

Why so Unhappy? The Effect of Union Membership on Job Satisfaction*

Alex Bryson

Policy Studies Institute and Centre for Economic Performance

Lorenzo Cappellari

Università del Piemonte Orientale

Claudio Lucifora

Università Cattolica

This version: January 2003

Abstract

We investigate the effect of union membership on job satisfaction. We account for the endogenous selection induced by the sorting of workers into unionised jobs and use different methodologies to address the question of how the membership decision is related to overall job satisfaction and to satisfaction with pay. We analyse linked employer-employee data from the 1998 British Workplace Employee Relations Survey (WERS) using Propensity Score Matching and Instrumental Variables techniques. Both the linked structure of the data analysed and the estimation strategies employed allow us to deal with the limitations characterising the existing literature. We find that once the endogeneity of membership is accounted for, the unhappiness of union members disappears, suggesting that their reported dissatisfaction stems from individual unobservable attributes rather than from the union status, i.e. unionisation has no causal effect on satisfaction. When satisfaction with pay is considered, the union/non-union satisfaction differential drops considerably, consistently with the existence of wage premia in favour of members.

JEL classification: J28, J51

Keywords: Job satisfaction, Trade unions, Linked employer-employee data

*We are grateful to Alison Booth, Mark Bryan and seminar participants at the EALE conference 2002, the Università Cattolica and the Policy Studies Institute for useful comments. Bryson thanks the Leverhulme Foundation ("The Future of Trade Unions in Modern Britain") and the Regent Street Polytechnic Trust and Cappellari and Lucifora thank MIUR for funding. The authors acknowledge the Department of Trade and Industry, the Economic and Social Research Council, the Advisory, Conciliation and Arbitration Service and the Policy Studies Institute as the originators of the 1998 Workplace Employee Relations Survey data, and the Data Archive at the University of Essex as the distributor of the data. None of these organizations or individuals bears any responsibility for the authors' analysis and interpretations of the data and usual disclaimers apply.

1. Introduction

In this paper, we investigate the effect of union membership on job satisfaction. Whilst it is common to study the effects of union status on satisfaction treating individual membership as given, in this paper we account for the endogenous selection induced by the sorting of workers into unionised jobs. In particular we use two methodologies to address the question of how the membership decision is related to overall job satisfaction and to satisfaction with pay. We analyse linked employer-employee data from the 1998 British Workplace Employee Relations Survey (WERS) using Propensity Score Matching (PSM) and Instrumental Variables (IV) techniques.

We argue that most previous studies, due to either data deficiencies or estimation problems, suffer from severe limitations that are dealt with in the present paper. In particular, this paper offers three main contributions to the previous literature. First, it uses a much richer data set which links employer and employee characteristics, providing the ideal set of information (and instruments) to investigate the effects of union membership and job satisfaction. Second, the endogeneity of unionisation decisions is explicitly modelled, showing the likely biases arising in assuming it to be given, and jointly estimated with overall job satisfaction and satisfaction with pay. Third, different complementary methodologies are used to address and put under close scrutiny the different hypotheses which are found in the literature. The latter also provides a more general framework of analysis to a wider class of economic problems in which unobserved heterogeneity and endogeneity are important issues.

Our main finding suggests that once the endogeneity of membership is accounted for by appropriate matching of workers, the dissatisfaction of union members falls significantly. When pay is considered, members' dissatisfaction disappears, suggesting that unions are quite successful in providing compensating differentials for undesirable characteristics of the job to their members. Hence the traditional view of union members being inherently dissatisfied is revisited in our paper to consider a wider role for union members in the firm.

The paper is organised in the following way. Section 2 reviews the relevant literature. Section 3 describes the data and presents some stylised facts. In Sections 4 and 5 we discuss the main features of the estimation methods used. Section 6 discusses the results, and concluding remarks follow.

2. The Economics of Satisfaction and Membership

There has been growing interest in the analysis of workers job satisfaction and union membership with particular concern to general working conditions and specific workplace attributes. Whilst union objectives may be broadly characterised as an attempt to improve workers' welfare and satisfaction, by providing a better working environment and better wage conditions, a number of papers have come to

the puzzling conclusion that union members are generally less satisfied as compared to non-unionised workers (Hammermesh, 1977; Kochan and Helfman, 1981).

2.1. Reporting subjective satisfaction

The analysis of reported subjective satisfaction has been the object of a long standing debate among economists. Besides the standard caveat of cross-individual comparability, inference is further complicated by the fact that researchers generally don't know the set of information and the characteristics (both personal and of the work environment) that are taken into account when reporting the level of happiness. While we acknowledge the existence of some of these caveats, still we believe that useful information can be retrieved and a better understanding of the behaviour of agents can be gained by analysing individual (subjective) satisfaction.

In the empirical analysis we use reported subjective satisfaction to study how union membership affects the level of satisfaction of workers, while conditioning the latter on a large set of information both on personal characteristics and job attributes. Conditional on individual characteristics and job attributes, membership has an independent effect that needs to be adequately accounted for. In this paper union membership is not just considered as a job attribute but a specific choice of the individual which can give rise to strategic behaviour.

2.2. A review of the literature

Different explanations have been advanced in the literature for the apparently puzzling result that union members are more dissatisfied with their jobs than non-members. One line of argument has emphasised the role of individuals' unobserved heterogeneity by suggesting that those who experience lower satisfaction are more likely to join the union and be involved in union activity. In other words, being a union member *per se* (once the endogenous self-sorting of workers has been accounted for) should bear no relation with job satisfaction. A similar view considers workplace heterogeneity and argues that union workers are generally more dissatisfied since where discontent is high unions are more likely to set up a formal organisation – i.e. net benefits of unionisation are higher. Again, reverse causation arises, albeit in a form that calls for an approach dealing adequately with differences among jobs and workplaces instead of individuals. In practice, most previous studies raise the above points but, due to data limitations (restricted controls, a lack of sample representativeness and small sample sizes) or inappropriate methodology, have not succeeded in dealing adequately with the issues. Bender and Sloane (1998), in their study, try to deal with the selectivity issue by correcting for unobserved heterogeneity in their satisfaction equations. Their analysis is problematic since by forcing membership status to be equivalent to coverage status -- i.e. dropping uncovered members and covered non-members --, they cannot distinguish between union membership effects and effects arising from

workplace unionisation and other workplace characteristics affecting satisfaction. The latter aspect is emphasized by Gordon and Denisi (1995) who also stress the importance of distinguishing between job satisfaction and satisfaction with the union and management. Although the different aspects of overall job (dis)satisfaction are probably not easily separated out, it is central when analysing the effect of individual union membership on satisfaction to control adequately for workplace characteristics as well as satisfaction with management. This raises a last point, often overlooked in the literature but discussed in Gordon and Denisi (1995), namely whether it is in the interest of the union to whip up dissatisfaction once recognised by the employer. In this context, being part of the union provides workers with the possibility of expressing their discontent to the management. This mechanism is regarded as a more efficient way of conveying dissatisfaction as compared to the standard (costly) individual quit behaviour that workers would have to resort to in the absence of union "voice" (Freeman and Medoff, 1984). Hence, the "voice" mechanism attributes to union membership a causal effect in reporting dissatisfaction. Duane and Leigh (1986) analyse job satisfaction and the desire for unionisation and find union workers express a greater desire for unionism than non-members at every level of job satisfaction. Consistent with the union voice hypothesis, dissatisfied and moderately satisfied union workers have approximately the same probability of desiring union representation as do very satisfied union workers. Conversely among non-union members, worker desire for unionisation rises with job dissatisfaction¹. This has a number of interesting implications: first, quit rates rise more moderately among union members than non-members as dissatisfaction rises; second, desire for unionisation rises with higher pay among members whereas desire for unionisation falls with higher wages among non-members. Finally, given the importance of union bargaining over wage levels and workplace pay policies, in the analysis of overall job satisfaction specific treatment should be given to satisfaction with wages². Union status may increase satisfaction by requiring the payment of compensating wage differentials for undesirable attributes of jobs; alternatively, the union may raise discontent as a result of wage policies designed to reduce wage dispersion (Borjas, 1979; Hersch and Stone, 1990).

It should be noted that the implications of the above hypotheses are rather different. In one case, it is workers' (unobserved) heterogeneity together with the job attributes that produce the dissatisfaction result. In the other case, it is union status *per se* which prompts workers to express discontent through their "voice". Thus, if the former were to be the relevant explanation, once the unobserved characteristics have been correctly identified no other role for union dissatisfaction should persist. On the other hand, if joining the union leads to a greater likelihood of voicing discontent, either

¹ On the desire of union representation see also: Farber, 2001; Bryson and Gomez, 2002.

² Another reason for analysing satisfaction with pay separately from overall is that it might be easier to report satisfaction with respect to pecuniary aspects of the job rather than with non-pecuniary aspects.

real or in a more strategic way (as predicted in the exit-voice model), unionisation and workers' (dis)satisfaction will not be independent.

3. Patterns of Membership and job satisfaction

3.1 The data

The data are linked employer-employee data from the British Workplace Employee Relations Survey 1998 (WERS). With appropriate weighting, they are nationally representative of British employees working in workplaces with 10 or more employees covering all sectors of the economy except agriculture (Airey *et al.*, 1999). The data are rich, allowing for the inclusion of a wide range of individual-level and workplace-level controls to estimate influences on job satisfaction. A complete list of the variables used in the analysis and their means is contained in Appendix Table A1.1.

We use two elements of the survey. The first is the management interview, conducted face-to-face with the most senior workplace manager responsible for employee relations. Interviews were conducted in 2,191 workplaces between October 1997 and June 1999, with a response rate of 80%. The second element is the survey of employees where a management interview was obtained. Self-completion questionnaires were distributed to a simple random sample of 25 employees (or all employees in workplaces with 10-24 employees) in the 1,880 cases where management permitted it. Of the 44,283 questionnaires distributed, 28,237 (64%) usable ones were returned.

