The Impact of Unions on Non-union Wage Setting: Threats and Bargaining*

David A. Green† Ben M. Sand‡ Iain G. Snoddy§ Jeanne Tschopp¶

February 6, 2025

Abstract

In this paper we provide new estimates of the impact of unions on non-union wage setting. We allow the presence of unions to affect non-union wages both through the typically discussed channel of non-union firms emulating union wages in order to fend off the threat of unionisation and through a bargaining channel in which non-union workers use the presence of union jobs as part of their outside option. We specify these channels in a search and bargaining framework that includes union formation and the possibility of non-union firm responses to the threat of unionisation. Our results indicate an important role played by union wage spillovers in lowering wages over the 1980-2010 period. We find that de-unionisation can account for nearly a third of the decline in the mean hourly wage between 1980 and 2010 in the US, with half of that effect being due to spillovers. Both the traditional threat and bargaining channels are operational, with the bargaining channel being more important.

*We are grateful to Thomas Lemieux, Craig Riddell, Raffaele Saggio and Anastasia Burya for their comments. We also greatly benefited from comments from participants at seminars at the University of Chicago, the University of Melbourne, Carleton University, the University of British Columbia, the University College London, the University of Mannheim, the TU Dortmund University, the Banque de France, and the Society of Labor Economists annual meeting.

†University of British Columbia and Research Fellow, IFS, London david.green@ubc.ca

‡Department of Economics, York University bmsand@yorku.ca

§Analysis Group, Montreal iain.snoddy@analysisgroup.com

¶University of Bern jeanne.tschopp@unibe.ch

1 Introduction

Private sector unionisation in the United States is very nearly dead. In 2020, only 6.3% of private sector workers belonged to a union (U.S. Bureau of Labor Statistics (2021)). Recently, however, there have been some glimmers of revival, including successful unionisation drives at Amazon and Starbucks, raising questions about whether a resurrected union movement could significantly impact wage levels across the economy. While we can't look ahead to future changes in unionisation, we can use the de-unionisation over the last 50 years to better understand union impacts. The most direct impact of decreased unionisation, of course, comes from the shifting of workers out of higher-paid union jobs. But it also has the potential to alter wage setting in non-union jobs. Such spillover effects are important since their existence would imply that the reach of unions is larger than it might first appear and larger than what is calculated based on standard shift-share decompositions. In this paper, we build a model of the impact of unions on wage setting in the non-union sector and use it in estimation based on Current Population Survey (CPS) data to re-assess the role of de-unionisation in movements in the wage structure in the U.S.

The idea that unions could impact non-union wage setting goes back at least to Lewis [1963]. The core idea raised in that book, and discussed in subsequent papers such as Rosen [1969], is that non-union firms raise their wages in response to the 'threat' that their workers will unionise, presumably imposing extra costs beyond direct wage increases.[1] Our model incorporates that threat effect plus an added union impact mechanism: a bargaining channel whereby the outside options of non-union workers and, through that, their bargained wages are affected by their ability to find high-paying union jobs. In a sense, both are threats, with one being the threat of workers leaving and finding a union job (what we will call a bargaining effect) and the other channel being the threat to unionise the non-union workplace (which we call a standard threat effect).

These two channels, however, have different implications for attempts to raise wages through policy tools. The threat effect is unique – it can only be harnessed by increasing the threat of unionisation. The bargaining effect, in contrast, is more general. It is about having more well-paid jobs in a location, improving the outside options for workers in all other jobs. As noted in Beaudry et al. [2012] and Caldwell and Danieli [2021], this can significantly increase wages for all workers in the location. Unions are one way to create a higher wage option, but other policies, such as eliminating non-compete arrangements, could also have such an impact [Johnson et al., 2020]. Our model clarifies the difficulties inherent in identifying these two effects separately while controlling for selection into the union/non-union sectors. Part of our contribution is to offer estimates of spillover effects through both channels, expanding our understanding of the impact of de-unionisation on the wage structure. Based on our estimates, we then assess how de-unionisation has contributed to changes in the wage structure in the U.S. over recent decades.

The existing literature estimates union wage spillovers by regressing non-union wages on the percent of organised workers in labour markets defined by location and/or industry.

[1] For instance, Starbucks recently offered wage increases to "company-operated stores" but not to "unionised stores, or to stores that may be in the process of unionising". The NLRB has designated the announcement as a threat, designed to have a "chilling effect" on impending union votes (New York Times, May 2022).

