9. The gender gap in workplace authority

In this chapter we will explore the intersection of gender inequality and one specific dimension of class relations – the authority structure within workplaces. No one, of course, would be surprised by the general fact that workplace authority is unequally distributed between men and women in all of the countries we examine. What might be surprising to most people, as we shall see, is the specific pattern of cross-national variation in the gender gap in authority. To cite just one example, in the United States the probability of a man in the labor force occupying an "upper" or "top" management position is 1.8 times greater than the probability of a woman occupying such a position, whereas in Sweden, the probability for men is 4.2 times greater than for women. The objective of this chapter is to document and to attempt to explain these kinds of cross-national variations in gender inequality in workplace authority in seven developed, capitalist countries – the United States, Canada, the United Kingdom, Australia, Sweden, Norway and Japan. In doing so we are particularly interested in revealing the extent to which these patterns reflect variations in gender discrimination in various forms.

9.1 Analytical strategy for studying the "gender gap"

The ideal data for analyzing gender discrimination in access to authority would include direct observations of the discriminatory acts that cumulatively shape the outcomes. Since such data are never available in systematic, quantifiable form, research on gender inequalities in labor market outcomes typically relies on indirect methods of assessing discrimination. We will adopt a strategy which can be called the "net gender gap" approach. The basic idea is this. We begin by measuring the "gross gender gap" in authority in a country. This is simply a measure of

the relative probabilities of a woman compared to a man having a particular kind of authority. We then examine what happens to these relative probabilities when we control for a variety of attributes of men and women (such as education or job experience). The relative probabilities of women compared to men having authority when these controls are included in the analysis will be called the "net gender gap" in authority. We will treat the magnitude of this net gender gap as an indicator of the degree of direct discrimination in the allocation of authority. In a sense, discrimination is being treated as the "residual explanation" when other nondiscrimination explanations (represented by the control variables in the equation) fail to fully account for gender differences in authority. Of course, even if the net gender gap were zero, this would not prove that discrimination is absent from the social processes generating overall gender differences in authority, since discrimination could systematically affect the control variables themselves. The net gender gap strategy, therefore, is effective only in assessing the extent to which discrimination operates *directly* in the process of allocating authority within organizations.

The net gender gap strategy of analysis is always vulnerable, either because of possible misspecifications of the equation (important nondiscrimination causes of the gender gap might be excluded from the analysis) or because of poor measurement of some of the variables. What looks like a residual "discrimination" gap, therefore, may simply reflect limitations in the data analysis. Nevertheless, if the gender gap in authority remains large after controlling for a variety of plausible factors, then this adds credibility to the claim that direct discrimination exists in the process by which authority is allocated.

The basic statistical device we will use to measure the extent of the gender gap in authority is derived from "odds ratios." We have already encountered these in the analysis of permeability of class boundaries in chapter 5. In that earlier chapter the issue was odds of a person from a particular class location having certain kinds of social ties across particular class boundaries. Here the issue is the odds of women compared to men having particular kinds of authority. The "gender gap coefficient" we will use is, technically, 1 minus the odds ratio of a woman compared to a man having authority. If the odds of having authority for women and men are equal (and thus the ratio of their respective odds is 1), we will say that the gender gap in authority is zero. If no women at all have authority, and thus the odds of a woman having authority is zero, the gender gap will be 1. If it should happen that the odds of women having

authority were greater than those of men, the gender gap will be negative.[1]

9.2 Empirical agenda

The data analysis in this chapter revolves around three main tasks: analyzing the net gender gap in authority *within* countries; examining whether the gender gaps in authority within countries take the form of a "glass ceiling"; and, exploring a variety of possible explanations of the cross-national variations in net gender gaps.

Authority variables

The analyses reported in this chapter will mainly revolve around a dichotomous measure of authority referred to as *overall authority dichotomy*. This variable is itself derived from three more specific measures of authority: *sanctioning authority* (the ability to impose positive or negative sanctions on subordinates); *decision-making authority* (direct participation in policy making decisions within the employing organization); and *Formal Position in the authority hierarchy* (occupying a job which is called a managerial or supervisory position in the official hierarchy of an organization). If a person has at least two of these three kinds of authority, then they will have authority on the *overall authority dichotomy*. (For details of the construction of these variables, see Wright 1997:

[1] The technical way of generating the coefficient for the *gross* gender gap is to first calculate, for each country, a logistic regressions in which gender is the only independent variable:

$$\text{Log}\,[\Pr(A=1)/\Pr(A=0)] = a + B_1\text{Female},$$

where $\Pr(A=1)$ is the probability of a person having authority as defined by our various measures, $\Pr(A=0)$ is the probability of a person not having authority, and Female is a dummy variable. The significance level of coefficient B_1 in this model is a test of whether men and women differ significantly in their chances of having managerial authority. Taking the antilog of this coefficient yields the odds ratio of women compared to men having authority. The gender gap is then calculated as 1 minus the antilog of B_1. To evaluate the *net* gender gap, we add the compositional control variables to this equation:

$$\text{Log}\,[\Pr(A=1)/\Pr(A=0)] = a + B_1\text{Female} + \Sigma_i B_i X_i$$

where the X_i are the firm attribute, job attribute and person attribute compositional variables. This enables us to test whether the bivariate relationship between gender and authority reflects other factors that are correlated with gender and managerial authority. See Wright (1997: 362–363) for definitions of these control variables.

361–367). I also analyzed all of the patterns using a more complex 10-point authority scale. None of the results were substantively different using this variable and thus I will only report the results for the simpler authority dichotomy.

Analyzing the net gender gap in having authority within countries

The core idea of the "net gender gap" approach is to specify plausible explanations of gender differences in authority that do *not* involve direct discrimination in promotions and then to see if the authority gap disappears when these nondiscrimination factors are held constant in an equation predicting authority. We will explore two explanations of this sort of the gender gap in authority: (1) the gender gap is due to gender differences in various personal attributes of men and women and their employment settings; (2) the gender gap is due to the self-selection of women.

