High Unemployment Yet Few Small Firms: The Role of Centralized Bargaining in South Africa

Jeremy R. Magruder*

May 8, 2009

Abstract

South Africa has very high unemployment, yet few adults work informally in small firms. One explanation is that unions extend arbitration decisions to non-unionized firms, raising wages. These agreements are enforced in a spatially discontinuous way; employment effects of these bargaining councils are identified through spatial fixed effects. This approach represents a methodological improvement over sample restrictions used in other spatial discontinuity studies. Bargaining councils are found to decrease employment by 6-11%, with larger decreases among small firms. These effects are not explained by resettlement to uncovered areas, and are robust to a wide variety of forms for average spatial heterogeneity.

1 Introduction

In determining the optimal level of worker protection, governments face a tradeoff between higher wages and better benefits for workers, and lower employment rates if the marginal cost of labor increases. If labor protection is extended to small firms with presumably low levels of capital and labor productivities, employment effects and entreprenuerial disincentives may be strong, suggesting that labor protection policies may have an effect not only on employment levels but on the distribution of firm sizes. Of course, the ultimate effect of any set of arbitration decisions or government policies is an empirical question, and the

*University of California, Berkeley. I thank Guojun He and Charles Seguin for excellent research assistance and John Bellows for alerting me to some of the data used in this study. I also thank Michael Anderson, Lori Beaman, Tim Conley, Ann Harrison, Paul Schultz, Chris Udry and seminar audiences at Berkeley, USC, UCSD, and the World Bank for many helpful comments. All remaining mistakes are naturally my own.

challenge posed to the empirical economist is to identify an adequate control group of regions or industries which would have experienced similar employment outcomes were it not for differential exposure to labor regulations.

While the tension between worker protection and employment opportunities is one that virtually every government faces, the case for a reconsideration of policy is made particularly frequently and it's consequences may be particularly severe in South Africa. In South Africa, much of the labor force is unionized. Wages are high for a country of its level of economic development. While not mandated, employer organizations and unions may opt to participate in bargaining councils, which extend arbitration agreements beyond the firms and unions which make them to all workers in an industry in a given political demarcation, regardless of firm size or participation in the arbitration process. Whether these agreements reduce employment, particularly in small firms, has been extensively debated. The immediacy of this debate is heightened by South Africa's severe and peculiar employment situation, where unemployment is extremely high (at 20-40% for prime-age adults), with long durations being common. Despite the shortage of formal employment opportunities, few individuals either start their own firms or work for small firms, informal or otherwise.

Other potential motivations for South Africa's unemployment without entrepreneurship abound. For example, high wages and an extensive social safety net may increase the demand for leisure or render long periods of unemployed search more palatable. Entrepreneurial opportunities may be limited by low skill levels, liquidity constraints, and high crime, while high capital stocks may make large firms competitive with relatively low labor inputs. Identifying the effects of labor regulations, therefore, requires a careful analysis which would hold these conditions constant. This challenge, combined with a previous lack of available data on these agreements, has prevented clean estimates of the impact of bargaining council agreements on employment or industrial structure. As an additional challenge, these agreements are outcomes of a complex bargaining process between unions and firms with unclear and likely anti-competitive motives. Since centralized bargaining is

not mandated, the firms which choose to pursue centralized bargaining may be those whose local labor market conditions are such that centralized bargaining would represent a large competitive advantage. Therefore, the identification assumptions of similarity in level, trend, or response to a bargaining council agreement which would be necessary for a simple OLS or difference-in-differences approach may not hold in practice.

However, local labor markets are spatially continuous across the intranational political boundaries in South Africa, while these agreements are enforced in a spatially discontinuous way. In fact, this context is near-ideal for spatial discontinuity, as agreements vary with space, across industries, and over time, creating enough policy groups to avoid the common error-component problems which have plagued existing labor regulation studies. This paper creates a database of industrial bargaining agreements and adopts the spatial fixed effects proposed in Conley and Udry (2008) and Goldstein and Udry (2008), and argues that this spatial-fixed-effects approach represents an improvement over restricting samples to border areas, which is used in previous spatial regression discontinuity estimations. This spatial discontinuity reveals large effects: industries which have an agreement in a particular town in a given year have about 8-14% lower employment and 10-21% higher wages then the same industry in uncovered neighboring towns. Firm sizes are also impacted, with 7-15% fewer employees in small firms, and 7-12% fewer entrepreneurs, while there are smaller and insignificant effects on large firms and single employee firms. Utilizing town-year and town-industry fixed effects, I illustrate that these spatial discontinuities are similar in magnitude and precision whether inter-industry variation (within a town-year) or intertemporal variation (within a town-industry) is utilized. I further illustrate that, while both large and small firms move across borders in order to avoid these agreements, this border-jumping does not drive the employment effects measured here, so that these reductions in employment represent a net loss for the economy. Finally, I introduce weighted spatial fixed effects to test the robustness of estimates to a wide variety of potential spatial heterogeneities, and find that estimates are quite robust.

2 South Africa's Missing Small Firms

Unemployment in South Africa is extremely high, particularly among non-whites. Table 1 reports data from the 2003 Labour force survey (described below), which indicates that only about 50% of 20-60 year old African men and 35% of 20-60 year old African women are actually working, while 20-44% of each group is unemployed, depending on whether the broad or narrow measure is utilized. Kingdon and Knight (2006) advocate the broad unemployment measure in this context, as local wages are more sensitive to that measure. A large number of potential reasons for this unemployment exist, and the unemployment numbers and potential contributors for them are surveyed more extensively in a series of papers by Kingdon and Knight (e.g. 2004, 2006, 2008) as well as Banerjee et al (2008). Wages are high, due to high capital/labor ratios, a strong union presence, and extensive governmental labor market regulation in addition to the industrial bargaining agreements which are the focus of this paper (e.g. Butcher and Rouse 2000, Schultz and Mwabu 1998). Second, entrepreneurial skills may be absent in the population, as informal employment was squashed under Apartheid (e.g. Kingdon and Knight 2004). Third, some unemployment may be voluntary; a generous non-contributory pension program combined with the high wages earned by the employed leave many unemployed individuals with networks capable of supporting them (e.g. Bertrand, Mullainathan, and Miller 2003 for labor supply effects; Edmonds, Mammen, and Miller 2005 for network effects of pensions on living arrangements).

While it is clear that many adults are unemployed in South Africa, and there are a variety of potential motivations for this unemployment, it is unclear what adults are in fact doing. Labor force surveys in South Africa go to great lengths to measure any economic activity, identifying as workers individuals who engage in unpaid household work or tend household plots "even for only one hour" in the past week; this approach yields the low employment numbers described above. A very natural response to this unemployment would be for many to be either self-employed or working informally[1]. Yet columns 5 and 6 of table 1 reveals that

[1] This is particularly true as unemployment durations are very long, and there is some evidence that social

only 5-6% of prime age black South Africans are self-employed. A number of women do work in domestic service, at around 8% of the female population, which may similarly represent entrepreneurial activity. Nonetheless, these numbers are tiny compared to countries with similar levels of unemployment (e.g. Charmes 2000, Kingdon and Knight 2004). Moreover, what is perhaps most striking is that there are relatively few employees of small firms in general. Table 2 reports the percent of employees in each firm size category in South Africa. Particularly for men, we see very few workers in firms of fewer than 5 employees. For comparison purposes, I also include similar data from the 1995-96 Brazillian LSMS survey[2]. We see that, while unemployment is a great deal higher in South Africa (particularly for men), the distribution of firm sizes looks fairly similar – with one big exception. What is missing in South Africa, compared to Brazil, are the small firms with 2-4 employees.

Of the above explanations for high unemployment, one in particular which may suggest minimal small-scale employment in a high unemployment context is labor regulation. The South African labor market is highly regulated, with a variety of legislated labor standards as well as privately bargained arbitration decisions. Unlike in many other developing countries, South Africa is successful in enforcing labor and tax regulations on many small (and potentially informal) firms; an influential study found that the average business with fewer than 5 employees pays nearly R14000 (about $2170) per employee in costs associated with tax and labor regulations[3](SBP 2005). Moreover, unions and firms can extend labor standard arbitration to all workers in a given political district through bargaining councils. Small businesses, in particular, have advocated aggressively against the extension of these labor arbitration decisions; in 2005 South African President Thabo Mbeki announced that small

connections may be important to find employment. Since jobs are scarce, job opportunities may be shared only among very close relations (Magruder 2008, Seekings and Nattrass 2005), leaving individuals with poor social connections with very limited opportunities to find work.

[2]It is not common for household surveys in developing countries to ask respondents about the size of the firm they work for. Fortunately, the Brazillian LSMS is an exception. Brazil represents a particularly good comparison for South Africa as a country with a broadly similar income level and similarly extreme level of inequality.

[3]This estimate is the average over complying and non-complying firms. The greatest contributor to this estimated cost is VAT regulations, though labor regulations are also important. Of course, small business respondents to this survey may overstate compliance.

businesses would be granted a blanket exemption from these bargaining council agreements within the year in his state of the union address(Mbeki 2005). However, under pressure from trade unions and employers organizations to the contrary, the government never enacted this blanket exemption (e.g. Cosatu Rejects 2005). The fact both that the government would consider a legal change to exempt small business and that it meant with strong opposition confirms the anecdotal and survey evidence that these regulations are enforceable.

The potential of labor regulations to affect employment has been extensively explored in economics; a summary of this literature through the late 1990s is available in Nickell and Layard (1999). Much of the recent literature (e.g. Bertrand and Kramarz 2002, Besley and Burgess 2004, Harrison and Scorce 2008) has adopted a difference-in-differences approach where a time series of data on the legislative environment in states is summarized by a before and after period. Difference in employment trends between "treatment" states which adopt a policy and "control" states which do not are then compared to get an estimate of the effect of the regulation on employment. A second approach is to utilize a spatial discontinuity (e.g. Holmes 1998; Dube, Lester, and Reich 2008), where neighboring counties or states are compared, under the assumption that geographically proximate counties share similar labor markets and incentives to form labor policy, but are differentially exposed. Many existant labor regulation studies use a variation which adopts some elements of each of these approaches (e.g. Card and Krueger 1994), so that changes in trends are compared across spatially proximate regions. The measure of each of these studies is how comparable of a control group can be developed without causing small sample problems; to determine which approach is best for South Africa will require a more careful description of the labor regulations to be studied.

3 Industrial Bargaining in South Africa

Unions in South Africa can bargain with employers in two primary ways. The 1995 Labor Relations Act codifies the right of employers to form employers organizations for their particular industry and region and bargain with unions centrally; the labor standards which result from this bargaining can then be applied to all employees working in the industry and region which the bargaining council presides over. That is, if employer organizations and unions decide to bargain centrally, than all employees – regardless of their union status – who work within that geographical region will work under the agreed-upon labor standards. Unions and employers may also choose to bargain unilaterally, resulting in plant level agreements(Bendix 2001). Both unilateral bargaining and centralized bargaining are observed in a wide variety of industries and areas in South Africa, so that different industries in the same location may be covered by different types of agreements, industries may be covered by unilateral agreements in some locations and centralized agreements in others, and industries in a particular location may be covered by centralized agreements in one year and not in another.