Evidence based on this approach is mixed and sensitive to the included control terms, with the preponderance of studies finding a small positive spillover effect.[2] In an important analysis, Farber [2005] carefully considers the role played by omitted variables, sequentially introducing industry and state-fixed effects. He finds great sensitivity in estimates to the source of variation used, providing some context for the disparity in estimates across earlier studies. When controlling for a wide range of potential omitted variables, his results indicate, at most, a small positive effect of union power on non-union wages.

In a recent paper, Fortin et al. [2021] estimate the impact of union threat effects on wage inequality, using industry×state-level variation in the unionisation rate as an additional covariate in their distribution regression approach. They find positive effects of the unionisation rate operating primarily at the part of the wage distribution just below the median, and their counterfactual exercise indicates that spillovers double the measured impact of de-unionisation in increasing wage inequality in the U.S.. While very useful, the paper shares with all of the early analyses a lack of an identification strategy for addressing the potential endogeneity of the union proportion – stemming from a lack of an effective instrument.[3] None of the papers in the literature even mention the twin problem of potential selectivity bias. As the proportion of unionised workers declines, the composition of non-union workers and firms will change.

In contrast to the existing literature on union spillovers that largely relies on reduced-form estimation, our approach formalises union spillovers in a search and bargaining framework, endogenising the union formation process and incorporating wage effects arising through differences in the bargaining process. In making clear what is being identified in the model and the variation used, we overcome the problems inherent in early studies of likely biases due to omitted characteristics and selection into the union sector, and we estimate an effect with a clear theoretical basis and interpretation.

Our model is based on that of Taschereau-Dumouchel [2020] (henceforth TD), whose work is informed by the contributions of Pissarides [1986], Açıkgöz and Kaymak [2014], and Krusell and Rudanko [2016], among others. The TD model is centred around union threat effects through the hiring channel. In the model, more skilled workers tend to dislike unionisation, and firms skew their hiring toward these workers to stack the unionisation vote. Though this effect is certainly interesting, we believe it is likely of second-order importance relative to a more direct firm response through raising wages to lessen the gains from unionising and direct union-busting actions, which raise the costs of unionisation. Our model focuses on these latter effects instead of the hiring channel.

Additionally, our framework is informed by papers, including Beaudry et al. [2012] (henceforth BGS), Tschopp [2017], Caldwell and Danieli [2021], Jarosch et al. [2024], and

[2] Both Freeman and Medoff [1981] and Donsimoni [1981] find a non-significant positive correlation between non-union wages and the proportion of unions. Conversely, Holzer [1982], Kahn [1980] and Dickens and Katz [1986] estimate large positive effects. Hirsch and Neufeld [1987] find a positive spillover effect at the industry level but insignificant effects at the local labour market level. Podgursky [1986] finds spillover effects exist only for large establishments, and Neumark and Wachter [1995] estimate a negative effect.

[3] Farber [2005] presents event studies of the enactment of Right to Work (RTW) laws in Idaho and Oklahoma. In an earlier version of Fortin et al. [2019], the authors extend Farber's analysis to include more states. They find evidence of reductions in non-union wages with the introduction of RTW laws but the estimates are poorly defined because few states switch RTW status in their time period.

Bassier [2022], which formalise the impact of changing outside options on wages. Following BGS, we model local labour markets composed of industries and firms with workers able to transition between jobs in proportion to job prevalence. As in BGS, we exploit cross-city, within-industry variation – in our case, to identify the effect of declining unionisation on non-union wages from 1980 to 2010. Our model is partial equilibrium in the sense that we treat prices, city level employment rates and transition rates at the national level as exogenous.

Combining these elements, we derive an empirical specification which incorporates spillover effects operating through both the bargaining and standard threat channels, formalises selectivity, and makes it straightforward to see barriers to identification. Specifically, changes in outside options associated with the union sector may be correlated with unobserved local productivity shocks. As in BGS, we overcome this problem using Bartik-style instruments related to worker outside options. For non-union workers, outside options are related to the probability the worker can transit to a union job (which we allow to vary by industry and over time) times the expected wage the worker could get in that job. It also depends on expected wages in non-union jobs in the local economy and the probabilities of transiting to those jobs. Our instruments use versions of these outside options based on the start-of-period industry and union employment composition in a locality interacted with changes in industry growth, industry premia, and the probabilities of moving to different types of jobs defined by industry and union status at the national level.