1. Compositional factors

We will explore three clusters of compositional factors: *firm attributes* (economic sector, state employment, firm size); *job attributes* (occupation, part-time employment, job tenure); and *personal attributes* (age, education, labor force interruptions). To the extent that women are concentrated in sectors with a lower proportion of managers, or have various job and personal attributes associated with low probabilities of managerial promotions, then once we control for these factors, the authority gap between men and women should be reduced and perhaps even disappear.

It could be objected that some of these compositional factors are in part *consequences* of discrimination in promotions rather than indirect *causes* of the gender gap, and therefore should not be included in the exercise. It could be the case, for example, that one of the reasons women are more likely to work part time is precisely because they are excluded from promotions to managerial positions. Exclusion from positions of authority could thus explain some of these compositional factors rather than vice versa. We have no way in the present data analysis to investigate this possibility. Nevertheless, if the inclusion of these diverse controls does *not* significantly reduce the gender gap in authority, this would add considerable weight to the claim that the gap is to a significant extent the result of direct discrimination in the allocation of authority positions.

2. Self-selection because of family responsibilities

For various reasons, it might be argued, women in similar employment situations and with similar personal attributes to men may simply not want to be promoted into positions of authority as frequently as men, particularly because of family responsibilities. Given the array of feasible alternatives, women may actually prefer the "mommy track" within a career because of the reduced pressures and time commitment this entails even though it also results in lowered career prospects, especially for vertical promotion. Again, this is not to deny that such preferences may themselves reflect the operation of oppressive gender practices in the society. The gender division of labor in the household or the absence of affordable high-quality childcare, for example, may serve to block the options women feel they realistically can choose in the workplace. Nevertheless, self-selection of this sort is a very different mechanism from direct discrimination by managers and employers in promotion practices.

The most often-cited form of gender self-selection centers around the choices women make with respect to family responsibilities and work responsibilities. We can therefore treat the presence of such responsibilities as additional "compositional factors." However, unlike in the simple compositional arguments which are based on *additive* models of compositional effects, the arguments for self-selection require an *interactive* model. For example, the self-selection hypothesis claims that the presence of children in the household leads women to select themselves out of competition for authority promotions whereas it does not for men. This means that in a model predicting authority, the coefficient for a variable measuring the presence of children would be negative for women but zero, or perhaps even positive, for men, if the presence of children increases the incentives for men to seek promotions because of increasing financial needs of the family. To assess the presence of such self-selection, therefore, we have to estimate a model that includes gender-interactions with the self-selection variables (as well as the additive compositional effects), and then assess the gender gap in authority at appropriate values for the interacting independent variables. For this purpose, we include three variables which are plausibly linked to self-selection: marital status, the presence of children in the household and the percentage of housework performed by the husband.

The glass-ceiling hypothesis

One of the most striking metaphors linked to the efforts of women to gain equality with men in the workplace is the "glass ceiling." The image is that, while women may have gained entry through the front door of managerial hierarchies, at some point they hit an invisible barrier which blocks their further ascent up the managerial highrise. In one of the earliest studies of the problem, Morrison et al. (1987: 13) define the glass ceiling as "a transparent barrier that kept women from rising above a certain level in corporations . . . it applies to women as a group who are kept from advancing higher *because they are women.*"

The glass-ceiling metaphor therefore suggests not simply that women face disadvantages and discrimination within work settings and managerial hierarchies, but that these disadvantages relative to men *increase* as women move up the hierarchy. Employers and top managers may be willing to let women become supervisors, perhaps even lower- to middle-level managers, but – the story goes – they are very reluctant to let women assume positions of "real" power and thus women are blocked from promotions to the upper levels of management in corporations and other work organizations. This may be due to sexist ideas or more subtle discriminatory practices, but, in any case, the glass-ceiling hypothesis argues that the disadvantages women face relative to men in getting jobs and promotions are greater in the upper levels of managerial hierarchies than at the bottom.

Casual observation seems to confirm this argument. There is, after all, a much higher proportion of bottom supervisors than of chief executive officers who are women. In the class analysis project data, at the bottom of managerial hierarchies perhaps 20–25% of lower level supervisors are women in the United States. In contrast, at most a few percent of top executives and CEOs in large corporations are women. According to Fierman (1990) fewer than 0.5% of the 4,012 highest-paid managers in top companies were women, while fewer than 5% of senior management in the Fortune 500 corporations were women and minorities. Reviewing the data on what they call the "promotion gap," Reskin and Padavic (1994: 84) report that "although women held half of all federal government jobs in 1992 and made up 86 percent of the government's clerical workers, they were only a quarter of supervisors and only a tenth of senior executives." Reskin and Padavic report similar findings for other countries: in Denmark women were 14.5% of all managers and administrators, but only between 1 and 5% of top managers; in Japan women

were 7.5% of all administrators and managers but only 0.3% of top management in the private sector. It is hardly surprising with such distributions that it is commonly believed by those working for gender equality that a glass ceiling exists in the American workplace.

However, things may not be what they seem. A simple arithmetic example will demonstrate the point. Suppose, there is a managerial hierarchy with six levels in which 50% of men but only 25% of women get promoted at each level to the next higher level (i.e. men have twice the probability of being promoted than women at every level of the hierarchy). In this situation, if roughly 25% of line supervisors are women, only 1% of top managers will be women. In spite of initial appearances, this example does not fit the story of the "glass ceiling." According to the glass-ceiling hypothesis, the obstacles to women getting managerial positions are supposed to increase as they move up the hierarchy. This could either take the form of a dramatic step function – at some level recruitment and promotion chances for women relative to men plummet to near zero – or it could be a gradual deterioration of the chances of women relative to men. In the example just reviewed, the disadvantages women face relative to men are constant as they move up the hierarchy. And yet, there are almost no women top managers but plenty of women bottom-level supervisors.

