Australia*. Cat. No. 6302.0 (August). Canberra: ABS.

Australian Electoral Commission (AEC). 2006. *Compulsory Voting in Australia*. Available at ⟨www.aec.gov/au/pdf/voting/compulsory_voting.pdf⟩.

———. 2007. *Electoral Milestones for Women*. Available at ⟨www.aec.gov.au/_content/when/history/milestone.htm⟩.

Australian Labor Party. 2004. *National Constitution of the ALP*. (As amended at the 43rd national ALP conference 2004.) Available at ⟨www.alp.org.au⟩.

Bean, Clive, and Ian McAllister. 1990. "The Electoral Performance of Parliamentary Candidates." Pp. 73–91 in Clive Bean, Ian McAllister, and J. Warhust, eds., *The Greening of Australian Politics: The 1990 Federal Election*. Melbourne: LongmanCheshire.

Berch, Neil. 2004. "Women Incumbents, Elite Bias, and Voter Response in the 1996 and 1998 U.S. House Elections." *Women and Politics* 26(1):21–32.

Blau, Francine D., and Lawrence M. Kahn. 1996. "Wage Structure and Gender Earnings Differentials: An International Comparison." *Economica* 63(250):S29–S62.

———. 2003. "Understanding International Differences in the Gender Pay Gap." *Journal of Labor Economics* 21(1):106–44.

Caul, Miki. 2001. "Political Parties and the Adoption of Candidate Gender Quotas: A Cross-National Analysis." *Journal of Politics* 63(4):1214–29.

Conway, Margaret M. 2001. "Women and Political Participation." *PS: Political Science and Politics* 34(2):231–33.

Darcy, Robert E., and Sarah S. Schramm. 1977. "When Women Run Against Men." *Public Opinion Quarterly* 41(1):1–12.

Ekstrand, Laurie E., and William A. Eckert. 1981. "The Impact of Candidate's Sex on Voter Choice." *Western Political Quarterly* 34(1):78–87.

Fox, Richard L., and Zoe M. Oxley. 2003. "Gender Stereotyping in State Executive Elections: Candidate Selection and Success." *Journal of Politics* 65(3):833–50.

Fox, Richard L., and Eric Ran Smith. 1998. "The Role of Candidate Sex in Voter Decision-Making." *Political Psychology* 19(2):405–19.

Herrnson, Paul. S., J. Celeste Lay, and Atiya Kai Stokes. 2003. "Women Running 'as Women': Candidate Gender, Campaign Issues, and Voter-Targeting Strategies." *Journal of Politics* 65(1):244–55.

Huddy, Leonie, and Nayda Terkildsen. 1993a. "The Consequences of Gender Stereotypes for Women Candidates at Different Levels and Types of Office." *Political Research Quarterly* 46(3):503–25.

———. 1993b. "Gender Stereotypes and the Perception of Male and Female Candidates." *American Journal of Political Science* 37(1):119–47.

Inglehart, Ronald, and Pippa Norris. 2003. *Rising Tide: Gender Equality and Cultural Change Around the World*. Cambridge: Cambridge University Press.

Inter-Parliamentary Union (IPU). 2007. *Women in National Parliaments: Situation as of 31 October 2007*. Available at ⟨www.ipu.org⟩.

Kelley, Jonathan, and Ian McAllister. 1983. "The Electoral Consequences of Gender in Australia." *British Journal of Political Science* 13(3):365–77.

———. 1984. "Ballot Paper Cues and the Vote in Australia and Britain: Alphabetic Voting, Sex, and Title." *Public Opinion Quarterly* 48(2):452–66.

Kenworthy, Lane, and Melissa Malami. 1999. "Gender Inequality in Political Representation: A Worldwide Comparative Analysis." *Social Forces* 78(1):235–68.

King, Amy, and Andrew Leigh. 2009a. "Are Ballot Order Effects Heterogeneous?" *Social Science Quarterly* 90(1):71–87.

———. 2009b. *Bias at the Ballot Box*. CEPR Discussion Paper 625. Canberra: Australian National University.

King, David C., and Richard E. Matland. 2003. "Sex and the Grand Old Party: An Experimental Investigation on the Effect of Candidate Sex on Support for a Republican Candidate." *American Politics Research* 31(6):595–612.

Kittilson, Miki Caul. 2006. *Challenging Parties, Changing Parliaments. Women and Elected Office in Contemporary Western Europe*. Columbus, OH: Ohio State University Press.

Mackerras, Malcolm. 1977. "Do Women Candidates Lose Votes?" *Australian Quarterly* 49: (September): 6–10.

Milyo, Jeffrey, and Samantha Schosberg. 2000. "Gender Bias and Selection Bias in House Elections." *Public Choice* 105: 41–59.

Paxton, Pamela, and Melanie M. Hughes. 2007. *Women, Politics and Power: A Global Perspective*. Los Angeles, CA: Pine Forge Press.

Sawer, Marian. 2001. "Women and Government in Australia." In *Year Book Australia, 2001*. Cat. No. 1301.0. Canberra: ABS.

Snooks, Graeme Donald. 1994. *Portrait of the Family Within the Total Economy. A Study in Longrun Dynamics, Australia 1788–1990*. Cambridge: Cambridge University Press.

Whip, Rosemary. 2003. "The 1996 Australian Federal Election and its Aftermath: A Case for Equal Gender Representation." *Australian Feminist Studies* 18(40):73–97.