It is encoded in law that bargaining councils must be representative of firms and employee unions in their jurisdiction; however, the extent to which this law is enforced is unclear. The official wording is that councils must be "sufficiently" representative, leading to a great deal of magisterial power and contention (primarily from small employers) as to whether the agreements represent all interests (Bendix 2001). South Africa's political structure is that about 350 magisterial districts are organized into one of 52 District Councils; these in turn comprise 9 provinces. In principle, there is not a strict criteria over which groupings of magisterial districts can form a bargaining council; in practice, most bargaining councils represent collections of magisterial districts which map to political boundaries, either national, provincial, or at the District Council level. In the model below, I follow the empirical trend in presuming that other magisterial districts within the district council are the natural bargaining partners in determining whether to form a bargaining council

agreement.

Existing studies on the effects of arbitration on wages and unemployment in South Africa have imperfect information on the presence of bargaining council agreements and treat the endogeneity of union membership via industry and occupational fixed effects, which may be an imperfect control; these studies find that unions receive very high wage premia, particularly at the bottom of the income distribution (Schultz and Mwabu 1998), and that bargaining councils exhibit a smaller, though still present, wage premium (Butcher and Rouse 2001). However, since the right to bargain centrally is one which must be exercized voluntarily, we may be concerned that bargaining council agreements exist systematically in the industries, towns, and years in which local labor markets make them particularly profitable for the firms which push for the bargaining agreement.

Moll (1996) outlines a theoretical model illustrating how larger, more capital-intensive firms may prefer the centralized bargaining structure as it obliges smaller, more labor intensive firms to pay similar wages and reduces competition. More specifically, we may imagine that there are three types of firms: Large Unionized Firms, Large non-Unionized Firms, and Small Firms. In the absence of a bargaining council agreement, large unionized firms pay privately bargained wages $\left(w^U\right)$, while large non-unionized firms and small firms pay market wages (w^*). Under a bargaining council agreement, all would pay the same bargaining council wage $\left(w^{BC}\right)$; following Moll (1996) in presuming that $w^U > w^{BC} > w^*$, it is clear that operating costs decrease for large unionized firms and increase for small firms and large non-unionized firms in the presence of a bargaining council agreement. Thus, large unionized firms have a direct incentive to advocate for a bargaining council agreement, as their wages may go down, while small firms and large non-unionized firms have a direct incentive to advocate against such an agreement. As the relative supply curves for the three types of firms shift, equilibrium changes as well. If small firms have the lowest marginal products of labor (due to low capital stocks), we may imagine that their supply curve shifts in by the largest margin, resulting in an increase in the residual demand faced by the two types of large

firms. Thus, large unionized firms benefit from less competition from small firms and lower wages, large non-unionized firms benefit from less competition from small firms but suffer from higher wages, and small firms lose by the greatest margins. The degree of these benefits, and the degree to which small firms and large non-unionized firms are punished by the bargaining council agreement, are functions of local demand, local labor supply, production technologies at each firm size, and other local labor market characteristics, as the changes in the demand faced by each type of firm will depend on anything which influences local supply and demand curves. In a simple model presented in an appendix on the author's website, I show more formally that, in equilibrium, large unionized firms will increase employment in response to a bargaining council agreement, while large non-unionized firms and small firms will decrease employment.

The differing profit incentives that employers face, outlined above, are clear. Therefore, the presence of a bargaining council agreement will clearly be related to some aggregation of the private incentives of large firms[4]. However, unions could adopt a bargaining position which is more or less hostile to bargaining councils, so the decision to pursue centralized bargaining may be an optimization not only of firms but also of unions. Most of South Africa's unions are aggregated into three large alliances which are important political entities. The Congress of South African Trade Unions (COSATU), the largest union alliance in South Africa, officially supports centralized bargaining council agreements for their effects on transforming labor standards, though both COSATU and NACTU, the National Council of Trade Unions (another large union alliance), have a history of also bargaining at the plant level (Bendix 2001). Both COSATU's official positions on bargaining council agreements and the discussion of commentators (e.g. Bendix 2001) suggest that unions have some support for these agreements due to the greater political support they receive from advocating for globally higher labor standards. We may also imagine that unions have varying incentives related to local labor market heterogeneity, for example, the amount of dues which can be

[4]This part makes the assumption that only large firms (and not small firms) initiate centralized bargaining.

received or local competition from uncovered workers. Here, I simply model the union's role in bargaining as a cost C of adopting the bargaining council agreement, though empirical analysis will be robust to heterogeneity in local labor markets and, in some specifications, the political value of a given town in a given year.

Suppose that, in the absence of a bargaining council agreement, large unionized firms in town t earn profits π_t^U, and that large non-unionized firms earn π_t^*. Further suppose that all large firms each earn profits π_t^{BC} in the presence of a bargaining council agreement before paying cost C to the union, and that fraction λ_t of the total Q_t large firms in town t are unionized. A bargaining council agreement is a collective result of the preferences of large firms throughout a district council, thus, if town t belongs to district council DC, bargaining council legislation is adopted if

$$\sum_{t \in DC} Q_t \pi_t^{BC} - C > \sum_{t \in DC} Q_t \left(\lambda_t \pi_t^U + (1 - \lambda_t) \pi_t^* \right) \tag{1}$$

Local labor demand, local labor supply, local production technologies, local unionization rates, and local product demand all determine the result of this relationship. We should expect to see bargaining councils in places where the majority of large firms would benefit from having a bargaining council – that is, places where prices would either increase a lot from a reduction in small firm production, where small firm production would decrease substantially from the presence of a bargaining council agreement, or where a large reduction in wages for unionized large firms may take place. In places where small firm production technologies are relatively inefficient, and large firms face little competition, the incentives to form a bargaining council agreement are weakened, while in places with a vibrant small firms sector, the incentives to enforce uniform wages may be high. Any econometric investigation into the effect of bargaining councils on employment and small firm employment would have to take these factors into account[5].

[5] This suggests, for example, that the effect of treatment on the treated is likely different from the effect of treatment on the untreatment, causing concerns over difference in difference estimates which do not control for spatial heterogeneity.

The focus of this paper will be on estimating employment effects, small firm employment effects, and wage effects of bargaining councils. That is, I want to estimate

$$Y_{ity} = \alpha + \beta_1 BC_{ity} + \Gamma X_{ity} + \nu_{ity} + \xi_i + \delta_y + u_{ity} \tag{2}$$

where Y_{ity} may be employment, or employment by firm size in industry i in town t during year y, or it may be some measure of wages in that industry-town-year; BC_{ity} denotes the presence of a bargaining council agreement, and X_{ity} are covariates including population. The model above suggests that the presence of a bargaining council agreement is related to many characteristics of local labor markets, including labor supply, small firm production technologies, etc. These are captured in ν_{ity}, which is unobserved and may be correlated with BC_{ity} and other explanatory variables. As such, the presence of a bargaining council agreement is an endogenous variable, and simple OLS estimates of this equation are likely to be biased.

To solve this endogeneity problem, I note that labor supply, production technologies, and other characteristics of local labor markets within a given industry are likely to be spatially continuous, so long as migration and trade are locally feasible. Formally, let $R(t)$ denote the set of all towns within radius R of town t, Z_{ity} be the vector $[X_{ity}; BC_{ity}, \xi_i, \delta_y]$ and $Z_{iR(t)y}$ denote the vector of $Z_{it'y}$, $\forall t' \in R(t)$. Then, spatial continuity suggests that $E\left[\nu_{ity}|Z_{iR(t)y}\right] = E\left[\nu_{it'y}|Z_{iR(t)y}\right]$ for R sufficiently small and $t' \in R(t)$. This assumption is identical to that made in standard regression discontinuity designs, which motivates the empirical strategy below.

Potential deviations from spatial continuity, and adjustments which should be robust to them, are discussed in section 4.1, below.

4 Spatial Regression Discontinuity

Spatial discontinuity analysis in economics has largely followed Card and Krueger (1994) who use Eastern Pennsylvanians as a control group for New Jersey residents. Card and Krueger estimate the effect of a change in minimum wage laws through a difference in difference – under the assumption that Eastern Pennsylvanians are similar to New Jerseyans, changes in employment rates of New Jerseyans should be similar to that of Eastern Pennsylvanians, except because of the differential changes in minimum wage laws. Other papers have followed this approach in restricting analysis to individuals who live near the border of states or countries which enact particular laws; for example, Holmes (1998) compares manufacturing production in counties just north of "right-to-work" states to those just south of the border, finding large effects which quickly diminish with distance from the border.

The approach of restricting analyses to border regions is appealing as the spatial discontinuity is transparent – towns on one side of the border are presumably similar to those on the opposite side. However, there are several limitations to analysis of this sort. First, restricting samples to border areas in practice reduces space from having two dimensions to having just a single one. A border analysis not only compares individuals, towns, or counties to proximate ones, it compares all towns on one side of the boundary to all towns on the other. If the border is long, this may in effect cause some pairs of towns which are geographically distant to represent a control-treatment difference. These towns need not be proximate, and while using border-region fixed effects can eliminate some of this heterogeneity, it remains an imperfect approach as it introduces a discontinuity into continuous space[6].

Secondly, we throw away a lot of information when we restrict the sample to border towns. Restricting the sample to border regions is analogous to cutting the sample to observations

[6] A different approach which solves this concern is presented in Dube, Lester, and Reich (2008). That study restricts the sample to border regions and uses contiguous county-pair fixed effects, a similar differencing approach to that used below. However,as that study also restricts the sample to border counties, it loses the capacity of non-border regions to improve the precision of covariate estimates, as discussed below.

within a small bandwidth of the discontinuity in the running variable in a more standard regression discontinuity. However, just as in a more standard RD, this eliminates the potential of more distant observations to inform on the relationships between other covariates and the dependent variable. If we are studying employment, the most important of these is likely to be population, as larger cities mechanically employ more people. Observations which are distant from the border can help us identify the relationship between population and employment, while allowing the effect of space to remain fairly unparametized. This is analogous to allowing the relationship between the running variable and the dependent variable to be flexible away from the eligibility cutoff, but still using the observations which are more distant from the cutoff to learn about covariates.