The sample of workplaces is a stratified random sample with over-representation of larger workplaces and some industries (Airey, *et al.*, 1999). Employees' probability of selection for the survey is a product of the probability of their workplace being selected and the probability of the employee's own selection. To extrapolate from our analyses to the population from which the employees were drawn (namely employees in Britain in workplaces with 10 or more employees) we weight the analysis using the employee weights.³ Our estimating sub-sample is all employees with complete information on the variables used in the analysis, namely 18,012.

3.2 Indicators of satisfaction

The survey asked each employee to provide a rating, on a five-point scale from 'very satisfied' to 'very dissatisfied', of how satisfied they were on four aspects of their job: the amount of influence they had over their job; the pay they received; the sense of achievement they got from their work; and the respect they got from supervisors and line managers. As the cross-tabulations in Appendix Table A1.2 show, union members are significantly less satisfied than non-members on all four indices. To obtain

³ The weighting scheme used in this paper compensates for sample non-response bias which was detected in the employee survey (Airey *et al.*, 1999: 91-92).

an indicator of overall job satisfaction we sum the times an individual says she is either ‘satisfied’ or ‘very satisfied’ on each of the four job aspects. This produces an ordinal scale running from 0 to 4, with higher scores indicating greater dissatisfaction. Again, members are significantly less satisfied than non-members.

3.3. Union membership and job satisfaction: a descriptive analysis

To get a better understanding of the relationship between union membership and job satisfaction in WERS data Table 1 shows membership coefficients estimated from job satisfaction equations. Throughout the paper we focus on the overall satisfaction indicator described above and on satisfaction with respect to pay. Given the discrete ordinal nature of satisfaction indicators, we utilise ordered probit regression. In addition, we account for the survey design by using survey weights and we deal with the presence of repeated observations within the same establishment by applying a robust variance estimator via pseudo maximum likelihood.

[insert Table 1 about here]

We start by presenting estimates that are unconditional on the set of individual and workplace attributes available in WERS. As can be seen from the first column, union members report satisfaction levels that are significantly lower than non-members. The gap is wider for overall job satisfaction than it is for pay satisfaction, indicating that members see themselves better off when it comes to pay compared to other aspects of their job. In column (2) we include a set of controls for personal characteristics, namely gender, education, marital status, parental status, health status and ethnicity. Although these controls have a significant impact on job and pay satisfaction, their introduction does not affect the estimated membership coefficient in the overall satisfaction equation, indicating that the negative estimate in column (1) was not the result of differences in personal characteristics between members and non members.⁴ The gap on pay satisfaction, however, does narrow quite markedly, though it remains statistically significant. Estimates in column (3) also control for occupational and job characteristics. The membership coefficient in the overall satisfaction equation drops a lot, indicating that members tend to be concentrated in occupations and jobs where satisfaction is low. The inclusion of these controls eliminates any difference in satisfaction with pay between members and non-members. As suggested by the literature reviewed in Section 2, lower job satisfaction among union members relative to non-members may also reflect unobserved heterogeneity at the firm or workplace level if unionisation is more likely to occur in places where working conditions are poor. Column (4) exploits the employer dimension of WERS and adds a set of controls for workplace characteristics,

⁴ Tables with coefficients estimated on all the controls used in these regression are available from the authors.

such as industry, size, region, establishment characteristics, gender and occupational composition of the workforce, union presence and density. Controlling for these factors does not seem to alter the conclusions reached so far, indicating that workplace heterogeneity does not contribute to the lower satisfaction of members. Column (5) adds controls for individual perceptions of industrial relations climate and opinions on the trade union, which might help in separating job satisfaction from satisfaction towards unions (see Bender and Sloane, 1998). The inclusion of these variables reduces the size of the membership coefficient in the overall job satisfaction equation still further, although the drop is lower compared to the one occurred when occupational and job controls were included and the member-non-member satisfaction differential remains significant. The membership coefficient on pay satisfaction remains small and statistically non-significant. Finally column (6) adds pay and hours worked to the controls. Their effect is to increase the dissatisfaction of members with both the aspects considered; in particular, the coefficient related to pay satisfaction is now statistically significant at the 20% level. There is evidence from other WERS analyses that membership confers a wage premium (Bryson, 2002). This evidence suggests that the higher wages of members compensate for their lower dissatisfaction, so that after the effect of pay on satisfaction has been controlled for the estimated membership coefficients gains size and significance.

Besides providing estimates of the association between membership and satisfaction, evidence in Table 1 shows that the estimated membership coefficient varies considerably with the set of regressors included in the analysis, indicating that membership is not distributed randomly across the cells defined by the available controls. Clearly, the same remark might apply to unobservable factors that might affect both satisfaction and membership, leading to a problem of endogeneity that might bias attempts at estimating the effect of membership on satisfaction. The remainder of the paper is devoted to the assessment of these endogeneity problems. First, we seek to isolate the causal effect of membership using the set of observable characteristics available in WERS to match members with like non-members using propensity score matching techniques. Second, we model the processes of union membership and job satisfaction simultaneously in order to estimate the unobservable correlation between the two and deliver estimates of the effect of membership on job satisfaction that are robust to endogeneity. While the former approach unravels causal effects by controlling for observable variables available in WERS, the latter also controls for unobserved factors inducing a co-movement between membership and satisfaction. Besides distributional assumptions invoked by both methods, the two approaches rely on different assumptions to identify the membership effect: statistical matching relies on the fact that all relevant variables affecting membership and status are observed by the researchers, while the second approach needs identifying restrictions.

4. The effect of membership on satisfaction: propensity score matching analysis

4.1 The method of matching⁵

Although members appear less satisfied than non-members, one can not attribute this difference to the causal effect of union membership where members differ systematically from non-members in ways which might affect their job satisfaction independent of membership. To establish whether the satisfaction differential is attributable to union membership, or is due to systematic differences in personal, job and workplace characteristics across members and non-members, we need to isolate the causal effect of union membership on satisfaction. In this section we use a semi-parametric statistical matching approach known as propensity score matching (Heckman et al., 1999) to compare satisfaction outcomes for unionised workers with ‘matched’ non-unionised workers to infer the causal effect of union membership on satisfaction. As with all non-experimental estimators, causal inference relies on an untestable assumption. In this case, the assumption is that the selection process is captured with observable data. (In the evaluation literature this is referred to as the conditional independence assumption.) For this key identifying assumption to be plausible, one must be able to control for all characteristics affecting both union status and satisfaction. This requires very informative data such as those available with the WERS linked employer-employee data.

Results from the standard cross-sectional regression techniques presented in Section 3 indicate average differences in satisfaction between union and non-union workers. The regression techniques reported in Section 3 share the conditional independence assumption underpinning matching such that the regression coefficient can only be interpreted as the causal effect of membership on satisfaction if variables entering the regression equation account fully for endogenous selection into membership status. But matching has two distinct advantages relative to regression in identifying an unbiased causal impact of membership on satisfaction. First, it is a non-parametric technique, so it does not require the assumption of linearity in the outcome equation. There is no clear indication from previous theoretical or empirical research on the relationship between membership and job satisfaction that the linear functional form assumption is appropriate. Secondly, matching estimators highlight the problem of common support and thus the short-comings of parametric techniques which involve extrapolating outside the common support. Matching is thus able to eliminate two of the three sources of estimation bias identified by Heckman, Ichimura, Smith and Todd (1998): the bias due to difference in the supports of X in the treated and control groups (failure of the common support condition) and the bias due to the difference between the two groups in the distribution of X over its common support. The other source of bias is the one due to selection on unobservables. This highlights the importance of the conditional independence assumption since, if this holds, selection on unobservables ceases to be a problem. The appropriateness of the conditional independence assumption is dependent on the

⁵ The theory and empirical implementation of propensity score matching are explained in detail in Appendix A2.

richness of the available data. In applying this technique, we are taking advantage of the very informative linked employer-employee data available in WERS.

4.2 The empirical implementation of matching

As noted above, the regression-adjusted estimate of the effect of membership on overall job satisfaction was to reduce the probability of being satisfied. However, for reasons discussed above, these estimates involve extrapolating beyond the common support and can not be interpreted as the causal effect of membership on satisfaction. Here we describe the empirical implementation of propensity score matching in WERS to yield an unbiased estimate of membership's effects on the satisfaction of union members.

Since the propensity to be a union member is unknown, the first task in matching is to estimate the propensity to be a union member. We do this with a probit estimating a (0,1) variable identifying individuals' union membership status. The estimation accounts for the complex sample design, that is, sampling weights, clustering and stratification. The conditional independence assumption requires that all variables influencing membership and satisfaction should be included in the estimate. Our choice of variables is informed by previous empirical work on satisfaction, discussed earlier, and the theory and empirical evidence on the worker choice and queuing models of selection into union membership (Farber, 2001; Bryson and Gomez, 2002). Variables entering the model are demographics (age, gender, marital status, health, qualifications), job-related (occupation, gender segregation), workplace (size, activity, industry, ownership, location) and local labour market conditions.⁶ Although our linked employer-employee data provide much of the requisite information, it is arguable that we are missing some data. For example, we have no data on motivation, which may be positively correlated with membership and satisfaction, potentially biasing the membership coefficient upwards. Our data set does contain workplace tenure and the amount of employer-provided training undertaken, both of which may be correlated with this tendency. However, because these variables may be influenced by membership itself, and are thus endogenous with respect to membership, their incorporation in the estimation of the propensity score could undermine the interpretability of estimated effects (Heckman, LaLonde and Smith, 1999). The omission of tenure from the matching analysis inflates the negative coefficient of membership on satisfaction since membership is associated with longer workplace tenure and lower job satisfaction.⁷ We return to the fact that matching is unable to condition on variables

⁶ This specification is the product of much experimentation. Variables such as ethnicity, nature of contract, age of establishment, foreign ownership were excluded from the model because they were not statistically significant in membership and/or satisfaction equations. Other variables, such as hours worked, were excluded as potentially endogenous with respect to membership. Later we report sensitivity of our results to alterations in the specification of the matching equation.

⁷ Two-thirds (66%) of members had been working at the workplace for at least five years, compared with 37% of non-members. Conversely, 39% of non-members had been at their workplace for under two years compared with only 7% of

that are endogenous to membership in discussion of the IV estimation where we are able to make use of these data.