The outside option for non-union workers identifies the bargaining channel for union effects. We get extra power to identify the bargaining effect because improvements in outside options have the same effect on bargained wages, whether they stem from reduced probabilities of finding a union job or a high-rent non-union job. That means we get identification from both unionisation changes and industrial structure shifts in both the non-union and union sectors. We argue that the validity of our instruments depends on a random walk-type assumption that we show implies an over-identifying restriction. We test that restriction and cannot reject it. Working from the model, we identify the threat channel by the impact on non-union wages of the interaction of the probability a firm in a given industry×city cell would face a union election (which shows the size of the direct threat) with the outside option value for union workers (which captures the size of what the firm needs to respond to in order to prevent unionisation). We construct and implement similar Bartik type instruments related to this component.

The results from our estimation point to the importance of both spillover channels. Between 1980 and 2010, the mean real wage for job entrants in the U.S. fell 16% (holding composition in terms of education, experience, race and gender constant). A decomposition exercise based on our estimates shows that de-unionisation accounts for a third of the decline. A third of that impact arises from a standard shift-share effect (because workers shifted away from higher-paying union jobs), while over half comes from spillover channels. Unions have spillover effects on non-union wages, and they are sizeable. While both the traditional threat and bargaining effects enter significantly, our decomposition exercise indicates that the spillover effects are almost entirely due to the latter. The threat probability was too low, even in 1980, to play a substantial role. As we pointed out earlier, the dominance of the bargaining channel means the effects of unions in raising non-union wages could also be achieved through other policies that raise average worker rents. The effect is not unique to unions.

Our estimates imply that spillovers roughly doubled the standard shift-share effect of unionisation over the long run. Perhaps surprisingly, the 1980s account for only half of the total effect of de-unionization from 1980 to 2010, even though this decade saw the largest decline in unionization. This is because those declines were offset by increases in the union wage premium, increasing the value of the outside option of non-union workers while the declining probability of finding union jobs reduced it. Our model provides an explanation for the increased wage premium in the 1980s, which echoes an argument in Farber [2005]. While both union and non-union wages faced downward pressures from technological change, trade, etc., the substantial reduction in the risk of being unionised in the decade meant that, in addition, non-union firms no longer had to pay higher wages in order to stave off unionisation. As a result, non-union wages fell faster than union wages. After 1990, the threat of unionisation stabilised at a low level, causing the union wage premium to decline, and the outside option effect of unions began to reflect the falling unionisation rate alone. The potential lesson for any re-unionisation efforts is that spillover effects onto non-union wages may arise through the traditional threat channel but the implied increase in non-union wages will dampen the bargaining channel. Union jobs would be more plentiful but not pay as high a premium over non-union jobs as before re-unionisation. Eventually, as the unionisation threat stabilised, the extent of spillover onto non-union wages would increase, but that could take time to realise fully.

Our work is also related to the substantial literature investigating patterns in declining unionisation, estimating both movements in the union wage premium and the role of declining unionisation in driving increasing wage inequality. Card et al. [2004] and Card et al. [2018] provide comprehensive summaries of the research in this area following the early contribution of Freeman [1980]. Farber et al. [2021] provides the most comprehensive account of the relationship between union density and inequality in the U.S., introducing new survey data that allows them to push their analysis back to the 1930s. They find that increasing unionisation substantially impacted decreasing inequality after WWII, while the reversal in the unionisation trend had a smaller effect on increasing inequality in the last 50 years. Their estimates allow for spillover effects onto non-union wages, but they do not study spillovers directly. Our results imply that spillovers may have played an important role in their estimated inequality impacts from unions and explain why those impacts were less evident at the time of the big union decline in the 1980s.[4]

The remainder of the paper is organised as follows. In Section 2, we present our model. In Section 3, we derive our empirical specification and discuss the implementation and identification of challenges and solutions. We also present the construction of our key outside option variables and our instrumental variables. Section 4 describes the data and Section 5 contains our estimation results. In Section 6, we present a counterfactual exercise designed to demonstrate spillovers' impact on wage structure movements and the role played by our two channels. Section 7 contains conclusions.

[4]Other papers in this literature include an important contribution by DiNardo et al. [1996], which attributes 14% of the increase in wage inequality over 1979-1988 (for men) to declining unionisation. Extensions of this work are found in DiNardo and Lemieux [1997] and Fortin et al. [2021]. Further studies by Card [2001], Card et al. [2004], Gosling and Lemieux [2001], and Card et al. [2018] extend the analyses by sector, gender, and across countries. See also recent studies by Farber et al. [2021] and Firpo et al. [2018].