This paper adopts the spatial fixed effects suggested in Conley and Udry (2008) and Goldstein and Udry (2008). In both of these papers, these spatial fixed effects are used to control for unobserved soil quality variation which is presumed to be similar amongst nearby plots. The idea of this approach is identical to the standard fixed effects approach. For each observation, we can subtract off the mean of observations which are spatially proximate. Using spatial fixed effects, we have

$$\begin{aligned} Y_{ity} - \frac{1}{n_{R(t)}} \sum_{t' \in R(t)} Y_{it'y} &= \beta_1 \left(BC_{ity} - \frac{1}{n_{R(t)}} \sum_{t' \in R(t)} BC_{it'y} \right) \\ &+ \Gamma \left(X_{ity} - \frac{1}{n_{R(t)}} \sum_{t' \in R(t)} X_{it'y} \right) + \nu_{ity} - \frac{1}{n_{R(t)}} \sum_{t' \in R(t)} \nu_{it'y} + \tilde{\delta}_t + \tilde{u}_{ity} \end{aligned} \tag{3}$$

where $n_{R(t)}$ represents the number of towns in $R(t)$ and $\tilde{\delta}_t$ and $\tilde{\varepsilon}_{it}$ similarly represent spatial deviations. By assumption, $E\left[\nu_{it'y} | Z_{iR(t)y}\right] = E\left[\nu_{ity} | Z_{iR(t)y}\right]$ for $t' \in R(t)$. That is, this differencing removes the endogeneity contained in the ν_{ity}. We presumed earlier that u_{ity} was unrelated to BC_{ity} and X_{ity} (as endogeneity was summarized by ν_{ity}). If we extend that so that u_{ity} is strictly exogenous, i.e. $u_{ity} \perp Z_{iR(t)y}$, then this within estimator will consistently estimate β_1 and Γ.

This identification is valid under the assumption that the incentive to form bargaining councils, modeled above as a collection of characteristics of local labor markets, is spatially continuous. Identification is by spatial discontinuity: outcomes are compared only against those of proximate neighbors as are program status and covariates. This equation estimates whether, if your bargaining council status is greater than your neighbors' (i.e. you live on the bargaining council-side of a border), you have more employment than your neighbors. For interior magisterial districts, the spatial deviation is zero, but the employment effects of differences in population and time trends can still be estimated using the interior variation[7]. The anologous approach in more standard regression discontinuity is to examine differences in the dependent variable that exist only within a small bandwidth of the eligibility threshold, while allowing the eligibility-determining variable to have a non-parametric effect more generally.

Finally, while theory may tell us that local labor markets should be spatially continuous, it offers few predictions as to the rate at which spatial similarity dissipates. Fortunately, spatial fixed effects allow an easy test to determine whether estimated effects are robust to a wide variety of assumptions on average spatial heterogeneity. In particular, rather than equally weighting all observations within radius R to determine spatial means, we can allow a weighted spatial mean, where observations within different radii are weighted differently. Then, we can examine simultaneously which choice of weights would lead to similar point estimates and statistical tests, as well as examining which choice of weights best describes the spatial heterogeneity in the data in a sum of squared deviations sense.

[7]Dube, Lester, and Reich (2008) argue that time trends may be different in interior and border zones. A slight modification of the spatial fixed effects used here can accomodate that hypothesis: specifically, we can take difference off means of town-industry-year observations only for towns within radius $R(t)$ and which are in the same year. This approach yields qualitatively similar results, though they are much noisier. For small firm and self-employment effects, both the coefficient and standard errors blow up at about the same rate, so statistical precision is maintained (and point estimates are large; with a BC associated with a 30% reduction in small firm employment). However, standard errors are quite large when we discard altogether the intertemporal variation, which suggests that such an approach may not have the power for sensible inference in these data, and gives preference to the fixed effects employed here. This preference seems especially appealing given the consistency of the results when town-year fixed effects are added to spatial fixed effects; these allow time trends which affect all industries to vary non-parametrically across space.

In the spatial heterogeneity robustness section below, I describe these weighted spatial fixed effects (WSFE), and illustrate robustness to a wide variety of spatial weighting schemes.

4.1 Is Spatial Discontinuity Sufficient?

Several concerns about the spatial discontinuity identification in this context demand further consideration. First, as mentioned above, Bargaining Council Agreements usually apply to all magisterial districts which belong to a larger political entity, either a District Council, Province, or the entire nation. However, in a few cases individual magisterial districts are added or subtracted from these groups in the coverage of a bargaining council (usually either the biggest town or closest neighbors of an adjoining district council). Though these observations represent a small share of the data, we may still worry about the implications of these observations for analysis. This is somewhat a problem for spatial discontinuity – if bargaining councils are choosing precisely where the boundary of coverage should stop, we may worry that the industry-specific labor markets are indeed discontinuous at those coverage differences in an important way. In fact, this problem is directly analogous to the problem of imperfect enforcement and incomplete take-up in more traditional regression discontinuity. Traditional regression discontinuity avoids the problem of endogenous take-up by instrumenting program receipt with program eligibility. I follow this approach in this paper: I describe a magisterial district-industry-year observation as eligible for the program if it belongs to a district council where at least one magisterial district has a bargaining council agreement in that industry-year and use that measure of eligibility as an instrument for program receipt. All first stages are extremely strong; the minimum t-statistic of bargaining council eligibility on bargaining council status across specifications and samples is 9.9.

Second, spatial fixed effects may imperfectly control for characteristics of local labor markets. This may be true because great circle distances are imperfect measures of true economic distance. For example, travel times may be large between two physically proximate

towns due either to poor infrastructure or rugged terrain, or other labor regulations may vary discontinuously at political boundaries. Moreover, local labor markets may change over time in ways correlated with regulation (e.g. Dube, Lester, and Reich 2008). It is also possible that industry-specific labor markets relate to geography in different ways. To deal with these two concerns, I repeat all analysis both with town-year and town-industry fixed effects. Town-year fixed effects compare across industries within the same town and year, and in doing so control for any ways in which the town is different from its spatially proximate neighbors that year, including ways in which that town may be politically valuable to the union alliance, the possibility that local neighbors are in fact distinct labor markets due to geographical separation or legislation (presuming that all industries are similarly effected by the long travel times or other disruptions to spatial continuity), or the possibility of secular local time trends which effect all industries. Town-industry fixed effects examine how employment changes when bargaining council agreements are added or expire, and so control for any ways in which that industry's local labor market differs from its spatially proximate neighbors which is constant over time[8]. These represent very different identification assumptions, and the similarity of results under these two approaches is additional evidence that the true model is the hypothesized one.

However, we may remain concerned that labor markets for a particular industry may vary discontinuously and change endogenously over time with bargaining council agreements. As such, I propose the following: Equation 1 makes clear that the presence of a bargaining council is due to some collaboration of towns in the same political district. If local labor markets are relatively continuous, then nearby towns should have similar incentives to form a bargaining council and the spatial fixed effects approach solves the endogeneity problem.

[8]Some care is required combining the spatial fixed effects with other fixed effects when other dimensions of the panel are not balanced. In this paper, this happens when using the wage, tenure, and DC ratio subsamples. In this case, the standard within estimator is biased, and so demeaning sequentially along spatial and other dimensions is not consistent. Unfortunately, a simple adjustment such as that in Davis (2002) is not possible, as the spatial fixed effects cannot be represented as a projection onto the column space of a number of dummy variables. As a result, I conduct these estimates using the full set of dummy variables for the additional fixed effects whenever the panel is unbalanced.

If they aren't, then it indicates that something about industry i in town t is different from industry i in neighboring towns. If town t has much lower employment than other towns in its District Council, then, as equation 1 makes explicit, town $t's$ preferences should not be strongly reflected in the presence or absence of a bargaining agreement. In particular, if town t is discontinuously different from its political neighbors in its incentive to form a bargaining council, than it will not be able to enact its optimal choice. As such, our concern for endogeneity is minimized. However, if a dominate share of industry i is located in town t, then this concern may remain, and the presence of a bargaining council in town t's district council may be a reflection of discontinuous labor market trends in town t. I repeat all estimation with a sample of industries and towns where employment is no more than 20% of employment in that industry in that district council on average.

Finally, two potential sources of non-independence among observations are well-known and relevant to this context. A first challenge to evaluating programs which are implemented at aggregate levels is that if individuals in a political district have correlated error terms, or there is autocorrelation in the error, then OLS produces inconsistent standard errors (Bertrand et al 2004). The standard solution is to cluster at the policy group level. Since bargaining councils vary on the district council-industry level[9] (there are 208 district council-industry groups and 52 district councils in the estimation sample), this context avoids the small group number concerns which have challenged some past studies of governmental policies (e.g. Donald and Lang 2007). Secondly, it seems likely that the error term may be spatially autocorrelated (Conley 1999). Of course, the group of towns which are physically proximate is similar to the group which are politically proximate; however, as identification

[9]Several of the bargaining councils extend agreements to entire provinces, while others operate only on the District Council level (and a small fraction operate for even smaller units). This makes it difficult to know, for certain, how to categorize observations (particularly for industries and towns which are not covered by a bargaining council agreement). The results presented here presume that, since some District Councils unilaterally receive bargaining councils, this is the true observation level (implicitly, this assumes that bargaining councils which exist react to considerations at the District Council level). An alternate assumption would be to assume that the observation unit is the province-industry level. Results which make this assumption and cluster simultaneously at the spatial-industry, province-industry, and town level are similar and available from the author.

rests on the difference between these two groups, it is desirable to construct standard errors which are robust to correlation amongst both of them. This paper allows observations to be related if either they are close spatially or in the same district council. This is the more computationally intensive procedure outlined in Cameron, Gelbach, and Miller (2006), and also a special case of the Conley (1999) spatial errors if "economic distance" is defined as equal amongst individuals who live either within a given physical distance or a political demarcation.

5 Data

Data is drawn primarily from two sources. The South African Labour Force Surveys are a nationally representative rotating panel conducted twice yearly from 2000 through the present, each iteration surveying around 70,000 people. I use the September surveys from 2000-2003. Unfortunately, the panel aspect has not been well-maintained, with household identifiers not remaining consistent from wave to wave. As such, I aggregate data to the magisterial district level and use it as a panel at that level. These data are not intended to be representative at the magisterial district level and are not publicly released at that level to prevent mistaken inference (on, for example, the extent of the variation in employment in a particular town year to year). This concern, however, should not limit more robust econometric analysis, so long as the degree to which the data are not representative is unrelated to the variables of interest and local-level unobservable heterogeneity is properly controlled for. While magisterial district identifiers are not released, they can be inferred from personal identification codes. These identifiers remain unchanged since at least the 1997 October Household Survey, which published an association between number and local municipality names[10]. From this list, I determine the magisterial district of each sampling area, and determine the longitude and latitude for the population center of that magisterial

[10] Examining characteristics of magisterial districts between these two surveys reassures that the identifiers are in fact unchanged.

district. The unit of analysis in this paper will thus be the magisterial district; since sampling weights are not designed to be representative at this level I do not use them. Therefore, I measure employment in a given industry in a given town as the number of people surveyed in that town who work in that industry[11]. We may be concerned that very large towns have different labor markets from their neighbors, and that we get little useful information out of small towns where relatively few individuals were surveyed. I exclude the top and bottom two percent of towns in terms of population from the analysis. Summary statistics of the variables which will be used are included in table 3.