The model estimating the probability of union membership for the whole private sector is presented in Appendix Table A2.1. Among non-members, the predicted probability of union membership ranges from .0004 to .9895, with a mean of .26 and a median of .17. Among members, the predicted probability ranges from .0040 to .9917, with a mean of .58 and median of .61. The zone in which there is no common support given by non-members is above .9895: enforcing common support at the extremes results in the loss of only 4 of the 7,444 members. Thus, the sub-group of members for whom we are unable to estimate the membership satisfaction effect is very small.⁸ The resultant distributions of propensity scores are presented in Graph 1. This shows that, although non-members' scores are bunched in the lower quartile of the distribution, they nevertheless offer support for members throughout the distribution.

As noted earlier, matching operates by constructing, for those participants with support, a counterfactual from the non-participants. There are a number of ways of defining this counterfactual using the propensity score. We use nearest neighbour matching which involves taking each treated individual (member) and identifying the non-treated individual (non-member) with the most similar propensity score. The matches were made with replacement so that, in some cases, a non-treated individual provides the closest match for a number of treated individuals, whereupon they feature in the comparison group more than once.⁹ The advantage of nearest neighbour matching is that the match is as good as it is possible to achieve in the sense that the bias across the treatment and comparison groups is minimised. In our case, the matches are very close: the mean difference in propensity scores between treated individuals and their matched comparators is .000132, and ranges between 0 and .008504. On the other hand, nearest neighbour matching disregards potentially useful information by not considering any matches of slightly poorer quality. Over-reliance on a reduced number of observations can result in effects being less precisely identified. Of the 10,568 non-members who could potentially have been matched to our 7,444 members, 3,070 were used as matched comparators. In half of cases (50.9%) these matched comparators have a match weight of 1 because they are matched to a single treated case. The largest weight is 31, and in only 57 cases is a non-member used as a match for 11 or more members. The mean match weight for non-members is 2.42.

members. Regression analysis revealed an independent association between membership and tenure (results available from the authors). The empirical literature suggests that membership increases tenure by reducing the likelihood of voluntary quits. These findings are consistent with the theory that unions provide a 'voice' alternative to quitting for dissatisfied workers (Freeman and Medoff, 1984).

⁸ As noted in Appendix A2, the treatment parameter we are estimating is the effect of treatment on the treated, that is, how members' satisfaction compares with what they would have received had they not been members, on average. In this case, only the condition $\Pr(D = 1 | X) < 1$ is required in enforcing common support.

⁹ Dehijia and Wahba (1998) find that allowing the non-treated to be used more than once as comparators improves the performance of the match.

To be effective, matching should balance characteristics across the treatment and comparison groups. Appendix Table 2.2 presents comparisons of the means in the characteristics used to match members and non-members, as well as a measure of the ‘distance’ of the marginal distributions of relevant characteristics in both groups (Rosenbaum and Rubin, 1985). For a given covariate, the standardised difference after matching is defined as the difference of the sample means in the treated and matched non-treated sub-samples as a percentage of the square root of the average of the sample variances in the treated and non-treated groups (Sianesi, 2001). Overall, the quality of the match seems good, the mean absolute standardised bias for all covariates being 0.68. Standardised bias for each variable tends to range from -6% to $+6\%$, and only twice does it exceed 10% .¹⁰

[insert Table 2 about here]

4.3 Results

In Table 2 we present proportions of members and non-members who expressed themselves satisfied on two of the five indices presented in Section 3, namely the overall index of job satisfaction and pay satisfaction. In both cases, we constructed a dummy variable identifying the percentage who said they were either ‘very satisfied’ or ‘satisfied’. In the case of the overall measure, those expressing themselves ‘very satisfied’ or ‘satisfied’ on at least two of the four constituent job satisfaction measures were said to be ‘satisfied’ on the dummy overall measure.¹¹

Column 2 of Table 2 shows that non-members were significantly more satisfied than members on the overall measure, with almost half (48.9%) declaring themselves either ‘very satisfied’ or ‘satisfied’ on at least two of the four constituent job satisfaction dimensions. A gap of around 8 percentage points separates members from non-members. Post-matching this gap closes to around 5.5 percentage points – a gap which remains statistically significant at a 99.9% confidence level.¹² This tells us that observable differences in personal, job and workplace characteristics between members and their matched counterparts accounts for around one-third of the satisfaction differential between members and non-members. This pattern of results, in which an appreciable gap between members’ and non-

¹⁰ Although achieving a reasonable balance on the X 's entering the participation equation is an indicator of how good the match is on observables, it cannot provide an indication as to whether the conditional independence assumption is plausible.

¹¹ Our estimates of satisfaction are for the population from which our sample was drawn, taking account of the complex survey design when comparing mean differences across members and matched non-members. In these population estimates, the survey sample weight of each treatment group member is applied to the corresponding matched comparator(s) (Frölich et al., 2001: 12). Hence, population estimates of the member-non-member differential are based on a weight incorporating both the matching weight and sampling weight. Population differences in satisfaction between members and their non-member comparators also account for variance arising from sample stratification and clustering.

¹² A complication that arises in the case of matching is that the estimation of the propensity score and the matching itself both add variation beyond normal sampling variation (Heckman, Ichimura and Todd, 1998). We run bootstrap estimates to account for this in the calculation of standard errors. Results are presented in footnotes to tables.

members' satisfaction pre-matching closes somewhat with matching, while remaining statistically significant, is repeated for all three of our non-pecuniary job satisfaction components.

A different picture emerges with respect to pay satisfaction. As noted in Section 3, employees are generally less satisfied with pay than with non-pecuniary aspects of their jobs, whether they are members or non-members. Around a third of employees express satisfaction with their pay. Furthermore, at 2.8 percentage points, the pre-matching gap between members' and non-members' satisfaction with pay is closer than it is on other satisfaction indices, although it remains statistically significant at a 95% confidence level. Post-matching, this gap falls to 2.1 percentage points, a gap that is no longer statistically significant. There are two possible interpretations of these findings.

- (a) The conditional independence assumption has been met, so the results reflect the causal impact of union membership

If we accept that membership is independent of satisfaction conditioning on those variables entering our propensity score estimator, they suggest that members are generally less satisfied with their jobs than they would be if they had not become members – except with respect to pay where their satisfaction levels are not significantly different from what they would have been as non-members. This is consistent with unions providing a wage premium for their members relative to similar non-members, perhaps by way of compensation for the types of jobs they perform. In essence, members display greater dissatisfaction than 'like' non-members about most aspects of their job but, because the union procures a wage premium for members relative to matched non-member counterparts, this attenuates their underlying propensity to express dissatisfaction.¹³ This interpretation is in line with previous research which showed the satisfaction deficit among members relative to non-members was appreciably lower on pay than on non-pecuniary aspects of the job (Berger et al., 1983; Evans and Ondrack, 1990; Meng, 1990). This wage differential is nevertheless insufficient to compensate for other aspects of their job, which is why a dissatisfaction gap remains on the overall job satisfaction index post matching.

If one accepts that our analysis fully accounts for endogenous selection into union membership, then we should have accounted for individual and job-specific characteristics generating greater dissatisfaction among members. One explanation for a causal impact of union membership which might explain the remaining dissatisfaction of members is the 'voice' effect of unionisation which engenders greater dissatisfaction than would otherwise be the case among members – for instance, through the politicisation of members to strengthen the union's bargaining hand, or through the greater

¹³ Using a similar matching technique on the same data set Bryson (2002) confirms that there is indeed a wage premium attached to membership, even among covered workers.

flow of information about the workplace which is available to members relative to non-members (Freeman and Medoff, 1984). The similarity between pay satisfaction among members and their matched non-member counterparts may be explained by unions' effectiveness in tackling pay issues. Certainly, there is evidence in the job satisfaction literature that unions are particularly effective in tackling employees' key job concerns – pay and job security – but less effective in tackling other 'softer' issues such as job quality (Meng, 1990).

(b) The conditional independence assumption has not been met

The second possibility is that, although the satisfaction deficit diminishes with matching, there are a number of factors which the analysis has not conditioned on, and which might conceivably affect both the propensity to unionise and job satisfaction. Failure to account for these factors means that results based on the matching estimator are biased. The direction of that bias depends on the influence of those factors missing from the analysis. It is possible that their inclusion could reduce the satisfaction deficit still further – perhaps even changing the sign of the union effect.

We shall return to (b) in detail when we discuss results using the IV methodology. However, in the light of the conjectures in (a) and (b) we conducted sensitivity analyses for the matching estimates which are reported in the next section.

4.4 Sensitivity analyses

Union recognition and union density

Earlier studies assessing the causal effect of union membership on job satisfaction have been criticised for failing to control for important differences between the work environments of union members and non-members. This led to studies based on samples of members and non-members drawn from the same working environment (eg. Gordon and Denisi, 1995). It is arguable that the presence of a union recognised for bargaining purposes may influence employees' job satisfaction through the union's ability to enhance employees' terms and conditions. The strength of the union may also be pertinent: as Borjas (1979: 24) suggests, only the strongest unions will be able to resist employer efforts to compensate for a union wage premium with a corresponding decrease in non-pecuniary job rewards. If this is so, one might expect job satisfaction to be highest where union strength maintains what Borjas (1979: 22-25) refers to as the 'full wage', comprising both the money wage rate and non-pecuniary job components. Since the presence and strength of a recognised union is also likely to influence employees' decisions to join a union – for instrumentalist or 'bandwagon' reasons (Bryson and Gomez, 2002) – we tested the sensitivity of our results to the inclusion of union recognition and union density

in our matching estimator.¹⁴ In fact, conditioning on union recognition and union density does not alter the results. As in the results presented above, matching closes the gap between member and non-member satisfaction. However, members remain significantly less satisfied on the overall job satisfaction measure than their matched counterparts who share a similar union environment. The satisfaction differential on pay becomes statistically non-significant post matching (closing by 1.9 percentage points to 0.9 percentage points ($P > F = 0.6591$)).