2 The Model

2.1 Model Set-up

Our goal with our model is to derive an estimable specification for non-union wages that captures key channels through which those wages can be affected by changes in unionisation. Our model is based on that of Taschereau-Dumouchel [2020] (TD), which places union formation and wage setting in a search and bargaining model. Unions are able to bargain a higher wage because they can threaten to take the whole workforce out of production, while an individual, non-union worker can only threaten to withdraw her own labour. As mentioned in the introduction, TD focuses on firms responding to the threat of unionisation by altering the skill composition of their hiring while we focus on a response through paying higher wages. Through the rest of the paper, we will refer to non-union firms' wage responses to resist unionisation as standard threat effects (to reflect that these are what have been discussed in the previous literature).

In addition to standard threat effects, we allow for unionisation levels to affect non-union wages through a bargaining channel. Since unions can bargain higher wages for their members, having more unionised jobs in the local economy improves the outside option for all workers – even workers in firms not directly threatened with unionisation or workers in different industries – thus raising their wages. We refer to any such effects as bargaining effects. To investigate whether this channel has sizeable effects, we alter the TD model by having only one skill level but, following Beaudry et al. [2012](BGS), multiple industries.[5]

In the model, there are C cities indexed by c, and we are interested in differences in non-union wages across cities with different unionisation levels. There are also I industries, indexed by i, which are assumed to produce tradeable goods. Worker-firm matches die at an exogenous rate, δ^m, and all agents face a common discount rate, ρ. Firms face an additional probability of closing down, δ^e, with new firms born at the same rate to keep the number of firms fixed. Workers search for jobs while unemployed. The model is partial equilibrium in the sense that we treat prices, the number of firms, the meeting rates between workers and vacancies, and the local employment rate as exogenous.[6] The model is centred on workers and firms (endogenously) ending up in one of three types of arrangements: simple non-union firms, non-union firms that emulate union wages, or union firms. Which arrangement is implemented is determined through a strategic interaction between workers and firms made under an assumption of perfect information.

To understand the intuition underlying our model, it is helpful to go through its timing.

1. Firms are all born non-union. At the time of birth, both their productivity and the value of an idiosyncratic amenity that workers would create should they unionise the firm are revealed. The values of the amenity and firm productivity – assumed to be independent of each other – along with unionisation costs, will determine which firms become unionised.

[5] We bring differing skill levels back in through a model-consistent route in our empirical specification.

[6] Working in partial equilibrium in this way eliminates a channel through which de-unionisation could affect wages by lowering labour costs, causing firms to post more vacancies and, through that, increasing labour market tightness. This channel would have effects that are opposite to those of the channels we emphasize. We return to this channel in the empirical work.

2. Following TD, firms first unilaterally determine their optimal level of employment and open vacancies to meet that target. The target employment will depend on the anticipated wage that will be bargained with workers, which, in turn, depends on which of the three arrangements is relevant to the firm. In steady state, the firm knows which arrangement it will experience.

3. Next, workers and firms meet according to a matching technology that is allowed to vary by the industry and union status of the vacancy as well as those of the previous job held by the worker. For example, a worker formerly in a unionised construction job may find it easier to be hired at a unionised auto plant than another worker who was formerly a non-unionised retail employee.

4. Once matches are made, union status is determined, and wages are bargained accordingly. We represent the union arrangement determination game in Figure 1.

   - We assume that the workers initially have a straw poll amongst themselves, determining whether they are interested in proceeding with the formal (costly) process of unionising by majority vote. They then signal the outcome of that vote to the firm. Their decision will be determined by a comparison of the present value of being non-union to the present value of being union minus the cost of unionisation for the median voter. There will be some values of the combination of firm-specific amenities, firm productivity, and the costs of unionising that imply that the workers will vote not to unionise in their straw poll, and they will signal that decision to the firm (branch *1*). Given perfect information, they cannot credibly threaten to unionise in order to try to get concessions from the firm. Given a decision by the workers not to unionise, they proceed to bargain wages individually with the firm. We call this the Simple Non-union arrangement.
   - If the workers threaten to unionise (branch *2*), the firm can either respond or not respond. If the cost of unionisation is larger than the value of the amenity, unionisation will cause a reduction in the total surplus of the match between workers and the firm. For some values of the amenity, productivity, and the cost of unionising, the enhanced bargaining power from unionising means the amount of the surplus the workers capture is still larger from unionising than remaining non-union, even though the total pie is smaller. In that case, the workers and the firm would jointly be better off if the workers do not unionise but are given a wage that gives them a bigger total slice than what they get in the Simple Non-union arrangement.[7] Recognizing this, the firm will post a wage in the workplace that is sufficient to make the workers at least indifferent to unionising. The workers will take a straw vote on that offer and if they reject it, the firm will make a new offer. This repeated offer process will continue, reaching a final result that can be represented by the Nash bargaining solution with the value of unionised outcomes for both the workers and the firm as the fallback options.[8] The resulting wage will

[7] In contrast, the Simple Non-union arrangement arises when unionisation reduces the total surplus so much that even with enhanced bargaining power, the workers end up with smaller pieces of pie than they can get from individual bargaining in a non-union setting.