The presence of bargaining council agreements in a given year is revealed by the South African Government Gazette, which publishes all agreements. A database compiled by the author reveals which industries in which magisterial districts were covered by an agreement at each year. This yields the outcome that 12 two-digit industries in South Africa are covered by bargaining council agreements for at least some of the sample period. Of these, 7 industries show at least cross-sectional variation in their coverage across the district councils of South Africa. All in all, 22% of prime-age African and Coloured workers in South Africa work in two-digit industries where, in their district council, some workers are covered by a bargaining council agreement[12]. Different industries have different minimum effective scales, limiting the potential for entrepreneurship in some industries. Table 4 reveals that 82.5% of the prime-age African and Coloured self-employed in South Africa work in two-digit industries which at least sometimes have bargaining councils – this suggests that these councils are being utilized more in industries where small scale firms are economically viable. In contrast, only about 40% of workers overall are working in these industries. Looking within industries which at

[11]In a related point, it is not immediately obvious how to treat observations of 0 employment in some category in a particular town (of which there are many). On the one hand, these observations give useful and important variation – if bargaining councils are brutally effective, we may expect to see 0 small firm employees in a particular town-industry. On the other hand, when I (ultimately) take log+1 as a measure of employment, the log operator strongly emphasizes observations which are 0. This concern is lessened by the use of the simple count data rather than weighted counts – the difference between log(1) and log(301) is a lot more than the difference between log(0) and log(1). I also present results using the fraction of the population who are employed, which does not weight zeros as strongly.

[12]The actual number of covered workers is probably lower, due to the aggregation at the 2 digit level. Aggregation challenges are addressed below.

least sometimes have a bargaining council, we see an even more interesting result. 53% of employees who work in one of these industries are covered by a bargaining council agreement. However, only 36% of self employed and 40% of small firm employees do, in contrast with 70% of large firm employees – that is, among industries which at least sometimes have bargaining council agreements, places with bargaining council agreements have limited small scale and self employment. The industries and the percentage covered are listed in table 5.

Industries with bargaining councils are aggregated to the two-digit level, while bargaining councils are often finer in their definition, so that my measure of coverage includes individuals who are actually working in distinct, uncovered industries as well as covered employees. In principle, this bias should result in conservative estimates due to measurement error, since the bargaining council agreements only cover a fraction of the workers in the two-digit industry. Two of the industries with variation end up in "other" categories. We might worry that these categories are more heterogeneous than other two-digit designations, and that the bargaining councils represented (hairdressing, laundry services, and contract cleaning) represent a smaller fraction of the workers in the "other services" and "other business activities" industries. Additionally, a third industry (electrical manufacturing) is very small in scale (with only 25 small firms employees measured in South Africa across the 4 survey years considered here), and covered almost everywhere. I exclude these three industries in the analysis below, although similar analysis including these industries is available from the author.

6 Results

I present first results for which spatial heterogeneity has not been corrected for comparability with earlier studies. Specifically, I estimate

$$Y_{ity} = \alpha + \beta BC_{ity} + \Gamma X_{ity} + \xi_i + \delta_y + \varepsilon_{ity}$$

where ξ_i and δ_y represent industry and year level fixed effects, respectively, and X_{ity} includes a quartic in log population. Y_{ity} variables include employment, large firm employment, self employment, and small firm employment (small firms are defined to be firms with fewer than 10 employees while large firms have more than 20), both measured as log(X+1) as well as by the ratio of each of variable to the population. I perform two types of unadjusted estimations here: simple OLS and a difference in difference specification, which conditions on District Council-Industry fixed effects; results are reported in table 6. I present results for both the full sample and the sample of observations who are no more than 20% of that industry's employment in their district council. All cells report the coefficient on bargaining council presence for a given sample and dependent variable. Across the board, the pooled OLS reveals that bargaining council agreements exist in towns that, if anything, have more employment in that industry. Even columns report the difference-in-differences analysis, which conditions on district council-industry fixed effects (as the policy variable is determined at the district council level). Here, effects are negative and significant in both samples, revealing that, as bargaining council agreements come into place, employment decreases by 7 %, as does small firm employment. Self employment appears to decrease by about 5%. Recall that our concerns over endogeneity were smaller for the town-industries which are too small to play a strong role in bargaining council status. A difference in difference over these towns (presented in column 4), reveals that employment and small firm employment are declining by 12-14%. However, we may be concerned that bargaining council agreements are being adopted or eliminated endogenously in places with specific time trends in employment or small firm employment. In the analysis that follows, I utilize spatial fixed effects to control for this endogeneity and achieve identification through spatial discontinuity.

6.1 Employment Results

Table 7 reports the coefficients on the presence of a bargaining council agreement on employment from several spatially differenced estimations, where the estimation equation is the

instrumental variables analogue of equation 3, and spatial deviations in bargaining council status are instrumented by spatial deviations in bargaining council eligibility. In all equations, the spatial fixed effect is taken at the 30-mile radius, so that each dependent and independent variable represents deviations of variables between the observation of interest and other observations in the same industry and within 30 miles, where distance is determined by the great circle method. All estimations are conditional on a quartic in log population and time fixed effects, and all errors are clustered among observations across all years of the same industry within 2 degrees of latitude or longitude, as well as among all industries, towns, and years in the same district council. The first two columns report the effects of a bargaining council agreement on log employment. Having a bargaining council agreement is not strongly associated with changes in log employment in the first row; once town-specific, town-year, or town-industry fixed effects are included, the estimate of the bargaining council effect becomes about 8-11%, and significantly different from zero. These coefficients are quite stable despite the very different identification assumptions: whether we look across industries at spatial deviations in employment, or across time within industries, we draw very similar inferences about the effect of bargaining councils. In general, the specifications in column three, which takes as the dependent variable the percentage of the population employed in a particular industry, are more precisely estimated,with bargaining councils reducing the fraction of the population employed in a particular two digit industry by 2-4 tenths of a percentage point.

As discussed in the identification section, we may worry about the identification assumption of spatial continuity for town-industry observations which are extremely important within the district council, as these towns have the capacity to uniquely determine bargaining council status. However, in towns which represent a relatively small share of that industry's District Council employment, even if their incentives to form a bargaining council are radically different from their neighbors, their wishes for bargaining council status are unlikely to be implemented. For this group, there are even greater incentives to believe in

the validity of the spatial discontinuity approach, and the even columns of table 7 implement this procedure using only these observations. Here, despite the smaller sample, precision increases and point estimates rise. Among town-industry groups which are too small to independently effect bargaining council policy, we see employment fall by 11-14% relative to neighbors and other industries or other years within the town[13].

6.2 Wage Results

Of course, the stated purpose of the bargaining council legislation is to improve working conditions rather then reduce employment. We can also ask if wages increase with bargaining council agreements. This analysis uses the subsample with at least one wage observation, which eliminates zero employment towns (and some with non-response to the wage question). One consequence of the smaller sample is that the 30 mile radius, in conjunction with various town-specific heterogeneity loses a lot of power; column 1 of table 8 indicates that we find a 10-21% effect on wages at this radius, though standard errors become large and the effect loses statistical signficance as we consider town-year or town-industry fixed effects. Column 2 repeats the analysis with a wider 50 mile spatial radius for the spatial fixed effects; at this larger radius the town-year effects regain precision. Overall, industries represented by a bargaining council in a town have 21% higher mean wages than the same industry in neighboring towns, and 13% higher wages if we hold constant mean deviations across industries in that year. Since the wage data appear not to be sufficiently dense for a 30 mile radius with town-year heterogeneity, I report the following wage and worker characteristic regressions using the 50-mile spatial fixed effects (30 mile radii give similar, but sometimes less precise, point estimates and are available from the author). The motivation above suggested that small firms should see larger wage increases than large firms, as large firms often must pay union wages anyway. We can examine mean log wages for small firms (with fewer than 10 employees) and large firms (with more than 20) separately, in columns 3 and

[13]It is worth noting that, for this group of towns, the standard difference-in-differences yields very similar point estimates. We may take that as supportive evidence that endogeneity is less important for this group.

4. Consistent with theory, wages in small firms are rising substantially, with (precisely measured) point estimates around 12-18%. In contrast, large firm wages are if anything decreasing in response to bargaining councils, consistent with the hypothesis that bargaining council wages are lower than privately bargained ones (though errors are too large to reject a null hypothesis of a zero effect). However, caution must be taken in interpreting wage estimates as a change in wages for individual workers, because the composition of employees is changing.

Table 9 tests the importance of worker composition effects. Column 1 reports wage estimates when we control for the fraction male, the average number of years of primary and secondary education, and a quadratic in average potential experience (age - education - six). We see that controlling for these observable characteristics attenuates the effect of bargaining councils on wages somewhat, with estimated effects dropping by 4 to 5 percentage points in each specification. Columns 2 through 4 looks at how each of these variables changes with bargaining council status, and we observe that the big difference is in the gender of employees. When a bargaining council is present in an industry, the fraction of the labor force which is male increases by 5 to 13 percentage points. Education and age of the labor force are not robustly associated with bargaining council status. The education result is consistent with other studies (e.g. Magruder 2008) which find that education is not a strong predictor of employment in South Africa. However, the age result is on surface somewhat surprising, as labor standards in these bargaining council agreements include hiring and firing regulations as well as wage standards. In a high unemployment context, age is a poor proxy for tenure, which we might expect to increase in the presence of hiring and firing restrictions. We can directly investigate the effect of bargaining council agreements on tenure; this requires using the sub-sample which responded to the tenure question. In column 5, I report the effect of bargaining councils on mean log tenure at the plant. Mean tenures are increasing by about 18-19% in response to bargaining council regulations[14].

[14]In the wage sample, we do not find precise estimates comparing within a town-industry over time, while in the tenure sample, we do not find precise estimates comparing within a town-year across industries. This

6.3 Firm Size Results

Here, I divide firms into four groups: large firms, with at least 20 employees; small firms, with fewer than 10 employees; self-employment, and single-worker firms. Many self-employed individuals thus are also represented in the small firms and the single-worker firms categories. From the model above, we expect the bargaining councils to have the largest effects on employment in small firms. Large firm employment should have, in principal, an ambiguous effect, and the effect on self employment will depend on how many entrepreneurs run larger small firms and the enforcement capacity of the bargaining council. If most single-employee firms aspire to grow to multiple employees, or if single employees are themselves paid a wage, it may be that bargaining council legislation reduces employment in single-employee firms. However, since most single employee firms are owner-operated, it seems likely that single-employee firms are primarily impacted through these dynamic incentives, and we may well anticipate that these incentives are weaker than the direct wage effects of the agreements. Therefore, we may anticipate smaller effects among single-employee firms. I again estimate the bargaining council effect in two ways, using the log counts of individuals in this group and the frequency of this group among the general population.