Conditioning on wage levels

If unions are successful in generating a compensating wage differential, we would expect members' underlying dissatisfaction with their jobs to emerge more strongly once wage levels had been taken into account. We therefore introduce wage levels and hours worked into our matching estimator to explicitly condition on wage rates.¹⁵ Wage levels were omitted from our baseline analyses because they may be viewed as endogenous with respect to union membership. However, their inclusion in sensitivity analyses is merited since the analysis helps 'net out' the effect of unions on wages.

[insert Table 3 about here]

As anticipated, members' pay dissatisfaction is actually accentuated post-matching once wages are entered in the matching estimator (Table 3). On the overall job satisfaction index, the gap in dissatisfaction between members and non-members is essentially unchanged post-matching. These findings are consistent with the proposition that unions increase members' satisfaction with pay by raising wages above their market level. Once this is accounted for, the underlying dissatisfaction of members with their jobs strengthens.

[insert Table 4 about here]

¹⁴ Union density is measured as a continuous percentage of all employees in the workplace belonging to a union. Union recognition is a simple (0,1) dummy indicator. Both were omitted from the baseline estimates for two reasons. First, it is arguable that union recognition is endogenous with respect to membership. Whether one treats it as exogenous or endogenous relies on assumptions regarding the sequencing of events, that is, whether recognition pre- or post-dates the individual's decision to join. Secondly, there are no theoretical reasons to assume that unionisation will affect non-members' job satisfaction. Unionisation will do so where union bargaining has spillover effects onto non-members' terms and conditions: these may be either positive or negative. Also, if membership secures higher rewards for similar jobs, the decline in non-members' relative position may result in lower satisfaction levels. Of course, if this were to persist over time, one would expect non-members to join the union.

¹⁵ Wages are the 12 banded categories for gross weekly pay. Hours is a continuous variable recording total weekly hours worked.

Alterations to the dependent variables to explore sensitivity of results to satisfaction in the tails of the distribution

Previous studies have indicated that the negative effect of union membership on job satisfaction is more precisely determined where the satisfaction measure is better able to distinguish between satisfied and dissatisfied workers (Borjas, 1979: 27). Table 4 confirms that results are a little sensitive to the construction of our dependent variables. Taking the overall job satisfaction measure first, members are less likely to be ‘very satisfied’ than non-members pre-matching. The gap closes a little after matching, but it remains statistically significant. However, matching reduces the membership differential among the ‘very dissatisfied’ by a half to 1.4 percentage points, a gap which is not statistically significant. Differences on pay satisfaction are not significant after matching, confirming earlier results. But it is notable that even pre-matching there is no difference between members and non-members saying they were ‘very satisfied’ with their pay. Rather the pre-matching differential referred to earlier is driven by members being more likely to say they are ‘very dissatisfied’ with their pay.

Is the union membership effect a ‘voice’ effect?

Although there is some debate as to whether it is in unions’ interest to raise employee dissatisfaction with their job once they have successfully organised a workplace (Gordon and Denisi, 1995), if there is a direct causal effect of union voice on job dissatisfaction, one might expect that effect to be more pronounced among union activists (Cappellari et al, 2002). This is because unionised workers tend to rely on those activists to convey their dissatisfaction to the employer (Bender and Sloane, 1998: 224). To test this proposition we compared job dissatisfaction levels for members with matched non-members having removed activists from the sample. If satisfaction differentials diminished substantially relative to the full matched sample – which includes activists – this might be construed as evidence of a voice effect driven largely by union activists. In fact, the member-non-member satisfaction differential was equally pronounced among non-activists post-matching, suggesting that greater job dissatisfaction among members can not be explained simply in terms of union voice effects.

5. The effect of membership on satisfaction: instrumental variables analysis

In this section we estimate the effect of union membership on job satisfaction while simultaneously modelling union membership status using instrumental variables. In this way we are able to control for the presence of unobserved correlation between union membership and job satisfaction, thus purging estimated effect from the bias induced by unobserved heterogeneity. We assume that the satisfaction propensity of individual i ($i=1, \dots, N$) is summarised by a continuous latent variable S_i^* which is a linear function of personal attributes represented by the vector X_i , a dummy variable M_i taking value 1 if the individual is a union member and 0 otherwise and an error term ϵ_i distributed as standard normal:

$$S_i^* = \mathbf{X}_i' \boldsymbol{\beta} + \delta M_i + \varepsilon_i \quad (5.1)$$

where $\boldsymbol{\beta}$ is a vector of coefficients associated to personal attributes and δ is the scalar coefficient associated to membership.¹⁶ The set of controls included in \mathbf{X}_i is the same featuring in the last column of Table 1, i.e. personal background, occupational and job characteristics, personal opinions on the climate of industrial relations and the trade union, gross hourly earnings and the set of workplace controls listed in Section 3. S_i^* is not observed; rather, in the WERS questionnaire we observe its discrete realisation S_i which takes values from 0 ('very dissatisfied') to 4 ('very satisfied') depending on S_i^* crossing a set of cut-off points τ_1, \dots, τ_4 .

In order to account for unobserved heterogeneity, we augment equation (5.1) with a probit equation for the probability of being a union member:

$$M_i^* = \mathbf{Z}_i' \boldsymbol{\gamma} + \mathbf{W}_i' \boldsymbol{\theta} + u_i \quad (5.2)$$

where M_i^* is a continuous latent propensity underlying the dummy M_i , \mathbf{Z}_i is a vector of observables, \mathbf{W}_i is a vector of variables that have no effect on satisfaction after union membership has been controlled for, $\boldsymbol{\gamma}$ and $\boldsymbol{\theta}$ are vectors of coefficients associated with personal attributes and u_i is an error term distributed as standard normal. We model the link between u_i and ε_i by allowing them to be distributed as bivariate normal with unrestricted correlation $\rho \equiv \text{corr}(\varepsilon_i, u_i)$. By simultaneously estimating equations (5.1) and (5.2) we are able to separately identify the correlation between unobservables – the coefficient ρ – and thus to purge the coefficient δ in (5.1) from the bias induced by unobserved heterogeneity.¹⁷

The set of attributes in \mathbf{Z}_i is a subset of those in \mathbf{X}_i . In particular, some job attributes have been excluded since they might form the bargaining object and be distributed across workers according to their membership status. Similarly, in the light of the literature on union membership wage differentials, hourly gross earnings do not enter \mathbf{Z}_i . In addition, individual opinions on the climate of industrial relations and the trade union do not appear in \mathbf{Z}_i , since they might be determined by the membership status.

The ‘‘instruments’’ in \mathbf{W}_i are given by a set of indicators of managerial opinions on the climate of industrial relations. The identifying assumption is that after membership status and worker’s opinions on industrial relations climate have been controlled for in the satisfaction equation, managerial opinions have no additional effect on satisfaction. On the other hand, they are the sole exogenous indicator of industrial relation climate in the membership equation, since individual opinions are likely to be shifted

¹⁶ Throughout the paper we assume vectors to be column vectors.

¹⁷ The advantage of this full information maximum likelihood approach over two stage methods is that it delivers asymptotically efficient estimates. Moreover, as Nicoletti and Peracchi (2001) point out, two stage procedures applied to non-linear equations of interest such as the ordered probit might not solve the endogeneity problems when the *true* unobserved correlation is large in size.

by the membership status, and have an effect on individual membership probabilities. Basically, our identifying strategy takes advantage of the availability of linked employer-employee data: we observe the climate of industrial relations from two different perspective – the worker’s and the manager’s ones – and we claim that the latter has no effect on individual satisfaction after the former (together with membership) has been controlled for. Since our model is not over-identified, we test the validity of our claims using functional form as an identifying restriction

[insert Table 5 about here]

We report estimates of the endogenised membership coefficients in Table 5; in order to facilitate comparisons, we also report coefficients from analogous models in which the union dummy is treated as exogenous – i.e. the coefficients from column (6) in Table 1. Before considering main results, we can notice that the tests for instruments validity reported at the bottom of the table are in favour of the use of managerial opinions on industrial relations climate as instruments for union status in satisfaction equations. In fact, these variable are not significant in the job satisfaction equations, whereas they significantly affect membership probabilities. Thence, the data support our proposed instruments.

The estimated coefficient on union membership from the Overall Job Satisfaction column reveals that, after endogenisation of the membership status, union membership has no residual effect on satisfaction, as is signalled by a t-ratio of 0.53. On the other hand, the negative relationship between membership and job satisfaction is now absorbed by the correlation coefficient ρ which has approximately the same size as the estimated coefficients on union dummy in the exogenous model, while it is statistically significant at the 15% level. Thence, these results reveal quite clearly that the negative estimate obtained from exogenous models is purely the effect of unobserved heterogeneity. The pool of union members is composed by individuals who have a satisfaction propensity lower than that in the population – as indicated by the negative ρ – and once this compositional effect has been taken into account no causal effect of membership on satisfaction can be found in the data. From an economic perspective, our results suggest that the same factors that lead individuals to unionise, as could be the case if individuals with higher expectations and aspirations from their working life were more likely to both unionise and see their aspirations frustrated once in the job.

Evidence about satisfaction-with-pay tells a story which is pretty similar to the one emerged from the overall satisfaction indicator. The coefficients on union dummy loses significance after endogenisation, while its negative sign is absorbed by the correlation coefficient. We can notice that now the statistical significance of the correlation coefficient is lower compared to the right column of the table, suggesting that endogeneity is less relevant in this case. This fact is in line with the estimate from the model with exogenous unionism, which was smaller in size and significance compared to the

overall satisfaction indicator. This finding confirms the importance of treating overall satisfaction and satisfaction with pay separately, since the higher wages granted by the union to their members might offset their larger intrinsic dissatisfaction generated by higher expectations.