What this example illustrates is that the existence of a glass ceiling cannot be inferred simply from the sheer fact that there are many fewer people at the top echelons of organizations who are women than at the bottom levels. The cumulative effect of constant or even declining discrimination can still produce an *increasing* "gender gap in authority" as you move to the top of organizational hierarchies.

The Comparative Class Analysis Project data do not allow us to conduct a fine-grained test of the glass-ceiling hypothesis. Nevertheless, we will make a first cut at the problem by examining the gender gap in authority separately for those people who have made it into the authority hierarchy. If we find that the gender gap in amount of authority for people in the hierarchy is the same or smaller than for the sample as a whole, then this undermines the glass-ceiling hypothesis that gender discrimination is weaker at the port of entry into the hierarchy than in promotions within it. Of course, the glass ceiling could take the form of intensified discrimination only at the very apex of organizations. If this were the case, then we will not be able to observe a glass ceiling in our analysis because of limitations of sample size. If, however, the glass ceiling takes the form of gradually increasing discrimination with higher

levels of authority, then the gender gap in how much authority people have conditional upon them having any authority should be greater than the gender gap in simply having authority.

Explaining cross-national variations

We will pursue two different strategies for exploring possible explanations for the cross-national variations in the gender gap in authority. First, we will compare the differences across countries in the *gross* gender gaps in authority (i.e. the country-specific gender gaps not controlling for any compositional effects) with the differences across countries *net* of the various compositional factors. If a significant portion of the gender gap *within* countries is explained by such compositional factors, then these factors may also account for much of the difference across countries in the gender gap.

Second, if significant differences across countries in the gender authority gap remain after controlling for all of the compositional factors, we then examine in a somewhat less formal way a number of possible macro-social explanations by comparing the rank-ordering of the seven countries on the net gender gap in authority with the rank-ordering on the following variables (see Wright: 1997 351–359 for a discussion of the measures used in these analyses):

1. Gender ideology

All things being equal one would expect a smaller gender gap in workplace authority in societies with relatively egalitarian gender ideologies compared to societies with less egalitarian ideologies.

2. Women's reproductive and sexual rights

Developed capitalist societies differ in the array of rights backed by the state in support of gender equality with respect to sexual and reproductive issues, such as rights to abortion, rights to paid pregnancy and maternity leaves from work, and laws concerning sexual violence, abuse and harassment. While such state-backed rights and provisions do not directly prevent discriminatory practices in promotions, they may contribute to the cultural climate in ways that indirectly affect the degree of inequality in promotions and thus in workplace authority. It would therefore be predicted that societies with strong provisions of these rights would have a smaller gender gap in authority than societies with a weaker support for these rights.

3. Gender earnings gap

It might be expected that in societies in which there was a relatively small gender gap in earnings, the gender gap in workplace authority would also be relatively small. The argument is not that greater equality in earnings capacity between men and women is a *cause* of a smaller authority gap (if anything, a smaller gender gap in authority could itself contribute to narrowing the gender earnings gap), but rather that a society that fosters low levels of income inequality between men and women is also likely to foster low levels of authority inequality as well. Low gender differences in earnings would therefore be taken as an indicator of an underlying institutional commitment to gender equality as such.

4. Occupational sex segregation

The logical relationship between occupational sex segregation and gender inequalities in workplace authority is complex. Clearly, the probability of acquiring authority varies from occupation to occupation, and thus occupational sex segregation can reasonably be viewed as one likely cause of inequalities in authority. However, if norms against women supervising men are strong, then, in a limited way, occupational sex segregation might actually open up managerial positions for women in so far as it increases the chances of women being able to supervise only women. Furthermore, promotions into positions of authority often entail changes in occupational titles. This is particularly true for occupations that are formally called "managerial occupations." Barriers to acquiring workplace authority for women, therefore, are also likely to be a cause of occupational sex segregation. In examining variations across countries in occupational sex-segregation, I am thus not suggesting that this variation is itself a direct cause of variation in the net gender gap in authority. Rather, as in the case of the earnings gap, we will treat occupational sex-segregation as an indicator of underlying processes that shape gender inequalities in the society. It would be expected that countries with relatively high levels of occupational sex segregation would also have large gender gaps in authority.

5. The proportion of the labor force with authority

There are two reasons for expecting the gender gap in authority to be greater in countries in which a relatively small proportion of the labor force held positions of authority than in countries in which there are many authority positions. First, it is more difficult for employers and top executives adequately to fill the positions with men in countries in which

a high proportion of the employees of organizations have authority. In simple supply and demand terms, therefore, employers have an incentive to fill a higher proportion of authority positions with women in a country with a large proportion of managerial and supervisory positions in the job structure. Second, if as some scholars argue (e.g. Reskin 1988; Acker 1990; Bergman 1986), the gender authority gap is at least partially the result of the interests of men in maintaining male predominance in the authority hierarchy, then the incentive for them to try to do so would be stronger when there were relatively few such positions to go around. A proportionately large managerial structure, therefore increases the incentives for the heads of organizations to recruit women into managerial positions and it reduces the incentives for male managers to engage in restrictive practices to protect their positions.

6. The organized women's movement and political culture

If sex discrimination plays a significant role in the exclusion of women from positions of responsibility and power within work, then it would be expected that one of the determinants of the erosion of such sexist practices would be the extent and forms of women's organized challenge to these practices. Two issues in this respect would seem especially important. First, the *overall strength* of the women's movement is crucial for its ability to challenge the gender gap in workplace authority. Second, and perhaps less obviously, the specific *ideological orientation* of the women's movement may shape the extent to which it directs its energies towards problems of workplace discrimination. In particular, it may matter in the extent to which a women's movement is oriented towards equal *rights* or to the provision of services which benefit women.