Table 10 reports the result of this analysis for each of these dependent variables, where columns represent different fixed effects specifications (again, in addition to the spatial fixed effects). Here, effects for small firms and self-employment are larger and consistently reach standard statistical significance threshholds. Consistent with theory, bargaining councils reduce small firm employment substantially, with bargaining council employment being associated with a 7-15% decline (also estimated as a .2-.38 percentage point change in the fraction of the adult population working in small firms in that industry). This effect remains very similar when we examine how spatial differences vary within industries in a town or

may be a case of sticky wages, or constant turnover across industries within a labor market; however, it may also be a case of low power in these estimations, and we cannot rule out similarly-sized effects. Non-response in the tenure variable is not closely associated with non-response in the wage variable, and so the sample which has both of these is further reduced; examining the effects of bargaining councils on tenure and wages in this sub-sample produces similar, though sometimes noisier, estimates.

town year and within an industry over time. Self employment similarly declines by 7-15%. Large firm employment, in contrast, does not report a consistent effect. Coefficients are never significant and are always smaller than small firm employment estimates. Similarly, single-employee employment is not consistently related to bargaining council agreement status, though in one specification bargaining council agreements are negatively related to single employee employment rates in a precise way (though, again, even this estimate suggests a smaller effect than the analogous estimates for small firm employment and self employment). This suggests that single-employee firms are not driving the large effects on small firms, and that these bargaining councils are most effective against small firms, as suggested by theory.

Recalling the robustness arguments associated with the group of town-industries which represent, on average, small fractions of their industry's county-level employment, I repeat table 10 using only these observations. These results are presented in table 11. Consistent with the idea that endogeneity is minimized in this subsample, results here line up precisely with theory, with the largest effects being on small firms, smaller and marginally significant effects on self employment, and consistently small and insignificant estimates on large firm and single-employee firm employment. An industry in a town which represents a small fraction of it's county's employment can expect to see a 10-12% decline in small firm employment, a 5-8% decline in self employment, and no change in it's large firm or single-employee firm employment relative to it's neighbors, and relative also to the variety of potential unobserved components at town, town-year, or town-industry levels. Similarly, looking at rates, .31-.41 percentage points of the population are not working in a small firm in a given industry as a result of a bargaining council agreement.

7 Robustness: Border jumping

The spatial fixed effects approach employed here has the limitation common to regression discontinuity studies that only local trends at the point of discontinuity are identified. In

the case of firm employment, we may be concerned that firms could relocate to a town immediately on the opposing side of the border. That is, two possible regimes would result in similar analysis: in the first, all towns with bargaining council agreements could be employing fewer people than their potential. This would indicate that, if we could control for all variables which determine employment except for the presence of a bargaining council agreement, the residual employment would appear as a spatial plateau. An alternate regime would be one where locally at the border, employment is depressed on the side with a bargaining council agreement and increased on the side without one, where one town's disadvantage is another's advantage. While both of these regimes indicate that firms prefer to operate outside of the bargaining council restrictions, the former is clearly a more important issue for policy. Figure 1[15] presents a graphical illustration of these two regimes. Here, locations on the positive side of the x-axis are presumed to have the bargaining council in place, while those on the negative side do not. If firms are merely jumping a border at x=0, then the bargaining council's effect on employment may look something like the dashed line, while if employment is being eliminated, the bargaining council's impact may more closely be represented by the solid line. Both approaches would yield similar spatial-fixed-effects estimates.

A direct test of the hypothesis of border jumping (similar to Holmes 1998) can be found by reexamining level effects, and asking whether log employment is different in a magisterial district if it is on the border of a bargaining council agreement than otherwise. That is, we can define $Border^{+}_{k,ity}$ to be equal to 1 if an industry-town-year observation is covered by a bargaining council agreement but is within k miles of another town in that industry and year which is uncovered by a bargaining council agreement and 0 otherwise. We can similarly define $Border^{-}_{k,ity}$ equal to 1 if the observation is uncovered but within k miles of a town which is covered by a bargaining council agreement in that industry and year. Then, we can determine if border regions are different from their counterparts in the same bargaining

[15] This figure is similar to those presented in Holmes (1998).

council regime by regressing

$$Y_{ity} = \sum_k \beta_{1k} Border^{+}_{k,ity} + \sum_k \beta_{2k} Border^{-}_{k,ity} + BC_{ity} + \gamma_t + \alpha_i + \varepsilon_{ity}$$

However, border regions may differ from interior regions for many reasons which may lead to faulty inference. Fortunately, we can use two controls in this set-up to control for spatial heterogeneity. First, the magisterial district-industry fixed effects used earlier were collinear with the spatial fixed effects and can still be used in this setting. These fixed effects identify the coefficients of interest off of only time variation, so they control for any local labor market characteristics which remain constant over time. Here, we can identify any border jumping successfully presuming only that changes in local labor market characteristics are summarized by changes in bargaining council status, and not changes in border status. Moreover, this analysis allows us to look simultaneously at border effects and at the effect of bargaining councils on average. Second, we could consider fixed effects at the District Council-Industry-Year level. As this is the level at which bargaining council decisions are made, anything about the local labor market which led to a bargaining council agreement existing or not in that year is controlled for, and border effects will be identified off of magisterial districts within district councils which are closer to the border than other magisterial districts within the same district council.

Table 12 reports the results of this analysis. We see several results in this analysis. First, there does appear to be border jumping. When a bargaining council is formed near a given town but not including that town, that town sees a large increase in employment (column 1) and small firm employment (column 5). We similarly observe border jumping for large and small firms using only the spatial variation, which reveal that having a bargaining council in your district council-industry-year but being closer to the edge of the bargaining council regime is associated with some flight of large firms (Column 4, row 1), and that not having a bargaining council, but being near towns that do, is associated with an increase in

small firms (column 6, row 4) However, this fact is unrelated to the bargaining council effect documented in this paper, as coefficients on bargaining council status are virtually unchanged by controlling for border status (columns 1,3, and 5) for overall employment, large firms, and small firms. This is in part becuase border regions actually have more employment on the bargaining council side as well as the non-bargaining council side, and in part because there are some fairly complicated spatial dynamics as you examine the effect of 30-50 miles out from the border as opposed to being within 30 miles of it. Regardless, we can conclude two things. First, border jumping is taking place, suggesting that firms do prefer to resettle outside of the bargaining council regime and offering supporting evidence that firms (and especially small firms) prefer to avoid bargaining council agreements. Second, this effect is not inflating our estimates of the employment implications of bargaining councils.

8 Robustness: Average Spatial Heterogeneity

The spatial fixed effect estimator compares employment in a town to the average across all towns within some radius of it. However, while introspection may provide some guidance as to the proper choice of radius, any choice will remain somewhat ad hoc. Ultimately, any spatial discontinuity study will revolve around an assumption similar to the one set forth here, where endogeneity, represented by v_{ity} in equation 2, is assumed to be constant. Finite data sets do require that spatial bandwidths are not arbitrarily small, and researchers face a tradeoff between greater power through more observations per fixed effect, and greater specificity through increasingly local fixed effects. This means that, at the very least, a sensitivity analysis is desirable to see whether the identified results are sensitive to different assumptions on the relevant spatial bandwidth.

A pragmatic approach to identifying spatial discontinuities, then, would ask how estimates change as we change the weight which we put on observations which are at different distances in controlling for local heterogeneity. To motivate this formally, consider sev-

eral sets $R_g(t)$, each of which contains towns within some radius R_g of town t. Further suppose that $\nu_{ity} = \sum_{g=1}^{G} \nu_{ity}^g$, that is, that the endogeneity takes the form of several spatial processes which are similar in expectation at different local radii, and where I assume $E\left[\nu_{ity}^g | Z_{iR_G(t)y}\right] = E\left[\nu_{it'y}^g | Z_{iR_G(t)y}\right] = \rho_g E\left[v_{it'y} | Z_{iR_G(t)y}\right]$ for all $t' \in R_g(t)$ where $R_G(t)$ represents the largest radius under consideration. Each component ν_{ity}^g, then, will represent the part of the endogeneity which is similar among observations at radius g, and ρ_g is the relative weight of ν_{ity}^g in ν_{ity}. The conditional mean assumption suggests that while the overall disturbance ν_{ity} may be related to estimation parameters, the fraction of the disturbance which is constant over a given radius is unrelated to the estimation parameters[16]. This suggests that

$$E\left[\nu_{ity}^g | Z_{iR_G(t)y}\right] = E\left[v_{it'y}^g | Z_{iR_G(t)y}\right] = \rho_g E\left[v_{it'y} | Z_{iR_G(t)y}\right] \text{ for } t' \in R_g(t) \text{ and } g = 1, ..., G$$

We are then left with

$$E\left[\nu_{ity} | Z_{iR_G(t)y}\right] = \sum_{g=1}^{G} \rho_g E\left[\nu_{it_g y} | Z_{iR_g(t)y}\right] \text{ for } t_g \in R_g(t)$$

if we knew the ρ_g, we could simply difference off a weighted mean

$$\begin{aligned} Y_{ity} - \sum_g \frac{1}{\sum_k \rho_k n_{R_k(t)}} \sum_{t' \in R_g(t)} Y_{it'y} &= \beta Z_{ity} - \sum_g \frac{\rho_g}{n_{R_g(t)}} \sum_{t' \in R_g(t)} Z_{it'y} \\ &\quad + \nu_{ity} - \sum_g \frac{\rho_g}{n_{R_g(t)}} \sum_{t' \in R_g(t)} \nu_{it'y} + \tilde{\delta}_y + \tilde{u}_{ity} \end{aligned}$$

Once again, the conditional expectation of the ν_{ity} is equal to the properly weighted conditional expectations of spatial heterogeneity in nearby observations, allowing consistent identification of β. Of course, we don't know the relative weights. However, for robustness analysis, we can choose a few rings, and solve this equation for a relatively fine grid of

[16] This assumption seems logical if we view the ν_{ity} as changing at some mechanical rate as we move to larger and larger radii. The ρ_g, then, reflect the extent to which endogeneity remains similar across observations of different distances.

all possible weights over those rings. Then, we can infer whether coefficient estimates and stastical inference would be robust to a wide variety of assumptions on the spatial heterogeneity. An alternate approach would treat the ρ_g as parameters to be estimated, for example selecting the ρ_g which minimize the sum of squared error terms. In practice, putting the full weight on the most local rings always minimizes the sum of squared errors in this exercise. In the limit, this collapses to the town-industry fixed effects employed earlier, suggesting that we may prefer these estimates. In this paper, I take the robustness approach as the comparability between spatial and intertemporal estimates has been a strength of the analysis, and I illustrate that the effects highlighted here are robust to a wide variety of potential spatial weights.