The effect delivered by the IV estimator is free from the bias induced by endogenous selection of dissatisfied individuals into the pool of members. In order to provide further insights on the substantial validity of our instruments in solving the selectivity issue, we investigate the characteristics of those employees whose membership probabilities are shifted by the instruments. With this aim we use probit regressions to predict individual membership probabilities with and without the set of instruments among the regressors, remaining explanatory variables being the same as the ones underlying the estimates of Table 5. We next compute the individual shift in membership probability induced by the inclusion of the instruments: $\Delta P_i = \text{Prob}(M_i=1 | Z_i, W_i) - \text{Prob}(M_i | Z_i)$. This latter variable has an estimated mean of -0.0001 (median=-.0017) and first and third quartiles equal to -.0113 and 0.0077 respectively. We define individuals who had their membership probability reduced by the inclusion of the instruments as those below the first quartile of ΔP_i , and individuals who had it increased as those above the third one. Finally, we estimate the difference in the means of observable characteristics between the group with membership probability reduced and increased by the instruments, positive values of the difference indicating that the instruments have reduced the membership probabilities of individuals endowed with the observable attribute whose conditional means are investigated.

[insert Table 6 about here]

Table 6 reports the results of this experiment restricting the attention to those observable attributes for which the difference in conditional means is statistically significant at least at the 5% level. As it turns out, differences in means are positive and significant only for a set of variables indicating a good work environment, such as those signalling the possibility to discuss work issues with the supervisor or a good employee's perception of the industrial relation climate. The job dissatisfaction of these individuals is not due to bad working conditions, but, rather, is innate, leading them to both join the union and rank themselves as dissatisfied. By reducing the membership probabilities for this group of workers, the proposed set of instruments removes the effect of endogenous selection from the estimated membership coefficient in the job satisfaction equation.

6. Conclusions

Our analyses are consistent with previous studies in showing union members to be less satisfied with their jobs and pay than non-members. However, when we sought to isolate the causal impact of union

membership on satisfaction by matching members to 'like' non-members, we found that most of this satisfaction differential can be accounted for by differences between members and non-members in their demographic, job and workplace characteristics. These differences wholly account for the greater dissatisfaction of members for their pay, a finding that is consistent with unions procuring compensating pecuniary differentials for the poorer conditions faced by their members relative to similar non-members. Efforts to establish whether the remaining difference between members and non-members in overall job satisfaction were attributable to the voice effect of union membership indicated that voice effects could not account for the greater dissatisfaction of members.

This led us to investigate whether the differences between members and non-members in their expressed job satisfaction could be accounted for by unobserved differences between members and non-members which might account for selection into membership and lower job satisfaction. This seemed to be the case since, having conditioned on a wider range of variables than the matching estimates to endogenise union membership, we found that there was no significant difference between members' and non-members' job satisfaction. This suggests that there are features of the membership selection process that are not captured by the linked employer-employee data that are driving the satisfaction differential between members and non-members.

Bibliography

- Airey, C., Hales, J., Hamilton, R., Korovessis, C., McKernan, A. and Perdon, S. (1999) *The Workplace Employee Relations Survey (WERS) 1997-8: Technical Report*
- Berger, C. J., Olson, C. A. and Boudreau, J. W. 1983. 'Effects of Unions on Job Satisfaction: The Role of Work-Related Values and Perceived Rewards', *Organizational Behavior and Human Performance*, Vol. 32, pp. 289-324
- Binder, D.A. (1983), 'On the variances of asymptotically normal estimators from complex surveys', *International Statistical Review*, 51, 279-92.
- Borjas, G. J. (1979) 'Job Satisfaction, Wages and Unions', *The Journal of Human Resources*, Vol. 14, No. 1, pp. 21-40
- Bryson, A. (2002) "The Union Wage Premium: An Analysis Using Propensity Score Matching", Centre for Economic Performance Discussion Paper, forthcoming, London School of Economics.
- Bryson, A and R. Gomez. (2002). "You Can't Always Get What You Want: Frustrated Demand for Unionisation in Britain", Centre for Economic Performance Discussion Paper, forthcoming, London School of Economics.
- Cappellari, L., Lucifora, C. and Piccirilli, G. (2002): 'Union Activism, Workers' Satisfaction and Organisational Change', Università Cattolica di Milano, mimeo, February
- Dehejia, R. and Wahba, S. (1998). 'Propensity score matching methods for non-experimental causal studies'. NBER Working Paper #6829.
- Duane and Leigh 'Union preferences, job satisfaction and the union-voice hypothesis', *Industrial Relations*, Vol. 25, No. 1, 1986, 65-71
- Evans, M. G. and Ondrack, D. A. 1990 'The Role of Job Outcomes and Values in Understanding the Union's Impact on Job Satisfaction: A Replication', *Human Relations*, Vol. 43, No. 5, pp. 401-18
- Farber, H. (2001). 'Notes on the Economics of Labor Unions'. Working Paper #452. Princeton University, Industrial Relations Section.
- Freeman, R. B. and Medoff, J. L. (1984) *What Do Unions Do?*, New York, Basic Books
- Frölich, M., Heshmati, A. and Lechner, M. (2001). 'A Microeconomic Evaluation of Rehabilitation of Long-term Sickness in Sweden', St. Gallen Working Paper
- Gordon, M. E. and Denisi, A. S. (1995) 'A Re-examination of the Relationship Between Union Membership and Job Satisfaction', *Industrial and Labor Relations Review*, Vol. 48, No. 2, pp. 222-236
- Heckman, J., Ichimura, H., Smith, J. and Todd, P. (1998). 'Characterizing selection bias using experimental data'. *Econometrica* 66(5): 1017-1098.
- Heckman, J., Ichimura, H. and Todd, P. (1997). 'Matching as an econometric evaluation estimator: evidence from evaluating a job training programme'. *Review of Economic Studies* 64: 605-654.

- Heckman, J., LaLonde, R. and Smith, J. (1999) 'The Economics and Econometrics of Active Labor Market Programs' in Ashenfelter, O. and Card, D. (eds.) *The Handbook of Labour Economics*, Vol. III, Amsterdam: North Holland
- Hersch and Stone (1990) "Is union job satisfaction real?" *Journal of Human Resources*, XXV, 4, 736-751
- Holland, P. W. (1986). 'Statistics and Causal Inference', *Journal of the American Statistical Association*, December, Vol. 81, No. 396, 945-960
- Huber, P.J. (1967), 'The behaviour of maximum likelihood estimators under non-standard conditions', in *Proceedings of the Fifth Berkeley Symposium in Mathematical Statistics and Probability*, University of California Press, Berkeley CA.
- Meng, R. (1990) 'The relationship between unions and job satisfaction', *Applied Economics*, 22, 1635-1648
- Rosenbaum, P. and Rubin, D. (1983). 'The central role of the propensity score in observational studies for causal effects'. *Biometrika* 70: 41-50.
- Rosenbaum, P. and Rubin, D. (1985). 'Constructing a Control Group Using Multivariate Matched Sampling Methods that Incorporate the Propensity Score'. *The American Statistician* 39, 1: 33-38
- Sianesi, B. (2001). 'An Evaluation of the Active Labour Market Programmes in Sweden', IFAU Working Paper #2001: 5
- White, H. (1982), 'Maximum likelihood estimation of misspecified models', *Econometrica*, 50, 1-25.

A1. Appendix : data description

Table A1.1: Variables used in the analysis

	Mean	Appears in RHS of satisfaction equation	Appears in RHS of membership equation
Female	0.46	×	×
Aged less than 20	0.05	×	×
Aged 20-24	0.07	×	×
Aged 25-29	0.12	×	×
Aged 30-39	0.29	×	×
Aged 40-49	0.25	×	×
Aged 50-59	0.18	×	×
Aged 60 or more	0.04	×	×
Has children aged 0-4	0.15	×	×
Has children aged 5-11	0.20	×	×
Has children aged 12-18	0.20	×	×
Single	0.21	×	×
Widowed	0.01	×	×
Divorced	0.07	×	×
Has at most CGSE	0.12	×	×
Has at most O-levels	0.27	×	×
Has at most A-levels	0.15	×	×
Has at most graduate degree	0.16	×	×
Has at most post-graduate degree	0.05	×	×
Has none of these qualifications	0.24	×	×
Has vocational qualifications	0.39	×	×
Disabled	0.06	×	×
Non-white	0.03	×	×
Manager	0.10	×	×
Professional	0.13	×	×
Associate professional and technical	0.09	×	×
Clerical and secretarial	0.16	×	×
Craft and skilled service	0.11	×	×
Personal and protective service	0.08	×	×
Sales	0.10	×	×
Operative and assembly	0.13	×	×
Other occupations	0.11	×	×
Temporary job	0.04	×	
Fixed term contract	0.03	×	
Overtime always paid	0.47	×	
Job equally done by men and women	0.29	×	
Availability of flexible hours	0.76	×	
Availability of job sharing	0.15	×	

Availability of parental leave	0.06	×	
Availability of nursery	0.005	×	
Can take day off if needed	0.97	×	
Paid less than £50 per week	0.06	×	
Paid £51-£80 per week	0.07	×	
Paid £81-£140 per week	0.12	×	
Paid £141-£180 per week	0.09	×	
Paid £181-£220 per week	0.11	×	
Paid £221-£260 per week	0.11	×	
Paid £261-£310 per week	0.10	×	
Paid £311-£360 per week	0.08	×	
Paid £361-£430 per week	0.10	×	
Paid £431-£540 per week	0.08	×	
Paid £541-£680 per week	0.04	×	
Paid more than £681 per week	0.04	×	
Total hours worked on average week	36.65	×	
Has discussed with supervisor about how getting on with job	0.59	×	
Has discussed with supervisor about promotions	0.21	×	
Has discussed with supervisor about training	0.48	×	
Has discussed with supervisor about pay	0.30	×	
Thinks management understanding of employees' problems	0.54	×	
Thinks meetings management/employees are useful	0.59	×	
Thinks relations management/employees are good	0.54	×	
Thinks managers are in favour of trade unions	0.16	×	
Thinks trade unions take notice of members' problems	0.33	×	
Thinks trade unions taken seriously by management	0.26	×	
Thinks trade unions make a difference in work environment	0.21	×	
10 thru 24 employees	0.13	×	×
25 to 49 employees	0.15	×	×
50 to 99 employees	0.15	×	×
100 to 199 employees	0.15	×	×
200 to 499 employees	0.20	×	×
500 or more employees	0.22	×	×
Manufacturing	0.26	×	×
Electricity, gas water	0.01	×	×
Construction	0.03	×	×
Wholesales and retail	0.15	×	×
Hotels and restaurants	0.05	×	×
Transports and communication	0.06	×	×
Financial services	0.04	×	×
Other business and services	0.08	×	×
Public administration	0.08	×	×
Education	0.09	×	×