9.3 Results

The gross gender grap in authority

Figure 9.1 presents in graphic form the gender authority gap coefficients for the overall authority dichotomy variable, both without any control variables and with the compositional controls used to evaluate the net gender gap. Two results are especially striking about the gross gender gap results.

First, in every country, there is a significant gender gap in authority. In results not reported here (see Wright 1997: 338), this gender gap was also significant for each of the three underlying measures of authority

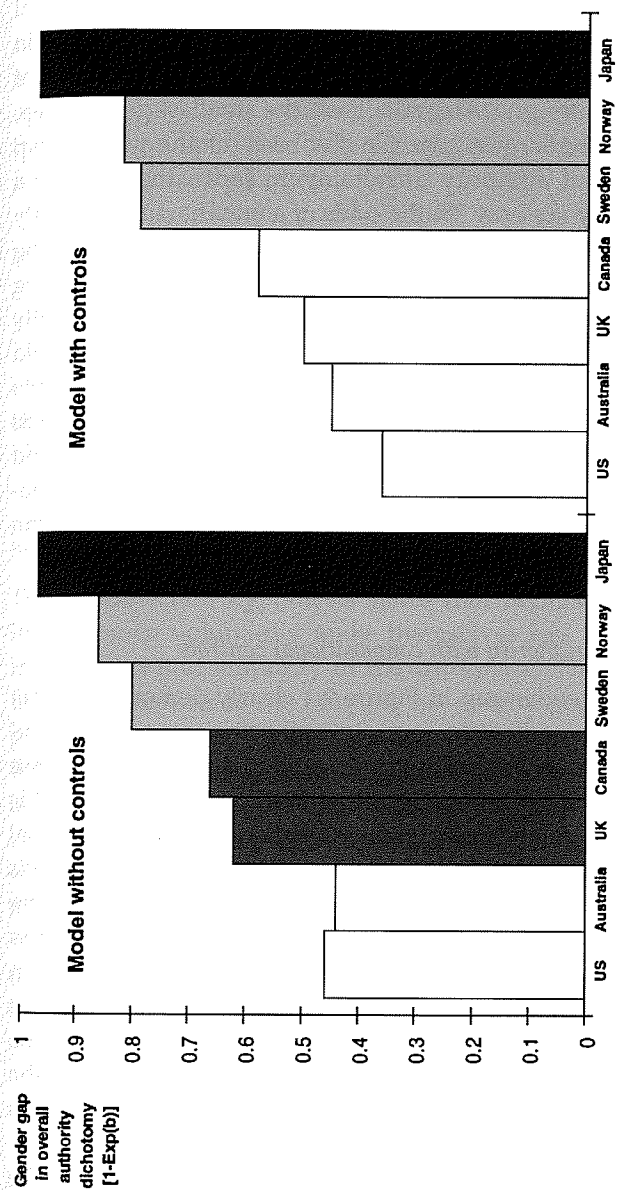

Figure 9.1 *The gender gap in authority in seven countries.*

used to construct the overall authority dichotomy. Women are less likely than men to be in the formal authority hierarchy, to have sanctioning power over subordinates and to participate in organizational policy decisions.

Second, there are statistically significant cross-national variations in the degree of gender inequality in authority. On all of the measures of authority, the United States and Australia have the smallest gender gap, and Japan has by far the largest gap. On the basis of the gender gap coefficients in the overall authority dichotomy, in Japan the odds of a woman having authority are only 3% the odds of a man having authority whereas the odds of a woman in the United States and Australia having authority are around 55% that of a man. The other two English-speaking countries – Canada and the United Kingdom – tend to have significantly greater gross gender authority gaps than does the United States and Australia, but – perhaps surprisingly – smaller gaps than the two Scandinavian countries, Sweden and Norway. While in many respects the Scandinavian countries are among the most egalitarian in the world both in terms of class and gender relations, with respect to the distribution of authority in the workplace, they are clearly less egalitarian than the four English-speaking countries in our analysis.

Net gender differences in authority with compositional controls

The results for the net gender gap in Figure 9.1 clearly demonstrate that relatively little of the overall differences in authority among men and women in any country can be attributed to gender differences in these control variables. One way of assessing this is to ask: by what percent is the gross gender gap in authority reduced when the compositional controls are added to the equation? The biggest compositional effects seems to be in the United States and the United Kingdom, where roughly 20% of the total gender authority gap is closed when the controls are added. In the other countries, the figures range from less than 1% in Japan to 12% in Canada. In both the US and the UK, virtually all of this modest reduction in the gender gap in authority comes from the two job attribute variables (occupation and full-time employment); the inclusion of the personal attribute variables in the equation has almost no effect on the authority gap.

The net gender gap results in Figure 9.1 also show that while the significance level of some of the cross national differences declines in the equations controlling for compositional effects, the basic patterns of the

results are essentially the same as in the equations for the gross gender gap. In particular, the only change between the model without compositional effects and the model with compositional effects is that in the latter the gender coefficients among the four English-speaking countries no longer differ significantly. For the net gender gap, therefore, we have a very clear grouping of our seven countries: the four English-speaking countries have the smallest net gender gaps in authority, the two Scandinavian countries have significantly larger net gender gaps, and Japan has by far the largest.

While it is always possible that we have omitted some crucial compositional variable from the analysis which might affect the results, nevertheless, these results are strongly supportive of the claim that gender differences in authority, and cross-national patterns of such differences, are not primarily the result of differences in the distributions of relevant attributes of men and women and their employment situations. This adds credibility to the claim that direct discrimination or self-selection in the promotion process itself are likely to be important.

Self-selection models

The self-selection hypothesis states that because of family responsibilities, women voluntarily make themselves less available for promotion into positions of authority in the workplace. The way we will examine this hypothesis is to see how the net gender gap in authority varies for people in different family situations. If selection is a significant factor, then we would expect that the gender gap in authority would be greater among married people than among single people, and greater still among married people with children under sixteen. The gender gap would be especially large among married people with children in which the husbands do very little housework.