[8] This generalizes the classic response to a unionisation threat seen in Rosen [1969], where the wage increase to avoid unionisation is a unilateral firm decision.

be higher than the Simple Non-union wage. We call this arrangement the Emulating Non-union Arrangement and it corresponds to branch *2.2*.

- Finally, if the value of the amenity exceeds the cost of unionisation then the total value of the surplus is larger under unionisation than under either non-union arrangement. In this case (branch *2.1*) the firm will not respond to the worker threat to unionise. The workers will then proceed to unionise (bearing the cost of doing so) and wages will be set in bargaining between the firm and the union (i.e., the wholes set of workers). As in TD, the worker outside option remains the value of being unemployed, but for firms, a breakdown in bargaining means a complete shutdown in production, which is why unions can bargain higher wages. We will refer to these firms as Union firms.

Ultimately, the arrangement is chosen to maximize the joint surplus (of the firm and the entire workforce) and the wage allocates that surplus based on bargaining power, with no alternative agreement able to make one party better off without harming the other.

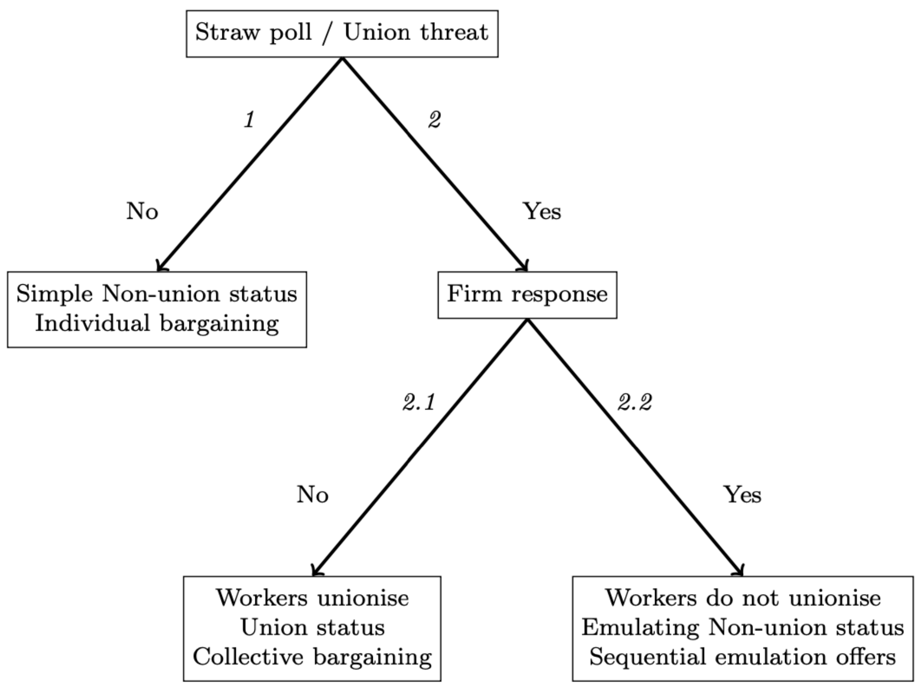

Figure 1: Sequence of events after hiring.

In this model, an increase in the cost of unionisation reduces the union threat and, thus, both the number of Emulating Non-union firms and the wages that they pay. It also reduces workers' outside options in Simple and Emulating Non-union firms because there are fewer higher-paid Union and Emulating Non-union firms for them to move to. The reduction in the outside option is relevant for workers in all sectors, lowering their bargained wages.

With this structure in mind, we next fill in the details needed to derive our estimating wage equations. A complete derivation of the model can be found in Appendix A.

In what follows, we index union arrangements by τ with: $\tau = u, n$ and e corresponding to Union firms, Simple Non-union firms, and Emulating Non-union firms, respectively. We index 'jobs' by j with $j = \{\tau, i\}$, i.e., jobs are combinations of union arrangements and industry. We use the subscript k for the potential destination jobs, with destination union status and industry, denoted τ' and i', respectively, such that $k = \{\tau', i'\}$.