Health	0.12	×	×
Other community services	0.03	×	×
East Anglia	0.06		
East Midlands	0.09	×	×
London	0.10	×	×
North	0.07	×	×
North West	0.10	×	×
Scotland	0.10	×	×
Rest of the South East	0.18	×	×
South West	0.08	×	×
Wales	0.04	×	×
West Midlands	0.10	×	×
Yorkshire & Humberside	0.08	×	×
Share female employees	0.47	×	×
Share part time employees	0.24	×	×
Share managers	0.08	×	×
Share professionals	0.12	×	×
Share technicals	0.09	×	×
Share clericals	0.16	×	×
Share crafts	0.11	×	×
Share personal and protective services	0.08	×	×
Share salers	0.10	×	×
Share operatives	0.14	×	×
Share other occupations	0.12	×	×
Single establishment	0.24	×	×
Head establishment	0.32	×	×
Foreign owned firm	0.15	×	×
Public owned firm	0.28	×	×
Workplace aged less than 3 year	0.08		
Workplace aged 3 to 4 years	0.06	×	×
Workplace aged 5 to 9 years	0.15	×	×
Workplace aged 10 to 20 years	0.20	×	×
Workplace aged More than 20 years	0.51	×	×
Workplace produces goods or services for consumers	0.53		
Workplace supplies goods or services to other companies	0.23	×	×
Workplace supplies of goods or services to other parts of organisation	0.08	×	×
Workplace does not produce goods or services for sale in the open market	0.13	×	×
Workplace is an administrative office only	0.03	×	×
TTWA unemployment rate between 5 and 7 percent	0.33	×	×
TTWA unemployment rate between 7 and 7.5 percent	0.12	×	×
TTWA unemployment rate above 7.5 percent	0.06	×	×
Workplace union density	34.68	×	×

Trade union recognised	0.56	×	×
Manager thinks trade unions improve performance	0.36		×
Manager thinks employees fully committed	0.67		×
Manager thinks IR climate is good	0.86		×
Manager thinks management against trade unions	0.12		×

Appendix Table A1.2: Distribution of members and non-members across five measures of job satisfaction, WERS 1998

	Achievement		Pay		Respect		Influence		Overall	
	M	NM	M	NM	M	NM	M	NM	M	NM
V satis.	12.6	15.3	3.1	3.5	10.8	15.5	8.8	12.7	16.8	23.0
Satis.	45.9	50.4	30.3	32.8	40.5	46.0	44.4	49.6	24.0	25.9
Neither	22.3	21.3	22.7	24.5	21.3	19.7	26.6	25.1	19.0	20.2
Dissat.	13.0	9.2	30.4	27.7	15.6	11.7	15.2	10.5	19.4	15.7
V.dissat	6.1	3.8	13.4	11.5	11.8	7.0	4.9	2.2	20.8	15.2
P	0.0000		0.0205		0.0000		0.0000		0.0000	

Notes: (1) Figures are column percentages which take account of survey design. M = union members. NM = union non-members. (2) Tables are based on our estimation sample of 18,012. (3) P indicates the significance of a design-based Pearson correlation coefficient testing whether distributions in satisfaction are significantly different across members and non-members.

A2 Appendix: Causal inference through statistical matching

To establish whether the union job satisfaction deficit is due to membership, or is due to systematic differences in personal, job and workplace characteristics across members and non-members, we need to isolate the causal effect of union membership on satisfaction. Let us conceive of union membership as if it were a ‘treatment’ that the individual receives. We wish to evaluate the causal effect of this treatment (treatment 1) relative to non-membership (treatment 0) on an outcome variable, Y , job satisfaction. Let Y_1 be satisfaction if the individual received treatment 1 (that is, where the individual is a union member) and Y_0 be the satisfaction that would result if the same individual received treatment 0 (non-membership). Let us denote the binary indicator of the treatment actually received as $D \in \{0,1\}$, while X is a set of attributes which are not affected by the treatment (demographic, job and workplace-related).

The effect of treatment 1 on individual i as measured by Y and relative to treatment 0 is:

$$(A2.1) \quad \Delta = Y_{1i} - Y_{0i}$$

which is simply the difference between the individual's potential outcome if 'exposed' to membership and the individual's potential outcome from non-membership. To estimate the impact of membership on members' job satisfaction, it is necessary to know what the outcome would have been if the individual had *not* been a member. The problem is that we can not observe the counterfactual, namely the outcome which would have resulted if an individual had made an alternative choice (that is, if members had chosen non-membership, and *vice versa*). Either Y_{1i} or Y_{0i} is missing for each i . Thus our problem is one of estimating missing data. This counterfactual cannot be inferred directly from the outcomes of non-members since they are likely to differ substantially in their characteristics from members.

To overcome this selection problem, researchers must choose from a range of evaluation methods, the choice being determined by a number of factors including the richness of the data and the nature of the treatment. Because it is impossible to observe the individual treatment effect, each method relies on generally untestable assumptions to make causal inferences (Holland, 1986). In order to identify individual treatment effects, it is necessary to make very strong assumptions about the joint distribution of Y_{1i} and Y_{0i} . However, the *average* treatment effect at the population or sub-population level can be identified under generally less stringent assumptions, some of which are set out below. Among the parameters that only depend on the marginal distributions of Y_{1i} and Y_{0i} is the parameter most commonly estimated and the one estimated in this paper, namely the mean impact of treatment on the treated:

$$(A2.2) \quad \theta = E(Y_1 - Y_0 \mid D = 1, X) = E(Y_1 \mid D = 1, X) - E(Y_0 \mid D = 1, X)$$

where $D=1$ denotes treatment (membership), $D=0$ denotes non-treatment (non-membership) and X is a set of conditioning variables. In assessing the expected treatment effect for individuals who are union members, we are addressing the question of how members' job satisfaction compares with what it would have been had they not been members, on average.¹⁸

For members we observe Y_1 so that the average observed outcome for participants is an unbiased estimate of the first component of the effect of treatment on the treated $E(Y_1 \mid D = 1, X)$. The evaluation problem arises from the term $E(Y_0 \mid D = 1, X)$. This is the mean of the counterfactual which, since it is unobservable, must be identified and estimated on the basis of some usually untestable identifying assumptions justifying the use of the observable pairs $(Y_1, D = 1)$, $(Y_0, D = 0)$.

Members may not be a random sample of all employees. If there are systematic differences in characteristics across members and non-members that are likely to influence job satisfaction, failure to take account of these will bias any estimate of the union membership effect on satisfaction. Thus, $E(Y_1 \mid D = 1) - E(Y_0 \mid D = 0)$ would in general be biased for the effect of treatment on the treated. An

¹⁸ To obtain the average treatment effect on the non-treated $E(Y_1 - Y_0 \mid D = 0)$ the procedure is applied symmetrically. The average treatment effect $E(Y_1 - Y_0)$ is a weighted average of the treatment effects for the treated and non-treated.

exception is when the independence assumption $Y_0 \perp D$ can be invoked. This is credible where the random assignment of individuals to treatment ensures that potential outcomes are independent of treatment status. In this situation, $E(Y_0 | D = 1) = E(Y_0 | D = 0) = E(Y | D = 0)$ so that the treatment effect can be consistently estimated by the difference between the observed mean of the outcome variable for the treatment group and the observed mean for the non-treatment group.

In the absence of random assignment, one option is to construct a comparison group based on statistical matching. Matching estimators try to resemble an experiment by choosing a comparison group from all non-participants such that the selected group is as similar as possible to the treatment group in observable characteristics. Matching can yield unbiased estimates of the treatment impact where differences between individuals affecting the outcome of interest are captured in their observed attributes. This assumption, which is often referred to as the Conditional Independence Assumption (CIA), is the key identifying assumption underpinning the matching methodology. The precise form of the CIA depends on the parameter being estimated. For the treatment on the treated parameter, the CIA requires that, conditional on observable characteristics, potential non-treatment outcomes are independent of treatment participation. Formally,

$$(A23) \quad E(Y_0 | X, D = 1) = E(Y_0 | X, D = 0)$$

Thus, CIA requires that the chosen group of matched controls does not differ from the group of treated by any variable which is systematically linked to the non-participation outcome Y_0 , other than on those variables that are used to match them. This permits the use of the matched non-participants to measure how participants would have fared, on average, had they not participated.

The plausibility of the CIA depends on the informational richness of the data since the set of X 's should contain all the variables thought to influence *both* participation (that is, membership) *and* the outcome (job satisfaction) in the absence of participation.

Under CIA,

$$(A2.4) \quad E(Y_1 | D = 1) - E(Y_0 | D = 1) = E_{X|D=1} \{E(Y|X, D = 1) - E(Y|X, D = 0)\}$$

Hence, after adjusting for observable differences, the mean of the no-treatment (potential) outcome is the same for those receiving treatment as for those not receiving treatment. This allows non-participants' outcomes to be used to infer participants' counterfactual outcomes.

However, this is only valid if there are non-participants for all participants' values of X (this is known as the support condition):

$$(A2.5) \quad Pr(D = 1 | X) < 1$$

This ensures that all treated individuals have a counterpart in the non-treated population for each X for which we seek to make a comparison. If there are regions where the support of X does not overlap for the treated and non-treated groups, matching can only be performed, and the treatment

parameter, θ , retrieved, over the common support region. If treated individuals have no support in the non-treated population, they are dropped from analysis and the estimated treatment effect is redefined as the mean treatment effect for those treated falling within the common support.

Matching operates by constructing, for those participants with support, a counterfactual from the non-participants. There are a number of ways of defining this counterfactual.¹⁹ Once the counterfactuals are identified, the mean impact of the programme can be estimated as the mean difference in the outcomes of the matched pairs.