8.1 Weighted Spatial Fixed Effects: Results

In what follows, I assume there are three different rings to the heterogeneity, relevant at 5, 30, and 50 miles from the source observation. I present figures depicting the coefficient estimate, the t-statistics, and the sum of squared residuals at every possible .05 weight for each of these rings for several specifications on employment and small firm employment. I further restrict all weights to be weakly positive, and $\rho_5 + \rho_{30} + \rho_{50} = 1$. In each picture, the vertical (Y) axis represents the dependent variable, while the Z-axis represents the relative weight (out of 20) on the 5-mile fixed effect, and the X-axis represents the relative weight on the 30-mile fixed effect. Therefore, the evaluation at (0,0) represents the coefficient estimate, t-statistic, or sum-of squared residuals resulting from a 50-mile spatial fixed effect, the evaluation at (20,0) represent the evaluation with a 5-mile fixed effect, the point (0,20) represents the evaluation with a 30-mile fixed effect, and interior points feature weighted fixed effects. Figures 2-4 represent these evaluations for log employment after controlling for a quartic in log population and time fixed effects, while figures 5-7 also estimate employment results adding in town-year fixed effects. Figures 2-4 show fairly stable point estimates and standard errors, and, just as in table 7, we see that at 30 or 50 mile fixed effects employment

effects do not reach conventional significance levels. As we approach more local fixed effects, point estimates become bigger and significant, just as they did with a town-industry fixed effect in table 7. When we add town-year fixed effects, we see that point estimates are stable, large, and precise, regardless of the spatial bandwidth. In both cases, the data prefers very local fixed effects, in a sum of squared residuals sense (figures 4 and 7), suggesting that our preferred estimates may be the town-industry fixed effects estimation used earlier. We can perform a similar analysis with small firm employment. Results for small firms only including the population and time controls are presented in figures 8-10, while results also including town-year fixed effects are presented in figures 11-13. For small firms, we see stable and significant point estimates both with and without the town-year fixed effects. Once again, the data prefer extremely local fixed effects, again lending credence to the town-industry analysis. Regardless, it appears that small firm effects are quite robust to different forms of spatial heterogeneity, and that employment is as well if we control for town-year unobservables.

9 Conclusions

Bargaining council agreements are the outcome of a complex bargaining process. Their location is related to a variety of local labor market characteristics, which may lead to biases in OLS and difference-in-differences specifications. Under the assumption that labor markets are spatially continuous, this paper argues that spatial fixed effects represent an improvement over traditional spatial regression discontinuity designs and determines that bargaining councils are associated with about 8-14% lower employment in a particular industry, 10-21% higher wages, and 7-15% less employment in small firms. We can control for potential heterogeneity by removing variation common to towns, town-years, or town-industries; whichever of these specifications are used result in similar (and similarly precise) point estimates. That is, an industry with a bargaining council has about 8-14% less employment than it's neigh-

bors without a bargaining council. This is true if we compare it to how different industries in the same town compare to their neighbors, or if we compare how employment in that town and industry changes over time with bargaining council status. Industry-town observations which employ a relatively small fraction of the employees in their District Council experience the most severe consequences; that is, towns whose voices should receive little weight in the decision to form a bargaining council are the most severely impacted by its existence. The identification assumptions of spatial continuity can be weaker for these towns – if they differ substantially from their neighbors in their incentives to form bargaining councils they will be unable to implement their desired bargaining council status and so these estimates are particularly compelling. Moreover, while both small and large firms appear also to prefer avoiding these restrictions, and hence resettle on the opposite side of the border, this effect is unrelated to the estimated employment effect of bargaining councils.

Eight to fourteen percent is a large decrease in employment in a given industry. By means of comparison, Bertrand and Kramarz (2002) estimate that French restrictions on new retail firms led to a 3% decrease in retail employment, Besley and Burgess (2004) estimate that labor regulation reduced manufacturing employment in India by 7%, and Harrison and Scorce (2008) find that a 50% increase in the Indonesian minimum wage is associated with a 6% employment reduction. The bottom end of the point estimates, then, is as large as these effects of labor regulation found in other contexts. However, bargaining councils cannot explain all of the unemployment problem in South Africa. 22% of employees work in two-digit industries in places with bargaining council coverage. If each of these two-digit industries were to increase employment by 6%, it would cause a 1.33 percentage point total increase in employment. Accepting the largest point estimates of 14%, the elimination of bargaining councils may result in a 3.08 percentage point increase. These effects are large and should be of interest to policy makers. However, the South African unemployment situation is severe enough that a 3 percentage point increase in employment would leave South Africa with a severe unemployment problem. So while the unemployment effects

of these policies are as big or bigger than other estimated labor regulation effects, other problems still contribute to such high unemployment in South Africa. Spatially continuous aspects of union behavior, labor market policies other than bargaining council agreements, and the other voluntary and structural stories which may lead to high unemployment levels may play an important role. Similarly, the larger small firm effects is 4-6 times larger than the difference-in-differences effect of French entry regulations (Bertrand and Kramarz 2002). This policy is thus having its intended effect of restricting small firm profitability, in a context where the small firms sector was already anemic. Once again, however, the small firms sector in South Africa is so minimal that this 7-15% increase in these industries would leave small firm employment substantially below global norms. Further research remains important to learn about the other potential contributors to this problem.

References

(2005): "Cosatu Rejects Plan to Bypass Centralized Bargaining," *The Star*, February 15, 6.

Banerjee, A., S. Galiani, J. Levinsohn, Z. McLaren, and I. Woolard (2008): "Why Has Unemployment Risen in the New South Africa?," mimeo, MIT.

Bendix, S. (2001): *Industrial Relations in South Africa*. Juta and Co., Ltd, 4th edn.

Bertrand, M., E. Duflo, and S. Mullainathan (2004): "How Much Should We Trust Difference-in-Differences Estimates?," *Quarterly Journal of Economics*, 119(1), 249–275.

Bertrand, M., and F. Kramarz (2002): "Does Entry Regulation Hinder Job Creation? Evidence From the French Retail Industry," *The Quarterly Journal of Economics*, 117(4), 1369–1413.

Bertrand, M., S. Mullainathan, and D. L. Miller (2003): "Public Policy and Extended Families: Evidence from Pensions in South Africa," *World Bank Economic Review*, 17(1), 27–50.

Besley, T., and R. Burgess (2004): "Can Labor Regulation Hinder Economic Performance? Evidence from India," *The Quarterly Journal of Economics*, 119(1), 91–134.

Butcher, K. F., and C. E. Rouse (2001): "Wage Effects of Unions and Industrial Councils in South Africa," World Bank Policy Research Working Paper No. 2520.

Cameron, A. C., J. B. Gelbach, and D. L. Miller (2006): "Robust Inference with Multi-Way Clustering," mimeo, University of California, Davis.

Card, D., and A. B. Krueger (1994): "Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania," *The American Economic Review*, 84(4), 772–793.

Charmes, J. (2000): "Informal Sector, Poverty and Gender: A Review of Empirical Evidence," http://www.wiego.org/papers/charmes3.pdf, Background Paper for the World Development Report 2001.

Conley, T. G. (1999): "GMM Estimation with Cross-Sectional Dependence," *Journal of Econometrics*, 92(1), 1–45.

Conley, T. G., and C. R. Udry (2008): "Learning About a New Technology: Pineapple in Ghana," mimeo, Yale University.

Davis, P. (2002): "Estimating Multi-Way Error-Components Models with Unbalanced Data Structures," *Journal of Econometrics*, 106(1), 67–95.

Donald, S. G., and K. Lang (2007): "Inference with Difference-in-Differences and Other Panel Data," *The Review of Economics and Statistics*, 89(2), 221–223.

Dube, A., T. W. Lester, and M. Reich (2008): "Minimum Wage Effects Across State Borders: Estimates Using Contiguous Counties," Institute for Research on Labor and Employment Working Paper Series, University of California, Berkeley.

Edmonds, E. V., K. Mammen, and D. L. Miller (2004): "Rearranging the Family? Income Support and Elderly Living Arrangements in a Low-Income Country," *The Journal of Human Resources*, 40(1), 186–207.

Goldstein, M., and C. Udry (2008): "The Profits of Power: Land Rights and Agricultural Investment in Ghana," *The Journal of Political Economy*, forthcoming.

Harrison, A., and J. Scorce (2008): "Multinationals and Anti-Sweatshop Activism," *The American Economic Review*, forthcoming.

Holmes, T. J. (1998): "The Effect of State Policy on the Location of Manufacturing: Evidence from State Borders," *The Journal of Political Economy*, 106(4), 667–705.

Kingdon, G. G., and J. Knight (2004): "Unemployment in South Africa: The Nature of the Beast," *World Development*, 32(3), 391–408.

——— (2006): "The Measure of Unemployment When Unemployment is High," *Labour Economics*, 13(3), 291–315.

——— (2008): "Unemployment in South Africa, 1995-2003: Causes, Problems, and Policies," *Journal of African Economies*, 16(5), 813–848.

Magruder, J. R. (2008): "Intergenerational Networks, Unemployment, and Persistent Inequality in South Africa," mimeo, Berkeley.

MBEKI, T. (2005): "Address of the President of South Africa, Thabo Mbeki, at the Second Joint Sitting of the Third Democratic Parliament, Cape Town.," http://www.info.gov.za/speeches/2005/05021110501001.htm.

MOLL, P. (1996): "Compulsory Centralization of Collective Bargaining in South Africa," *The American Economic Review*, 86(2), 326–329.

NICKELL, S., AND R. LAYARD (1999): "Labor Market Institutions and Economic Performance," in *Handbook of Labor Economics, Volume 3*, ed. by O. Ashenfelter, and D. Card. Elselvier-North Holland.

SBP (2005): "Counting the Cost of Red Tape to Business in South Africa," Discussion paper, Small Business Project.

SCHULTZ, T. P., AND G. MWABU (1998): "Labor Unions and the Distribution of Wages and Employment in South Africa," *Industrial and Labor Relations Review*, 51(4), 680–703.

SEEKINGS, J., AND N. NATTRASS (2005): *Class, Race, and Inequality in South Africa*. Yale University Press, New Haven and London.

STATISTICS SOUTH AFRICA (PRODUCER) (1997): "October Household Survey, 1997," South African Data Archive (distributor). www.nrf.ac.za/sada.

——— (2000-2003): "South Africa Labour Force Surveys, September 2000, September 2001, September 2002, September 2003," Pretoria, South Africa. South African Data Archive (distributor). www.nrf.ac.za/sada.

Population Group	Employed			Unemployed	
	All types	Self Employed	Domestic Work	Narrow	Broad
Black Males	49.85	5.84	2.05	20.89	35.73
Black Females	35.19	6.27	8.71	20.33	44.03
Coloured Males	66.02	3.94	1.05	14.56	21.18
Coloured Females	48.66	1.37	8.4	14.23	26.11
White Males	84.72	24.92	0.04	3.61	5.37
White Females	62.16	10.06	0.12	4.26	8.64

Table 1: Percentages of Adults in Employment Categories
Data are from the September 2003 labour force survey. Prime-aged adults are 20-60 years old.