2.2 Matching

Firms and workers operate in a labour market with frictions, meaning that workers and firms do not find each other and form a match perfectly easily. We assume that match formation depends on both the job type in which the vacancy is posted and the job type in which the worker was last employed. In particular, we employ a matching function of the form:

$$M_{kc|jc} = \theta_{jc} M(U_c, \Omega_c) \phi_{kc} \chi_{kc|j}(\varphi_{k|j}) \tag{1}$$

where $M_{kc|jc}$ is the number of matches of unemployed workers whose last job was of type j to vacancies of job type k in city c; θ_{jc} is the proportion of unemployed workers who were formerly in j; ϕ_{kc} is the proportion of vacancies that come from k; $M(U_c, \Omega_c)$ is the total number of matches observed in a city, with U_c being the total number of unemployed workers and Ω_c, the total number of vacancies in the city; and $\chi_{kc|j}(\varphi_{k|j})$ represents the specific frictional costs of moving from j to k. Thus, the number of matches of workers from j to vacancies in k equals a purely mechanical component (the total number of matches times the proportion of unemployed workers who come from j and the proportion of vacancies of type k) times a component representing the fact that there are barriers to forming some j, k matches.[9] For example, a match between a worker formerly in a unionised construction job and a vacancy posted by a unionised steel firm may be particularly easy to consummate while a match between that same worker and a non-union legal services firm may be less likely to actually happen. $\chi_{kc|j}(\varphi_{k|j})$ represents these frictional costs and takes the form:

$$\chi_{kc|j} = \frac{\varphi_{k|j}}{\sum_{k'} \eta_{k'c} \varphi_{k'|j}} \quad \forall k \tag{2}$$

where $\varphi_{k|j}$ represents the specific mobility frictions in moving from type j to type k jobs (regardless of city), and $\eta_{k'c}$ is the proportion of employment in job k' in city c, which equals $\phi_{k'c}$ in steady state. We shift to using $\eta_{k'c}$ from here on because we observe employment shares but not vacancies in our data. Assuming (as is standard) that $M(U_c, \Omega_c)$ is constant returns to scale (CRS), $M_{kc|jc}$ is also CRS. As we show in Appendix A.1, in steady state, θ_{jc}, ϕ_{kc} and $\chi_{kc|j}(\varphi_{k|j})$ all adjust to maintain a constant matching rate and sectoral composition.

The probability that a firm fills a vacancy of job type k is $q^v_{kc} = \frac{M_{kc}}{\Omega_{kc}}$, where $M_{kc} = \sum_j M_{kc|jc}$ and $\Omega_{kc} = \phi_k \Omega_c$. Appendix A.1 shows that given the CRS assumption, in steady state, $q^v_{kc} = \frac{M_c}{\Omega_c} = q^v_c$. Hence, the probability that a firm fills a job is independent of the specific job and only depends on the local matching process. In a similar vein, the probability an unemployed worker from j makes a successful match with a vacancy in k equals $q^u_{kc|j} = q^u_c \eta_{kc} \chi_{kc|j}(\varphi_{k|j})$, where $q^u_c = \frac{M_c}{U_c}$. Thus, the probability an unemployed worker who last worked at a job, j, matches to a vacancy in k is a function of the overall average probability unemployed workers make matches, the proportion of employment in job k, and

[9] Characterizing differential match rates as reflecting differential frictional costs follows Tschopp [2017]. Bassier [2022], alternatively, refers to differences in worker movements across firms as reflecting differences in 'consideration sets'. Caldwell and Harmon [2019] discusses differences based on personal networks.

the mobility friction, $\varphi_{k|j}$. Thus, in our example, a worker whose last job was in the unionised construction sector may have a high probability of consummating a match with a unionised steel firm in general (i.e., have a high value of $\varphi_{k|j}$) but have a low probability of actually making that match if either the local labour market is very slack or there are very few steel firms in the city (either q_c^u or η_{kc} is low).

2.3 Firms

We assume the number of firms operating with type j jobs (a combination of union status τ and industry i) in city c is fixed, leaving the endogenisation of firm formation for future work. For notational clarity, we drop the (firm-job-city-specific) subscript on firm employment and vacancies. All firms operating in a given industry have a common production function:

$$y_{fjc}(n) = \epsilon_{fic} n - \frac{1}{2}\sigma_i n^2,$$

where ϵ_{fic} is a firm-specific productivity draw, n is the number of employees, and $\sigma_i > 0$ is a parameter reflecting the potential span of control issues.[10] It will prove useful to write $\epsilon_{fic} = \epsilon_{ic} + u_{fic}$, where ϵ_{ic} corresponds to a sector-wide productivity shock (at the city level) and u_{fic} is a mean zero, firm-specific shock. This specification implies that technology is common across cities within an industry but that there are differences across cities in comparative advantage in producing each good, captured in the ϵ_{ic}'s. The technology also does not vary by union status. The literature on union productivity effects seems to us to be inconclusive, and so we assume that unions affect firm activity by affecting wages (and employment) but not through technological adaptations.[11] We assume that the σ_i's are sufficiently smaller than 1 that, combined with the assumption of a fixed number of firms in each ic cell, they imply that production of any good is spread across cities.