A refinement to the matching approach was introduced by Rosenbaum and Rubin (1983). If the CIA is met and there is common support then:

$$(A2.6) \quad Y_0 \perp D \mid P(X) \text{ for } X \text{ in } \mathbf{X}$$

where $P(X)$ is the propensity score, the conditional probability of participating in the programme – in our case, the probability of being a union member – given a vector of observed characteristics \mathbf{X} . Formally,

$$(A2.7) \quad P(X_j) = Pr(D_i = 1 \mid X_j)$$

Rosenbaum and Rubin show treatment and the observed covariates are conditionally independent given the propensity score, that is:

$$(A2.8) \quad D_i \perp X_i \mid P(X_j)$$

The advantage of Rosenbaum and Rubin's innovation is that the dimensionality of the match can be reduced to one. Rather than matching on a vector of characteristics, it is possible to match on just the propensity score. This is because, as Rosenbaum and Rubin show, by definition treatment and non-treatment observations with the same value of the propensity score have the same distribution of the full vector of regressors X . Having matched on the propensity score, the mean impact of the programme is estimated as the mean difference in the outcomes of the matched pairs.

If the CIA is satisfied, matching offers an attractive means of identifying the impact of union membership on job satisfaction. The main attraction is that it is non-parametric, avoiding the need to define a specific form for the outcome equation, selection process or unobservables in either equation. In addition, it avoids extrapolation beyond the common support which occurs with simple linear estimators. Heterogeneous treatment effects are allowed for, so no assumption of constant additive treatment effects for different individuals is required. Effects for sub-groups can be estimated by running the match on sub-populations (see Section 5.1). Matching estimators also highlight the problem of common support and thus the short-comings of parametric techniques which involve extrapolating outside the common support. Matching is thus able to eliminate two of the three sources of estimation bias identified by Heckman, Ichimura, Smith and Todd (1998): the bias due to difference

¹⁹ See, for example, Heckman et al. (1997) for a comparison of alternative matching schemes.

in the supports of X in the treated and control groups (failure of the common support condition) and the bias due to the difference between the two groups in the distribution of X over its common support. The other source of bias is the one due to selection on unobservables. This highlights the importance of the CIA since, if this holds, selection on unobservables ceases to be a problem. The appropriateness of the CIA is dependent on the richness of the available data.

Appendix Table A2.1: Individual union membership status

Demographics:	
<i>Female</i>	-0.129
	(3.18)**
<i>Age (ref.: under 25 years)</i>	
25-29 years	0.394
	(5.51)**
30-39 years	0.543
	(7.61)**
40-49 years	0.686
	(9.70)**
50-59 years	0.566
	(7.11)**
60+ years	0.253
	(2.38)*
<i>Qualifications (ref.: none, low or high)</i>	
GCSE or 'A' level	-0.087
	(2.27)*
<i>Health problem</i>	0.104
	(1.57)
<i>Married or living as married</i>	0.096
	(1.89)
Job-related:	
<i>Occupational classification (ref.: operative)</i>	
Manager/senior administrator	-0.956
	(9.04)**
Professional	-0.294
	(2.94)**
Associate Professional and technical	-0.305
	(3.12)**
Clerical and secretarial	-0.754
	(7.85)**
Craft and skilled service	-0.036
	(0.43)
Personal and Protective service	-0.754
	(6.70)**
Sales	-0.408
	(3.76)**
Other unskilled occupations	-0.483
	(5.43)**
<i>Gender segregation (ref.: roughly equal)</i>	
Type of work performed solely by men	0.146
	(2.39)*
Type of work performed mainly by men	0.064
	(1.24)
Type of work performed mainly by women	0.021
	(0.43)
Type of work performed only by women	-0.200
	(2.39)*
Workplace:	
<i>Size (ref.: 200+ employees)</i>	
10-24 employees	-0.459
	(4.51)**
25-199 employees	-0.439
	(7.12)**
<i>Single independent establishment</i>	-0.037
	(0.58)

Public sector	1.162
	(12.96)**
Workplace activity (ref.: supplier to other companies)	
Producers of goods/services for customers	-0.367
	(3.74)**
Produces goods/services for other parts of the organisation	-0.028
	(0.22)
Non-producer	0.050
	(0.74)
Administrative office only	-0.322
	(2.54)*
Industrial classification (ref.: manufacturing)	
Utilities	0.810
	(6.33)**
Construction	-0.585
	(4.57)**
Wholesale and Retail Distribution	-0.545
	(4.67)**
Hotel and Restaurants	-0.631
	(3.56)**
Transport and Communication	0.317
	(2.42)*
Financial Services	0.387
	(2.66)**
Other Business Services	-0.840
	(5.93)**
Public Administration	-0.110
	(0.78)
Education	-0.567
	(4.11)**
Health	-0.384
	(2.89)**
Other Services	-0.517
	(3.75)**
Location (ref.: rest of the South East)	
East Anglia	0.112
	(0.74)
East Midlands	0.202
	(1.85)
London	0.025
	(0.22)
North	0.539
	(3.33)**
North West	0.465
	(4.15)**
Scotland	0.312
	(2.92)**
South West	0.464
	(3.61)**
Wales	0.291
	(1.93)
West Midlands	0.333
	(2.47)*
Yorkshire and Humberside	0.303
	(2.09)*
Local labour market conditions:	
<i>TTWA unemployment rate of 5%+</i>	0.214
	(2.83)**

Constant	-0.475
	(3.20)**
Observations	18012
F-stat	(51,1532) = 36.14 Prob>F = 0.0000

Note: * = significant at a 95% confidence level. ** = significant at a 99% confidence level or above.

Appendix Table A2.2: Imbalance in means between treated and matched comparators, plus standardised differences (%)

	Non- members pre-match	Non- members matched	Members	% bias before match	% bias after match
Age					
25-29 years	.15	.12	.10	-13.42	-4.43
30-39 years	.28	.31	.30	6.23	-0.42
40-49 years	.23	.32	.33	23.51	2.33
50-59 years	.16	.18	.20	9.42	4.12
60+ years	.04	.03	.03	-5.12	-1.94
Female					
Married or living as married	.67	.76	.77	23.79	1.65
Health problem	.05	.05	.07	7.07	5.53
Mid-level qualifications	.45	.38	.40	-10.57	4.16
Occupational classification					
Manager/senior administrator	.15	.07	.08	-21.53	3.38
Professional	.11	.19	.20	25.32	1.82
Associate professional and technical	.08	.14	.11	9.36	-12.03
Clerical and secretarial	.25	.20	.20	-12.40	-1.75
Craft and skilled service	.07	.11	.12	16.71	2.42
Personal and protective service	.10	.08	.09	-2.64	4.79
Sales	.10	.04	.03	-27.66	-1.29
Other unskilled occupations	.08	.07	.06	-4.75	-1.27
Gender segregation					
Occupation performed solely by women	.08	.04	.04	-15.07	0.67
Occupation performed mainly by women	.27	.24	.26	-6.22	-2.80
Occupation performed mainly by men	.20	.21	.23	6.77	3.64
Occupation performed only by men	.13	.16	.17	11.85	4.80
Workplace size					
10-24 employees	.10	.04	.05	-19.41	3.48
25-199 employees	.63	.45	.52	-21.45	14.94
Public sector					
Single independent establishment	.17	.52	.52	78.34	0.85
	.25	.16	.15	-25.78	-2.30
Workplace activity					
Administrative office only	.07	.06	.05	-9.97	-4.70
Non-producer	.10	.17	.19	26.78	5.95
Supplier to rest of organisation	.06	.11	.09	12.04	-9.01
Supplier to other companies	.29	.14	.13	-42.04	-3.28
Industrial classification					
Utilities	.02	.08	.08	29.40	3.87
Construction	.06	.04	.05	-7.36	2.02
Wholesale and retail distribution	.19	.06	.05	-44.23	-1.69
Hotels and Restaurants	.06	.01	.01	-29.06	-0.86
Transport and communication	.04	.09	.10	21.00	3.96
Financial Services	.06	.07	.06	1.39	-2.92
Other business services	.14	.03	.03	-41.31	0.49
Health	.09	.12	.13	12.71	5.11
Education	.08	.14	.15	23.03	3.30
Public Administration	.05	.19	.16	37.32	-7.67
Other	.05	.03	.04	-7.05	3.49

Location					
London	.15	.12	.11	-10.69	-1.82
Yorkshire and Humberside	.08	.08	.08	2.78	3.36
North	.04	.07	.08	15.75	4.03
North West	.08	.12	.12	12.69	0.63
Scotland	.09	.11	.12	10.54	3.11
South West	.08	.09	.09	0.70	-1.40
Wales	.04	.05	.05	8.12	2.97
West Midlands	.09	.09	.09	-1.50	-2.27
East Midlands	.09	.08	.08	-1.03	0.19
East Anglia	.05	.04	.04	-6.51	-1.27
Unemployment rate of 5%+	.48	.54	.56	16.98	5.62
Average absolute standardised bias pre-match, whole sample					16.77
Average absolute standardised bias post-match, whole sample					3.62
Average absolute standardised bias pre-match, matched sample					0.36
Average absolute standardised bias post-match, matched sample					0.44
Absolute bias reduction					49.65

A3. Appendix : IV estimation

The model estimated in Section 5 is an Ordered Probit with Endogenous Dummy variable. Endogenisation of the membership dummy is achieved by simultaneously estimating the job satisfaction equation and a union membership probit. The two equations of interest are equations (5.1) and (5.2), which we report below for convenience:

$$S_i^* = X_i' \beta + \delta M_i + \varepsilon_i \quad (\text{A3.1})$$

$$M_i^* = Z_i' \gamma + W_i' \theta + u_i \quad (\text{A3.2})$$

see the main text for notation. The two error terms are allowed to be jointly distributed as bivariate normal with 0 means, unit variances and correlation coefficient ρ :

$$(\varepsilon_i \ u_i) \sim N_2[(0,0);(1,1,\rho)] \quad (\text{A3.3})$$

where N_2 denotes the bivariate normal density function.