	South Africa		Brazil	
	Male	Female	Male	Female
Num. Employees				
1	12.07	34.92	17.39	31.06
2-4	13.64	12.59	29.57	24.98
5-9	10.99	8.5	8.72	6.63
10-19	13.9	11.79	7.08	5.53
20-49	17.08	13.93	6.82	7.02
50 +	32.32	18.27	30.43	24.78
Percent Working	52.1	37.1	83.7	49.7

Table 2: Percentages of Employees in Firm Size Categories
South Africa Data are from the September 2003 labour force survey, while Brazillian data are from the 1995-95 LSMS survey. Satistics are for 20-60 year old adults.

	mean	sd	min	max	N
Log Employment	0.966	1.015	0.000	4.564	4984
Log Large Firm Emp	0.333	0.601	0.000	3.664	4984
Log Small Firm Emp	0.736	0.895	0.000	4.127	4984
Log Self Emp	0.502	0.765	0.000	3.932	4984
Log Single Firm Emp	0.346	0.651	0.000	3.829	4984
Log Population	4.825	0.974	2.773	6.868	4984
Bargaining Council	0.352	0.478	0.000	1.000	4984
Mean log wage	5.220	0.787	1.569	9.798	2698
Large Firm log wage	5.724	0.670	3.178	8.476	1256
Small Firm log wage	5.029	0.832	1.569	9.798	2232
Fraction Male	0.548	0.371	0.000	1.000	2698
Potential Exper	23.925	9.218	0.000	66.000	2698
Worker Education	8.510	2.699	0.000	12.000	2698
Mean Log Tenure	1.194	0.772	0.000	3.761	2505

Table 3: Summary Statistics

	Never a BC	BC Industry
All	58.36	41.64
Self	17.47	82.53
Small Firm	56.8[17]	43.2
Large Firm	60.21	39.79

Among Bargaining Council Industries		
	BC absent	BC present
All	46.97	53.03
Self	64.01	35.99
Small Firm	60.03	39.97
Large Firm	30.57	69.43

Table 4: Employment and Firm Size by BC status
Data are from the September 2003 labour force survey. Prime-aged adults are 20-60 years old.

Two Digit Industry	Fraction of all Workers	Fraction Covered
Textile Manufacturing	3.21%	100.0%
Metal Product Manf	2.05%	100.0%
Electical Machinery Manf	0.21%	93.0%
Transport Equip Manf	0.66%	100.0%
Furniture Manf	0.65%	77.2%
Construction	5.67%	35.7%
Retail Trade	13.28%	21.3%
Hotels and Restaurants	3.13%	36.1%
Land Transport	3.03%	100.0%
Other Business Acts	3.47%	19.5%
Public Administration	4.52%	100.0%
Recreational/Cultural Act	0.70%	100.0%
Other Service Activities	1.05%	51.4%

Table 5: Bargaining Council Coverage

	(1)	(2)	(3)	(4)
Employment	0.087	-0.071**	0.060	-0.137***
	(0.088)	(0.031)	(0.077)	(0.048)
Large Firm Emp	0.106	-0.047	0.068	-0.017
	(0.074)	(0.037)	(0.058)	(0.027)
Small Firm Emp	-0.015	-0.066**	-0.026	-0.122***
	(0.062)	(0.032)	(0.054)	(0.036)
Self Emp	-0.044	-0.062*	-0.057	-0.053
	(0.062)	(0.033)	(0.055)	(0.034)
Emp Rate	0.004	-0.0029*	0.002	-0.0051**
	(0.002)	(0.002)	(0.001)	(0.002)
Large Firm Rate	0.0014	-0.0001	0.0008	0.0002
	(0.001)	(0.000)	(0.001)	(0.000)
Small Firm Rate	0.0000	-0.0023*	-0.0007	-0.0039**
	(0.001)	(0.001)	(0.001)	(0.002)
Self Rate	-0.0007	-0.0014	-0.0012	-0.0015
	(0.001)	(0.001)	(0.001)	(0.002)
N	4984	4984	3551	3551
Sample	Full	Full	DC Ratio	DC ratio
Estimation	OLS	Diff-in-Diff	OLS	Diff-in-Diff

Table 6: No Spatial Fixed Effects: OLS and Difference-in-Differences

Table reports estimates of the coefficient on bargaining council presence for the dependent variable listed on the left. Spatial effects are not controlled for, and table rows represent similar specifications with different dependent variables. Bargaining council status is instrumented with bargaining council eligibility, and all results are conditional on time fixed effects and a quartic in log population. Difference-in-difference specifications include District Council-Industry fixed effects. All errors are clustered within the industry over space and time and among all industries, towns, and years in a given district council. DC Ratio include the subsample which have, on average, less than 20% of the employment in that industry in that district council.

Dep Variable	Employment		Employment Rate	
Fixed Effects Level				
None	-0.044	-0.116*	-0.0022***	-0.0049***
	(0.039)	(0.069)	(0.001)	(0.002)
Town	-0.083**	-0.146***	-0.003***	-0.0051***
	(0.035)	(0.049)	(0.001)	(0.001)
Town-Year	-0.106**	-0.143**	-0.004***	-0.0054***
	(0.044)	(0.060)	(0.000)	(0.001)
Town-Indus	-0.076**	-0.114***	-0.003*	-0.0041***
	(0.032)	(0.042)	(0.002)	(0.002)
N	4984	3551	4984	3551
Sample	Full	DC Ratio	Full	DC Ratio

Table 7: Employment Effects of Bargaining Councils

Presents coefficients of Bargaining Councils on log Employment (columns 1 and 2) or the Fraction of the population employed in an industry (columns 3 and 4). Rows consider different fixed effects. Bargaining council status is instrumented with bargaining council eligibility, and results are conditional on spatial-industry and time fixed effects, and a quartic in log population. All errors are clustered within the industry over space and time and among all industries, towns, and years in a given district council. The DC ratio sample has observations with on average, less than 20% of the employment in that industry in that district council

Dependent Variable	Wage	Wage	Small Wage	Large Wage
Fixed Effects Level				
None	0.195**	0.208***	0.183**	0.136
	(0.086)	(0.075)	(0.090)	(0.111)
Town	0.120**	0.172***	0.135**	-0.058
	(0.053)	(0.046)	(0.063)	(0.120)
Town-Year	0.074	0.126**	0.177**	-0.055
	(0.101)	(0.055)	(0.073)	(0.144)
Town-Indus	0.091		0.115	-0.069
	(0.059)		(0.114)	(0.064)
Radius	30	50	50	50
Sample	Wage	Wage	Small Wage	Large Wage
N	2698	2698	2232	1256

Table 8: Wage Effects of Bargaining Councils

Presents coefficients of Bargaining Councils on mean log wages, and mean log wages in small or large firms. Rows consider different fixed effects. Bargaining council status is instrumented with bargaining council eligibility, and results are conditional on spatial-industry (with a radius given in the radius row) and time fixed effects, and a quartic in log population. All errors are clustered within the industry over space and time and among all industries, towns, and years in a given district council.

Dependent Variable	Wage	Male	Educ	Age	Tenure
Fixed Effects Level					
None	0.160**	0.082***	-0.111	-0.404	0.190***
	(0.062)	(0.017)	(0.205)	(0.924)	(0.051)
Town	0.138***	0.095***	-0.260*	-1.151	0.193**
	(0.022)	(0.023)	(0.156)	(1.040)	(0.080)
Town-Year	0.076**	0.131***	-0.221	-1.486	0.117
	(0.032)	(0.037)	(0.200)	(1.262)	(0.071)
Town-Indus	0.080	0.041**	-0.283***	-0.784	0.181**
	(0.052)	(0.019)	(0.096)	(0.852)	(0.075)
Radius	50	50	50	50	50
Sample	wage	wage	wage	wage	tenure
N	2698	2698	2698	2698	2505

Table 9: Wage Effects: Employee Composition Controls

Presents coefficients of Bargaining Councils on mean log wages, the fraction male, mean education, age, and log tenure, with analysis restricted to observations with at least one wage (wage sample) or tenure (tenure sample) observation. Rows consider different fixed effects. Bargaining council status is instrumented with bargaining council eligibility, and all are conditional on 50-mile spatial-industry and time fixed effects, a quartic in log population, and worker composition variables. All errors are clustered within the industry over space and time and among all industries, towns, and years in a given district council.

	(1)	(2)	(3)	(4)
Large Firm Emp	0.018	-0.052	-0.040	-0.052
	(0.046)	(0.041)	(0.044)	(0.038)
Small Firm Emp	-0.103***	-0.113***	-0.150***	-0.070**
	(0.037)	(0.029)	(0.030)	(0.033)
Self Emp	-0.104**	-0.118***	-0.149***	-0.067**
	(0.044)	(0.039)	(0.052)	(0.034)
Single Emp	-0.058	-0.047	-0.091**	-0.017
	(0.042)	(0.036)	(0.041)	(0.022)
Large Firm Rate	0.001	0.000	0.000	0.000
	(0.000)	(0.000)	(0.001)	(0.000)
Small Firm Rate	-0.003***	-0.0029***	-0.0038***	-0.002*
	(0.000)	(0.001)	(0.001)	(0.001)
Self Rate	-0.002**	'-0.0021**	-0.0026***	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)
Single Emp	-0.001	-0.001	-0.001	0.000
	(0.001)	(0.001)	(0.001)	(0.001)
N	4984	4984	4984	4984
Fixed Effects	None	Town	Town-Year	Town-Indus

Table 10: Bargaining Council Effects by Firm Size

Presents coefficients of Bargaining Councils on log Employment, and employment rates for different firm size categories. Rows report the coefficient on bargaining council status for specifications with different dependent variables, while columns indicate the level of fixed effects. Bargaining council status is instrumented with bargaining council eligibility, and results are conditional on spatial-industry and time fixed effects, and a quartic in log population. All errors are clustered within the industry over space and time and among all industries, towns, and years in a given district council.

	(1)	(2)	(3)	(4)
Large Firm Emp	0.011	-0.029	-0.016	-0.015
	(0.030)	(0.028)	(0.039)	(0.022)
Small Firm Emp	-0.118**	-0.120***	;-0.124**	-0.101***
	(0.053)	(0.044)	(0.055)	(0.031)
Self Emp	-0.050	-0.070*	-0.081	-0.045
	(0.049)	(0.041)	(0.059)	(0.030)
Single Emp	-0.045	-0.031	-0.046	-0.019
	(0.046)	(0.040)	(0.055)	(0.018)
Large Firm Rate	0.000	0.000	0.000	0.000
	(0.001)	(0.000)	(0.001)	(0.000)
Small Firm Rate	-0.0041**	-0.0036***	-0.0036***	-0.0031**
	(0.002)	(0.001)	(0.001)	(0.001)
Self Rate	-0.002	-0.002	-0.002	-0.001
	(0.002)	(0.002)	(0.002)	(0.001)
Single Emp	-0.001	0.000	-0.001	0.000
	(0.001)	(0.001)	(0.001)	(0.001)
N	3551	3551	3551	3551
Fixed Effects	None	Town	Town-Year	Town-Indus

Table 11: Bargaining Council Effects by Firm Size: DC Ratio Sample

Presents coefficients of Bargaining Councils on log Employment, and employment rates for different firm size categories. Rows report the coefficient on bargaining council status for specifications with different dependent variables, while columns indicate the level of fixed effects. Bargaining council status is instrumented with bargaining council eligibility, and results are conditional on spatial-industry and time fixed effects, and a quartic in log population. All errors are clustered within the industry over space and time and among all industries, towns, and years in a given district council. Town-industry observations with more than a 20% employment share in the district council are excluded.