At the beginning of each period, firms choose the optimal number of vacancies (and, so, optimal employment) given the wage (specified as a function of firm employment) they know will be bargained with their workers later. To simplify, we assume that the flow cost of hiring is linear in the number of vacancies posted. Since δ^m matches are randomly destroyed in each period, a firm which had n_{-1} workers in the previous period enters the current period with $(1-\delta^m)n_{-1}$ workers. From this, it knows the number of vacancies, v, it must post in order to have n workers for production in the current period. Hence, the firm value function of filled positions is given by:

$$\begin{aligned} \Pi_{fjc}(n_{-1}) = \max_{v} \quad & [p_i y_{fjc}(n) - w_{fjc}(n)n - \kappa v + \rho^e \Pi_{fjc}(n)] \\ \text{s.t.} \quad & n = n_{-1}(1-\delta^m) + q_c^v v, \end{aligned} \tag{3}$$

where p_i is the price of the industry i good, $w_{fjc}(n)$ is the wage bargained at the firm for this type of job with n workers at the firm, and κ is the cost per vacancy posted. ρ^e is the firm effective discount rate, taking account of the firm death rate, i.e. $\rho^e = (1-\delta^e)\rho$. Note

[10]The production function includes a firm-job-city-specific subscript to account for the variations in firm employment, which differs across firms depending on their union status.

[11]Hirsch and Link [1984] and Addison and Hirsch [1989] summarise the early research in this area which finds largely inconclusive and mixed evidence on the effect of unionisation on productivity.

that we assume that firm-specific union amenities are created by the union and, so, do not enter the cost function of the union firm.

2.4 Workers

The value of employment for a worker in a job of type j in firm f is given by:

$$V^E_{fjc}(n) = w_{fjc}(n) + \psi_{fjc} + \rho[\delta V^U_{jc} + (1-\delta)V^E_{fjc}(w'_{fjc})] \tag{4}$$

where ψ_{fic} is a non-wage amenity for workers from being in a union in this particular firm and, so, equals zero in non-union firms. w'_{fjc} is the wage that will be paid by the firm in the next period if the job is not terminated and δ is the job destruction probability; i.e. $\delta = \delta^e + (1-\delta^e)\delta^m$. Following TD, there is implicit continual renegotiation each period, meaning agents understand that tomorrow's wage will be renegotiated independently of today's negotiation. Hence, in the bargaining game, both parties treat the future wage as predetermined and independent of today's bargaining outcome. V^U_{jc} is the value of unemployment for a worker formerly employed in a job of type j and is given by:

$$V^U_{jc} = b + \rho\left[q^u_c \sum_k T_{kc|j} V^E_{kc}(w'_{kc}) + (1-q^u_c)V^U_{jc}\right] \tag{5}$$

where, b is the flow value of being unemployed, $V^E_{kc}(w'_{kc})$ is the expected value of employment in job k and city c across firms, w'_{kc} is the average wage in job type k next period, and $T_{kc|j}$ is the probability a worker formerly in job type j in city c finds a job of type k, conditional on making any match. Based on our discussion of matching rates,

$$T_{kc|j} = \eta_{kc}\chi_{kc|j}(\varphi_{k|j}) = \eta_{kc}\frac{\varphi_{k|j}}{\sum_{k'} \eta_{k'c}\varphi_{k'|j}} \tag{6}$$

We do not model the source of the differences in $T_{kc|j}$ by k and j, treating them as exogenous facts from the workers' perspectives. Thus, this is a model of random search with probabilities of a worker meeting specific jobs given by $T_{kc|j}$.[12]

Equation (5) says that the value of unemployment is higher when b is higher, when the probability of making a match (q^u_c) is higher, and when the expected value of the match, $\sum_k T_{kc|j}V^E_{kc}$, is higher. In Appendix A.4, we show that in steady state, $\sum_k T_{kc|j}V^E_{kc}$ can be written as $\Gamma_c + \mathbb{B}_c \sum_k T_{kc|j} w_{kc}$, where $\Gamma_c > 0$ and $\mathbb{B}_c > 0$ are functions of the matching rates and model parameters, and w_{kc} are average wages across firms offering jobs of type k in city c. We will refer to $\sum_k T_{kc|j} w_{kc}$ as the 'outside option value' of the worker, though this is a slight abuse of terminology (since the full outside option includes Γ_c). Strictly speaking our estimated coefficients capture the effects of the expected wage part of the outside option alone (see Appendix B.3). This outside option is higher if the local economy has a greater concentration in high-wage jobs that the worker has a relatively high probability of matching with (i.e., with high associated $T_{kc|j}$ values).