The technical strategy for seeing how the gender gap varies across family situation is to examine equations which include interaction terms between gender and the different measures of family situation.[2]

[2] The logistic regression used to estimate these interactions is:

$$\text{Log}\,[\Pr(A=1)/\Pr(A=0)] = a + B_1\text{Female} + B_2\text{Married} + B_3\text{kids} + B_4\text{Husband's Housework} + B_5\text{Female} \times \text{Married} + B_6\text{Female} \times \text{kids} + B_7\text{Female} \times \text{Husband's Housework} + \Sigma_i B_i X_i$$

where the X_i are the compositional control variables used in the analysis of the net gender gap. If self-selection is a powerful force in shaping the gender gap in authority, then at least some of the interactive terms in these equations – B_4, B_5, B_6, B_7 – should be statistically significant.

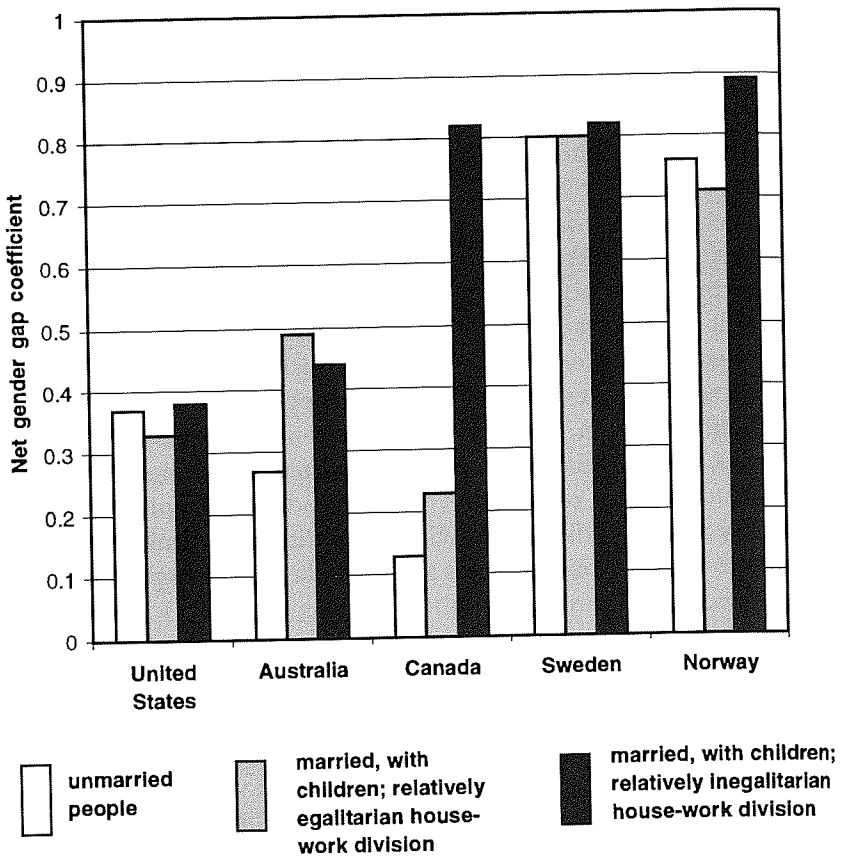

Figure 9.2 *Testing the effects of "self-selection" on the gender gap in authority.*

Figure 9.2 presents the results of our estimates of the gender gap in authority for people in three types of family situations: unmarried people; married people without children and with a relatively egalitarian distribution of housework; married people with children and with an inegalitarian distribution of housework (Japan and the UK are not included because they lacked the housework data). In the United States, Sweden and Australia, none of the interactions are significant. In Norway and Canada, however, some of the interaction terms are significant, indicating that the gender gap in authority does vary with family situation (see Wright 1997: 345–347 for details). In particular, in these two countries, as the proportion of housework done by *husbands* increases, the likelihood of married women having workplace authority

also increases. Only in Canada, however, does this interaction term generate a substantively large effect on the gender gap in authority. As Figure 9.2 indicates, the gender gap in authority in Canada for married women without children in the home, living in a relatively egalitarian household (a household in which husbands do 40% of the housework) is 0.23, whereas the gap for married women with children in the home living in an *in*egalitarian household (in which husbands do only 10% of the housework) is 0.82, comparable to the levels in Sweden and Norway. I can offer no explanation for why the patterns in Canada are so different from the other countries. For Canada, therefore, these interactions are consistent with the claims of the self-selection hypothesis that when women have high levels of domestic responsibility they frequently select themselves out of the running for positions of authority. Of course, the negative association between housework inequality and women's workplace authority in Canada could mainly reflect a causal impact of having authority on housework rather than of housework on the likelihood of getting authority and thus not support the self-selection hypothesis. In any case, there is little or no support for the self-selection hypothesis for the other countries in the study.

The glass-ceiling hypothesis

So far we have only discussed the differential likelihood of men and women *having* authority, but not the *amount* of authority that they have if they have any authority. This is a central issue for the "glass-ceiling hypothesis" – the idea that the gender gap in authority increases as one moves up authority hierarchies.

The strategy for evaluating the glass-ceiling hypothesis involves restricting the analysis to respondents in the authority hierarchy and then examining gender gaps in authority within this subsample. To do this, I use the formal position in the authority hierarchy variable as the criterion for restricting the sample: all persons who say that they are at least a supervisor on this question will be treated as in an authority hierarchy. On this restricted sample, we then examine the gender gap in authority for three dependent variables: being a middle-manager or above in the formal hierarchy; having sanctioning authority; and value on the 10-point amount of authority scale. The results for the first of these variables are presented in Figure 9.3. (The results for the other variables are in Wright 1997: 349).