The two continuous variables on the left hand side of (A.1) and (A.2) are unobserved. Rather we observe a discrete ordered variable S_i taking values from 0 to 4 depending upon S_i^* crossing a set of cut off points τ_1, \dots, τ_4 , and the membership dummy M_i . Thence the event space is partitioned into the ten outcomes defined by the intersection of S_i and M_i . From the viewpoint of the construction of sample likelihood there are three relevant cases to take into account. First of all let us define an index $k_i \equiv 2M_i - 1$ which takes value 1 for members and -1 for non members. For an individual reporting the lowest satisfaction score, the relevant probability is:

$$L_i = \Phi_2[\tau_1 - (X_i' \beta + \delta M_i), k_i(Z_i' \gamma + W_i' \theta); -k_i \rho] \quad (\text{A3.4a})$$

where Φ_2 is cumulative density function (c.d.f.) of the bivariate normal distribution.

For an individual scoring an intermediate level of satisfaction, say with $\tau_{j-1} < S_i^* < \tau_j, j=1 \dots 4$, the relevant probability is:

$$L_i = \Phi_2[\tau_j - (X_i' \beta + \delta M_i), k_i(Z_i' \gamma + W_i' \theta); -k_i \rho] - \Phi_2[\tau_{j-1} - (X_i' \beta + \delta M_i), k_i(Z_i' \gamma + W_i' \theta); -k_i \rho] \quad (\text{A3.4b})$$

Finally, for individuals reporting the highest satisfaction score, the relevant probability is:

$$L_i = \Phi[k_i(Z_i' \gamma + W_i' \theta)] - \Phi_2[\tau_4 - X_i' \beta + \delta M_i, k_i(Z_i' \gamma + W_i' \theta); -k_i \rho] \quad (\text{A3.4c})$$

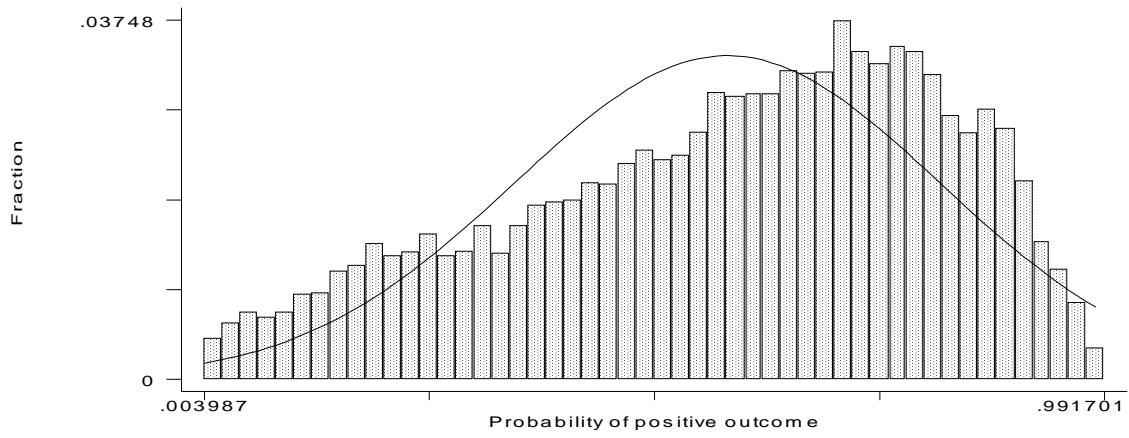
where Φ is the c.d.f. of the univariate standard normal.

By assuming that observations are independently and identically distributed (i.i.d.) across individuals the log-likelihood function can be derived as $l = \prod_i \log(L_i)$. However, the availability of repeated observations within the same establishment means that the i.i.d. assumption might be violated.

To account for this, we used a Pseudo Maximum Likelihood (PML) estimator that allows for arbitrary correlations between observations within the same establishment. See *inter alia* Huber (1967), Binder (1983) and White (1982). In addition, we weight each contribution to the pseudo likelihood in order to allow for differential sampling probability among establishment, using weights provided in WERS. Thence, our estimator is a Weighted Pseudo Maximum Likelihood (WPML).

Graph 1: Predicted union probability for members and non-members

Members



Non-members




Table 1: Estimated coefficients on union membership dummy from job satisfaction equations

Job Satisfaction	(1)		(2)		(3)		(4)		(5)		(6)	
	Overall	With Pay	Overall	With Pay	Overall	With Pay	Overall	With Pay	Overall	With Pay	Overall	With Pay
Union Member	-0.222 (9.43)	-0.091 (2.70)	-0.217 (8.77)	-0.068 (1.96)	-0.157 (6.07)	-0.007 (0.24)	-0.155 (4.88)	-0.034 (1.08)	-0.121 (3.80)	-0.016 (0.50)	-0.135 (4.17)	-0.047 (1.34)
Personal characteristics	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Job characteristics	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Workplace characteristics	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Opinions on industrial relations and trade unions	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes
Pay	No	No	No	No	No	No	No	No	No	No	Yes	Yes
Log likelihood	-28712.33	-25850.52	-28389.60	-25650.58	-27835.12	-25332.65	-27721.64	-25071.46	-25022.29	-24016.22	-24871.74	-23348.97
Model p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Note. N=18,012. Ordered Probit Weighted Pseudo Maximum Likelihood estimates use survey stratification weights and account for the presence of repeated observations on the same establishment. Asymptotically robust t-ratios in parentheses.

Table 2: Mean percentage satisfied among members and non-members pre- and post-matching

	Overall job satisfaction	Pay satisfaction
Pre-match		
Members (N = 7,444)	40.8	33.5
Non-members (N = 10,568)	48.9	36.3
Significance	F(1, 1582) = 53.34 P > F = 0.0000	F(1, 1582) = 4.87 P > F = 0.0275
Post-match		
Members (N = 7,440)	40.8	33.5
Non-members (N = 3,070)	46.3	35.6
Significance	F(1, 1392) = 10.61 P > F = 0.0012	F(1, 1392) = 1.30 P > F = 0.2552

Notes: (1) The 95% confidence interval for the post-match bootstrapped estimates are 2.2% to 8.8% for overall job satisfaction and -1.3% to 5.6% for pay satisfaction.

Table 3: Mean percentage satisfied among members and non-members pre- and post-matching: conditioning on wage levels, hours, recognition and density

	Overall job satisfaction	Pay satisfaction
Pre-match		
Members (N = 7,444)	40.8	33.5
Non-members (N = 10,568)	48.9	36.3
Significance	F(1, 1582) = 53.34 P > F = 0.0000	F(1, 1582) = 4.87 P > F = 0.0275
Post-match		
Members (N = 7,440)	40.8	33.5
Non-members (N = 3,070)	49.0	37.1
Significance	F(1, 1269) = 17.09 P > F = 0.0000	F(1, 1269) = 3.15 P > F = 0.0762

Notes: (1) In addition to the variables entering the propensity score estimator in the baseline analysis, these analyses incorporate banded gross weekly wages, total hours worked, union recognition and union density. (2) The 95% confidence interval for the post-match bootstrapped estimates are 4.3% to 12.5% for overall job satisfaction and 4.1% to 11.6% for pay satisfaction. (3) Diagnostics for matching: 12 out of 7,444 members lost through common support enforcement. Mean difference in propensity scores for treated and matched non-treated: .0001598. Maximum difference: .0014725. Mean match weight for non-members: 3.04, maximum = 41. Mean absolute standardised bias post-matching: 0.05.

Table 4: Mean differences in percentages very satisfied and very dissatisfied between members and non-members pre- and post-matching

	Pre-match	Post-match
Overall job satisfaction		
Very satisfied	-3.4%****	-2.5%**
Very dissatisfied	+3.1%****	+1.4%
Pay		
Very satisfied	-0.5	+0.1
Very dissatisfied	+1.9%**	+0.7

Notes: (1) * significant at a 90% confidence level. ** significant at a 95% confidence level. *** significant at a 99% confidence level. **** significant at a 99.9% confidence level. (2) These estimates are based on the baseline matching estimator which excludes wages, hours, recognition and density (3) Those identified as 'very satisfied' on the overall job satisfaction index are those saying they were 'very satisfied' on at least one of the four separate job components. They accounted for 10.5% of employees. Those identified as 'very dissatisfied' on the overall job satisfaction index are those saying they were 'very dissatisfied' on at least one of the four separate job components. They made up 6.6% of employees. The dummy variable identifying those who were 'very satisfied' with their pay is simply a recode of the pay satisfaction index, as is the dummy identifying those who were 'very dissatisfied'.

Table 5: Estimated coefficients on union membership dummy from job satisfaction equations

	Overall Job Satisfaction		Satisfaction with Pay	
	Exogenous	Endogenous	Exogenous	Endogenous
Union Member	-.135 (4.17)	.078 (0.53)	-.047 (1.34)	.123 (0.71)
Personal characteristics		Yes		Yes
Job characteristics		Yes		Yes
Workplace characteristics		Yes		Yes
Opinions on industrial relations and trade unions		Yes		Yes
Pay		Yes		Yes
Correlation of unobservables between satisfaction and membership equation (ρ)		-.124 (1.52)		-.101 (1.05)
Exclusion of instruments from satisfaction equation, p-value		0.5812		0.4540
Exclusion of instruments from membership equation, p-value		0.0001		0.0001
Log likelihood	-24871.74	-31263.589	-23348.97	-29759.439
Model p-value	0.0000	0.0000	0.0000	0.0000

Note. N=18,012. Ordered Probit with Endogenous Dummy Weighted Pseudo Maximum Likelihood estimates use survey stratification weights and account for the presence of repeated observations on the same establishment. Asymptotically robust t-ratios in parentheses.

Table 6: Conditional means of selected observable attributes of employees whose membership probabilities are affected by the instruments

	Mean if Prob(M)↓	Mean if Prob(M)↑	Difference	t-stat
Discussed with supervisor issues about				
Job in general	0.607	0.544	0.063	(2.956)
Promotion	0.225	0.188	0.037	(2.217)
Training	0.528	0.451	0.077	(3.403)
Management is understanding	0.546	0.473	0.073	(4.102)
Meetings with management are useful	0.629	0.538	0.091	(4.204)
Good relations with employees and management?	0.544	0.447	0.098	(4.661)

Note: only those attributes for which the difference in conditional means is statistically significant at least at the 5% level are reported.