Dependent Var	Employment		Large Firm Emp		Small Firm Emp	
BC 30 Mile Border	0.045**	-0.097	0.028	-0.203*	-0.034	-0.107
	(0.021)	(0.173)	(0.118)	(0.112)	(0.085)	(0.152)
Not BC 30 Mile Border	0.368**	0.134	0.084	-0.027	0.478**	0.253**
	(0.151)	(0.089)	(0.084)	(0.100)	(0.219)	(0.113)
BC 50 Mile Border	0.109	0.121	-0.079	0.051	0.125*	0.069
	(0.100)	(0.219)	(0.134)	(0.134)	(0.068)	(0.198)
Not BC 50 Mile Border	-0.149	0.048	0.073	0.057	-0.277*	-0.039
	(0.142)	(0.052)	(0.082)	(0.040)	(0.155)	(0.049)
BC	-0.119**		-0.001		-0.113**	
	(0.048)		(0.030)		(0.053)	
Fixed Effects	Town-Ind	DC-Ind-Yr	Town-Ind	DC-Ind-Yr	Town-Ind	DC-Ind-Yr
N	4984	4984	4984	4984	4984	4984

Table 12: Border Jumping

Regresses employment or employment by firm size on bargaining council status, as well as border status at different border lengths, where the effect of being on a border is allowed to be asymmetric by which side of the border a town is on (BC side or Not BC side). All results are conditional on time fixed effects and a quartic in log population

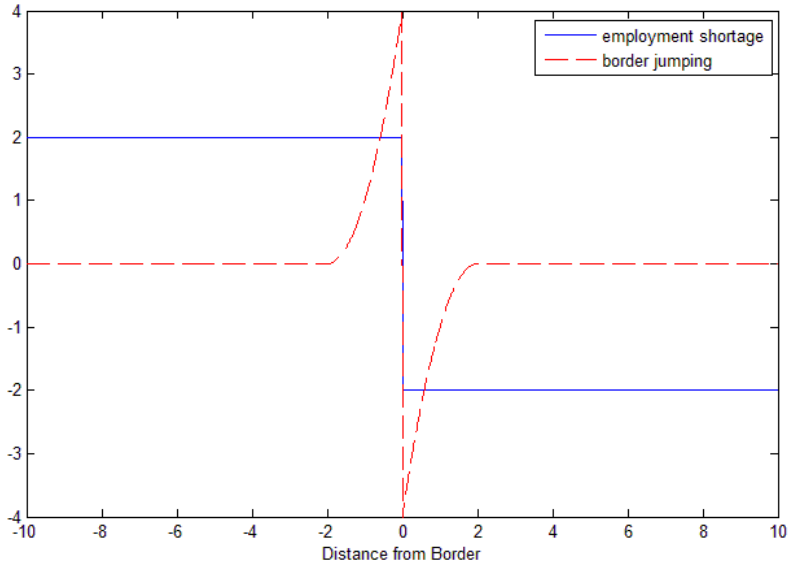

Figure 1: Employment shortage vs. border jumping

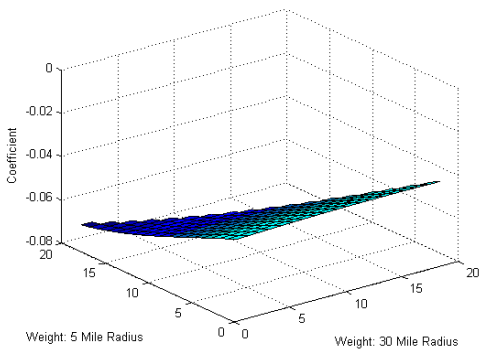

Figure 2: Employment Coefficients

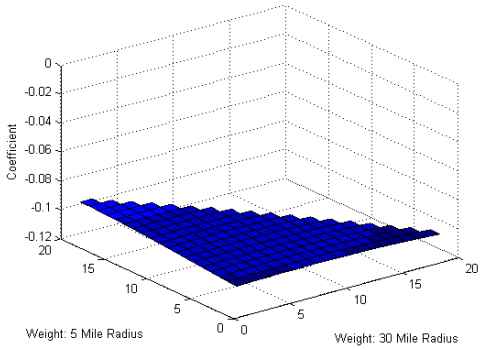

Figure 5: Employment Coefficients: Town-Year FE

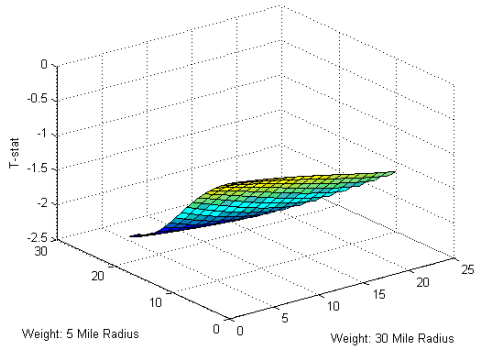

Figure 3: Employment t-statistics

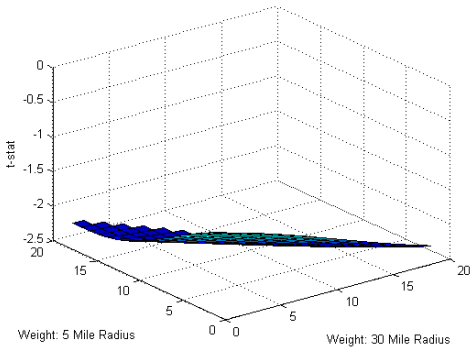

Figure 6: Employment t-stats: Town Year FE

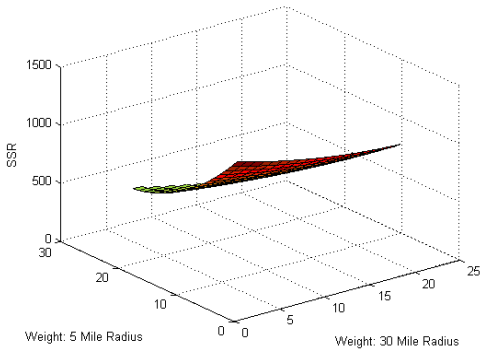

Figure 4: Employment Sum of Squared Residuals

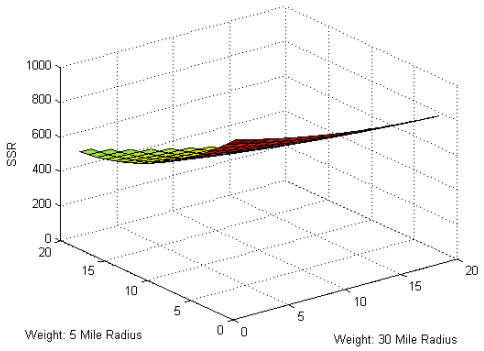

Figure 7: Employment Sum Squared Resid: Town-Year FE

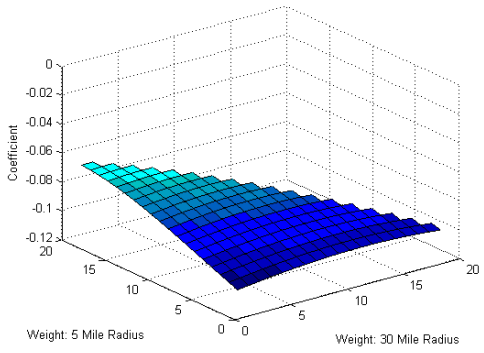

Figure 8: Small Firm Emp Coefficients

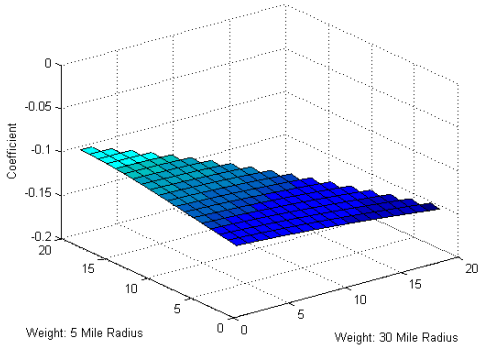

Figure 11: Small Firm Coefficients: Town-Year FEs

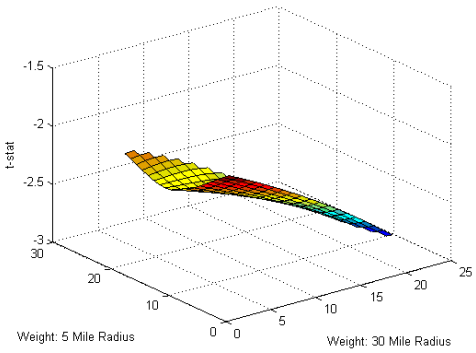

Figure 9: Small Firm t-stats

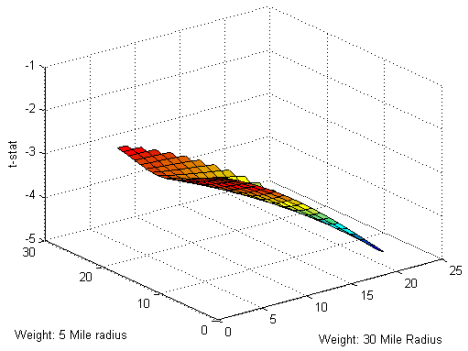

Figure 12: Small Firm t-stats: Town-Year FEs

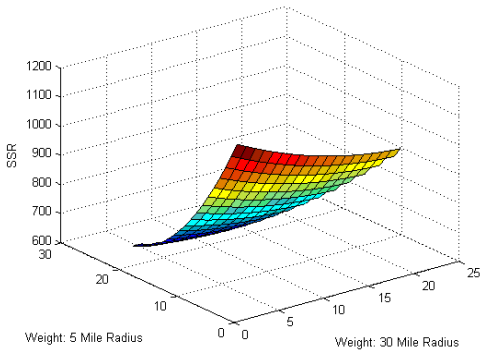

Figure 10: Small Firm Sum Squared Resids

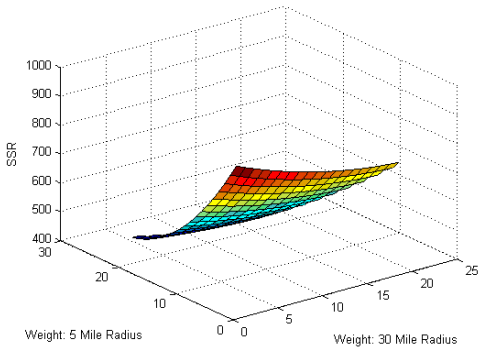

Figure 13: Small Firm Sum Squared Resids: Town-Year FEs