If there is a strong glass-ceiling effect, then the gender gap for people

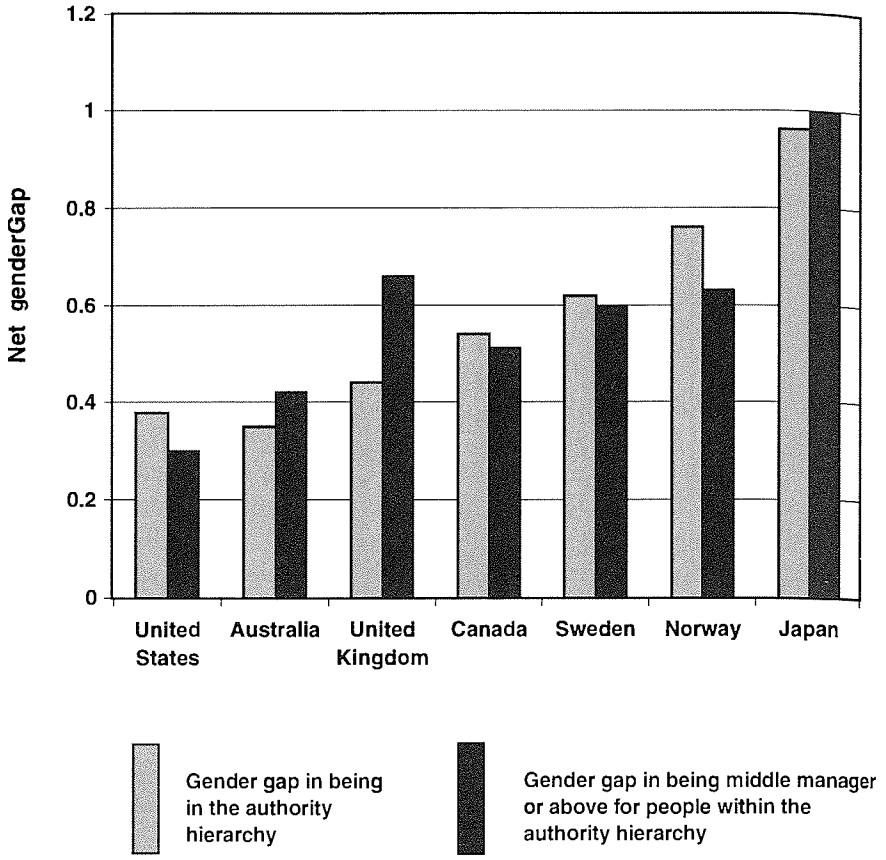

Figure 9.3 *Testing the "glass ceiling".*

within the authority hierarchy in the odds of a woman compared to a man being a middle-manager or above should be greater than the gender gap in simply being in the authority hierarchy. As the results in Figure 9.3 indicate, this is only nominally the case in three countries – Australia, the United Kingdom and Japan. And in only one of these, the United Kingdom, is the gender gap in authority substantially greater for people inside the hierarchy: in the UK the odds of a woman (net of the various control variables) already inside the hierarchy being a middle manager or above are 66% less than those of a man, whereas the odds of a woman being in the hierarchy altogether are only 44% less than those of a man. In all the other countries, there is either no difference between insiders and everyone in these gender gaps or the gender gaps for

people inside the hierarchy are somewhat less than for the labor force as a whole.

The lack of evidence for a glass ceiling is particularly strong in the US data. For all three of the measures of authority – middle manager or above, sanctioning authority, or the 10-point authority scale – the gender gaps in authority actually cease to be statistically significant for people inside of the hierarchy. In all countries except the United States, the gender gap in these measures of authority remain large and statistically significant when we restrict the sample to people in the hierarchy. Overall, therefore, these results do not lend support to the glass-ceiling hypothesis. Especially in the United States, it does not appear that once women are in the hierarchy, the barriers they face to promotion relative to men at least into the middle range of the hierarchy are greater than the barriers they faced in getting into the hierarchy in the first place.

Explaining cross-national variations

We have already examined, and rejected, one possible explanation for the differences across countries in the gender gap in workplace authority. These differences cannot be attributed to differences in the various compositional factors included in our analyses of the net gender gap since the basic pattern of intercountry differences is the same for the gross gender gap and the net gender gap in authority.

We will now explore somewhat less formally a number of general macro-social and cultural factors which might help explain the variations across countries in the gender gap. The results are presented in Table 9.1.

1. Gender ideology

The Comparative project on Class Structure and Class Consciousness contains a limited number of attitude items on gender equality. Respondents were asked how much they agreed, or disagreed with each of the following statements:

1 Ideally, there should be as many women as men in important positions in government and business.
2 If both husband and wife work, they should share equally in the housework and childcare.
3 It is better for the family if the husband is the principal breadwinner outside the home and the wife has primary responsibility for the home and children.

Table 9.1 *Rank ordering of countries from more to less egalitarian by the gender gap in authority and other relevant variables*

Net gender gap in having authority		Gender attitudes[c]		Legal gender egalitarianism[d]	
Rank order of countries[a]	Gender gap[b]	Rank order of countries	Mean score	Rank order of countries	Mean score
US	.36	Sweden	1.77	Norway	1.83
Australia	.45	Norway	1.82	Sweden	1.17
UK	.50	Canada	2.01	US	1.17
Canada	.58	Australia	2.05	Canada	−.48
Sweden	.79	US	2.17	UK	−.48
Norway	.82	Japan	2.43	Australia	−1.02
Japan	.98			Japan	−1.02

Gender earnings gap[e]		Occupational sex segregation (Index of dissimilarity)[f]	
Rank order of countries	Women's hourly earnings as percentage of men's	Rank order of countries	Mean score
Sweden	91.0	Japan	22.2
Norway	81.9	Australia	31.9
Australia	81.7	US	36.6
UK	74.0	Canada	41.0
Canada	66.0	Sweden	41.8
US	65.0	UK	44.4
Japan	51.8	Norway	47.2

Occupational sex segregation ("Ratio index of sex segregation")[g]		Proportion of the labor force in official managerial positions[h]	
Rank order of countries	Mean score	Rank order of countries	Proportion of labor force (%)
US	.65	Australia	15.8
Japan	.72	US	13.7
Canada	.75	Canada	12.2
UK	.92	UK	12.2
Australia	.95	Sweden	10.9
Sweden	.96	Norway	10.4
Norway	.99	Japan	5.9

Table 9.1 (*Continued*)

a. As indicated in Figure 9.1, the rank ordering of countries is virtually the same for the gross gender gap and the net gender gap.

b. The gender gap in workplace authority is defined as 1-Exp(b), where b is the coefficient for gender in the logistic regression predicting the overall authority dichotomy.

c. This is a simple index based on three Likert items concerning sex role attitudes. The lower the score the more egalitarian. The scores range from 1 to 4. The variable was not available for the United Kingdom.

d. This is a simple factor analytic scale of three legal rights for women: rights to abortion; rights to at least 12 weeks paid pregnancy leave; marital rape is a crime. See Charles (1992: 491–2).

e. *Sources.* Sweden, Norway, Australia, UK, Canada, US: National Committee on Pay Equity, "Closing the Wage Gap: an international perspective" (Washington, DC: National Committee on Pay Equity, 1988), pp. 10–14. *Japan: The Yearbook of Labor Statistics* (Geneva: International Labor Organization, 51st edn, 1992), pp. 798–804. There are some differences in the definitions for each country: Australia (1985), full time, average weekly earnings; Canada (1986), not specified; Japan (1984), average monthly earnings; Norway (1980), average hourly earnings in manufacturing; Sweden (1985), average monthly earnings, industry; United Kingdom (1985), average hourly earnings; United States (1987), median annual earnings.

f. Blau and Ferber (1990).

g. Charles (1992: 489).

h. This is defined as people in jobs which are described as "managerial positions" (but not supervisory positions) in the formal hierarchy variable.

A gender ideology scale was constructed by adding the responses to each item and averaging over the number of valid items. The scale ranges from 1, indicating a consistently strong egalitarian attitude towards gender roles, to 4, indicating a consistently strong conservative attitude.

As can be seen in Table 9.1, the rank ordering of countries in terms of their degree of ideological gender egalitarianism does not at all parallel the rank ordering for the gender gap in workplace authority. Sweden and Norway are the most ideologically egalitarian but have among the largest gender gaps in authority; the United States is exceeded only by Japan in the level of inegalitarianism ideologically, yet it has the smallest gender gap in authority.

2. Sexual and reproductive rights

Maria Charles (1992) has constructed an index of legally enforced gender egalitarianism based on three dummy variables: (1) abortions available on request, (2) marital rape is a crime, and (3) women are guaranteed at least 12 weeks of paid pregnancy leave from work. The scale values range from 1.83 to −1.02, where positive scores indicate more rights. As can be seen in Table 9.1 the rank order among the seven countries on this variable is both quite different from the rank ordering for gender attitudes and the rank order for the net gender gap in having authority.

3. Gender earnings gap

The gender gap in earnings is one possible indicator of institutional arrangements for gender equality within work which might impact on the gender gap in authority. Contrary to this expectation, however, the data in Table 9.1 indicate that there is no association between the level of the gender gap in hourly earnings and the gender gap in authority. Japan and the United States both have relatively large gender differences in earnings, yet the United States has a small gender gap in authority while Japan has the largest gender gap; Sweden and Norway are both relatively egalitarian in terms of gender differences in earnings, yet they both have relatively large gender gaps in authority.

4. Occupational sex segregation

As in the case of the gender gap in earnings, the expectation that the rank order of countries in sex-segregation of occupations should roughly mirror the gender gap in authority is not supported by the available data. Based on the data of two comparative studies of occupational sex segregation, Blau and Ferber (1990) and Charles (1992), the rank-ordering of our seven countries in terms of overall occupational sex segregation is not at all the same as the rank ordering in terms of the gender gap in authority (see Wright 1997: 354–355 for a more detailed discussion of these results, especially for Japan).

5. The proportion of the labor force with authority

The rank ordering for the size of the managerial category (as measured by the formal hierarchy variable) quite closely mirrors the rank ordering of the gender gap in authority: the four English-speaking countries have the largest proportion of their labor forces in managerial positions, followed by the two Nordic countries and then, with a much smaller

figure, Japan. It therefore does appear that the aggregate availability of managerial positions in the society may influence the size of the gender gap in the allocation of authority.

6. The organized women's movement and political culture

I know of no comparative research which systematically assesses women's movements in different countries either in terms of their organizational and political strength or in terms of the details of their ideological stance. Our analysis of these issues, therefore, will have to remain at a relatively impressionistic level.

In terms of the political strength of the women's movement one thing seems particularly clear: the women's movement in Japan is far weaker than in any other country. It is less obvious how to judge the relative strength of women's movements in the other six countries, although it seems clear that the politically organized women's movement in the United States would be among the strongest. On the basis of qualitative research by Katzenstein (1987), the US women's movement has generally been stronger and more powerful than movements in Europe. While this evidence is impressionistic, it seems fairly safe to say that the politically organized women's movement is probably weakest in Japan and strongest in the United States, with the other countries falling somewhere in between.

It is somewhat easier to make some judgments about the ideological orientations of different women's movements, at least if we are willing to assume that these women's movements are likely to reflect to a significant extent the broader political culture of their societies. Esping-Anderson (1990) classified capitalist democracies along a variety of dimensions characterizing the ideological principles within their welfare states. These are presented in Table 9.2. With the exception of the placement of Japan within these rank orderings, these political orderings closely parallel the rank ordering of the gender gap. Specifically, the four English-speaking countries score low on what Esping-Anderson terms "decommodification" (i.e. welfare state policies which reduce the dependency of workers on the market) and high on liberalism of regime characteristics whereas Norway and Sweden score high on decommodification and extremely low on liberalism.

How does this relate to the problem of the gender gap in authority? Liberalism is a doctrine which argues that markets are a legitimate and efficient means of distributing welfare so long as they are "fair." Eliminating ascriptive barriers to individual achievement in labor

Table 9.2 *Rank ordering of countries' political culture and institutions from liberal/commodified to socialist/decommodified*

Degree of "decommodification" in the welfare state [a]		Degree of liberalism in regime attributes[b]		Degree of socialism in regime attributes	
Rank order	Score	Rank order	Score	Rank order	Score
Australia	13.0	United States	12	United States	0
United States	13.8	Canada	12	Japan	2
Canada	22.0	Australia	10	Australia	4
UK	23.4	Japan	10	UK	4
Japan	27.1	UK	6	Canada	4
Norway	38.3	Norway	0	Norway	8
Sweden	39.1	Sweden	0	Sweden	8

a. From Esping-Anderson (1990: 52). This score is a measure of the extent to which the welfare state neutralizes the effects of the market through its welfare policies.

b. From Esping-Anderson (1990: 74). The score indexes the extent to which welfare state interventions follow the principles of classical liberalism. The socialism score indexes the extent to which the regime follows socialist principles. Because in Esping-Anderson's analysis there is a third form of regime, "classical conservatism," the rank ordering for socialism is not necessarily simply the inverse of the rank ordering for liberalism.

markets and employment relations is therefore a central objective of liberal politics. A women's movement animated by a liberal political culture, therefore, would be particularly concerned with equal rights and the elimination of such barriers. In keeping with this expectation, Goldberg and Uremen (1990: 28–30) have emphasized the relatively strong forms of antidiscrimination laws that have been passed in the United States and their relative effectiveness, at least compared to many other countries.

Social democracy, in contrast, questions the legitimacy of market-determined inequalities regardless of the equality of opportunity, and seeks to render human welfare at least partially independent of market mechanisms. A women's movement embedded in a social democratic political culture would be expected to be much less concerned with labor market mechanisms as such, and more concerned with state interventions which directly provide services and resources which enhance the welfare of women. Policy initiatives would therefore concentrate on such things as parental leaves, maternal health care, childcare services and child allowances. Women would certainly benefit in many ways from

such strategies, as many commentators on Scandinavian social democracy have stressed (Goldberg and Uremen 1990: 141–144; Moen, 1989), but these priorities would not directly impact on barriers to authority promotions in the workplace. Commenting on the contrast between American liberal feminism and European social democratic feminism, Nancy Fraser (1993) argues that the former adopt a "universal breadwinner" model of gender equality which emphasizes employment rights, whereas the latter adopt a "caregiver parity" model which stresses the provision of services and resources to equalize the conditions of life of women engaged primarily in domestic responsibilities. The relatively large gender gap in workplace authority in the social democratic Nordic countries, therefore, may in part be a by-product of the relatively lower priority placed on liberal goals of individual competition and achievement relative to more communal benefits.

Taking these various arguments together, I hypothesize that the variations across countries in the size of the gender gap in workplace authority is the result of the interaction between the relative scarcity or abundance of authority positions, on the one hand, and the capacity and interest of the politically organized women's movement to challenge the barriers to women being promoted into those positions on the other. Where there are relatively few managerial positions in the first place and the women's movement is particularly weak, as in Japan, the gender gap in authority will be very large. Where there are somewhat more managerial positions, but the women's movement is oriented towards collective goods and decommodified social provisions, the gender gap will still be relatively large. When there are relatively abundant managerial positions in the job structure and where the women's movement is relatively strong and oriented towards liberal individualist goals, the gender gap will be most effectively challenged.

Our evidence in support of these interpretations is rather sketchy, especially because we do not have cases of countries with a high proportion of the labor force in managerial positions combined with a weak women's movement, or countries with a strong, liberal women's movement and relatively few manager positions. Such cases would be needed to tease out the relative importance of these two factors and the nature of their interactions. Considerably more research is needed about the impact of women's struggles and the process by which the gender gap in authority changes over time within and across countries before these interpretations could be affirmed with confidence.

9.4 Conclusions

Several conclusions emerge from the research reported in this chapter. First, while a gender gap in authority exists in all of the countries we have studied, there is considerable cross-national variation in the magnitude of this gap: it is smaller in the English-speaking countries, especially in the United States and Australia, relatively large in the Scandinavian countries, and huge in Japan. These results appear quite robust across a variety of measures.

Second, the gender gap in authority within countries and the pattern of cross-national variations do not appear to be significantly the result of compositional factors among men and women in the labor force. Even when we control for a range of attributes of firms, jobs and individuals, the gap within every country and the basic pattern of cross-national differences remain. Furthermore, with the possible exception of Canada, there is little evidence that the gender gap in authority is attributable to self-selection processes by women. Much of the gender gap in workplace authority in the countries we have studied can thus provisionally be attributed to various forms of discrimination, at least some of which occur directly in the promotion process.

Third, the "glass-ceiling" hypothesis (at least in the relatively weak form we were able to investigate) is not supported in most of the countries in the study. While a gender gap in authority generally continues to exist when we restrict the analysis to people already in the authority hierarchy, this gap does not appear to be greater than the gap in acquiring authority in the first place. The commonly held view that the women's movement has been more successful in opening up positions at the bottom of the organizational hierarchy for women and less successful in moving women up the corporate ladder is not supported by these data.

Finally, and more tentatively, we have presented data which suggest that the variations in the gender gap across countries may be the result of the interaction between variations in the relative abundance of authority positions and the effectiveness of different women's movements in challenging barriers women face in moving into those positions. Both political and economic factors thus seem to be important in explaining variability in gender inequality in workplace authority, whereas cultural variations more specifically linked to gender ideology seem less significant.

Part III

Class structure and class consciousness