[12] Our model, therefore, abstracts away from issues related to workers queueing for union jobs (Abowd and Farber [1982]). This queuing mechanism could imply an additional spillover channel whereby the existence of union firms drives down vacancy-filling rates in the non-union sector, pushing up wages. The prevalence of queueing is likely driven by union wage premia and the relative likelihood of finding union work such that queuing effects are likely to enter through the outside option channel in our model.

2.5 Wage Bargaining

Recall that the job subscript j combines union status $\tau = \{u, n, e\}$ and industry i. To make the notation more intuitive as we derive wages for Union, Simple Non-union, and Emulating Non-union firms, we sometimes replace j with specific subscripts: ui for Union firms, ni for Simple Non-union firms, and ei for Emulating Non-union firms in industry i.

2.5.1 Union Firms

Wages in Union firms are given by the solution to the Nash Bargaining condition:

$$\beta S_{fuic}(n) = (1-\beta)n[V^E_{fuic}(n) - V^U_{uic}], \tag{7}$$

where β is the bargaining weight. Following TD, in a unionised setting, the firm surplus equals:

$$S_{fuic}(n) = [\pi_{fuic}(n) + \rho^e \Pi_{fuic}(n)] - [\pi_{fuic}(0) + \rho^e \Pi_{fuic}(0)], \tag{8}$$

where $\pi_{fuic}(n)$ denotes current-period profits, $\Pi_{fuic}(n)$ is the value of the firm with n workers, and $\pi_{fuic}(0)$ and $\Pi_{fuic}(0)$ are the flow profits and value of the firm with no workers, respectively, reflecting the fact that if bargaining breaks down, the union will remove all the workers. At the point of bargaining, the number of workers in the firm is fixed and the hiring cost is sunk. For this reason, current period recruitment costs do not appear in (8).

In Appendix A.3, we show that for a Union firm:

$$S_{fuic}(n) = p_i y_{fuic}(n) - w_{fuic}(n)n + \rho^e(1-\delta^m)\frac{\kappa}{q^v_c}n \tag{9}$$

That is, it equals the profits lost from a shutdown plus the additional cost of hiring the entire optimal workforce in the following period.

On the right-hand side of (7) is the sum of workers' surplus, which is given by the gain to employment for all workers hired by the firm. Since the workers are identical, we use a specification that focuses on the total surplus and assume that the union members will all get an equal share of the part of the surplus captured by the union. This ignores issues related to seniority (see, e.g., Abraham and Medoff [1984, 1985] and Abraham and Farber [1988]). Note that workers' surplus will depend on the value of unemployment and, through that, on $\sum_k T_{kc|ui} w_{kc}$, the outside option value of the worker.

In Appendix A.5, we solve for the steady state wage written as a function of firm size, then solve for optimal firm size, substituting it into the bargained wage equation to arrive at our expression for the Union wage. That expression is a non-linear function of q^u_c and q^v_c, the matching rates for unemployed workers and vacancies, respectively. BGS show that in a steady state, these matching rates can be written as simple functions of the city employment rate, ER_c, and we substitute those functions. To get to our empirical specification, we linearize the resulting wage expression with respect to the vector, $\mathbf{x} = \{b, p_i, \sum_k T_{kc|ui} w_{kc}, ER_c, \epsilon_{fic}, \psi_{fuic}\}$. We take the linear approximation around a point $\mathbf{x_0}$ where employment is equally spread across industries (which occurs when the national mobility frictions are constant, i.e. when $\varphi_{k|j} = \varphi \ \forall k, j$) and the employment rate takes the same value in all cities (see Appendix A.7).

The linearized union wage expression is:

$$w_{fuic} = \tilde{\gamma}_{0i} + \tilde{\gamma}_1 \sum_k T_{kc|ui} w_{kc} + \tilde{\gamma}_2 ER_c + \tilde{\gamma}_3 \epsilon_{ic} + \tilde{\gamma}_3 u_{fic} - \tilde{\gamma}_4 \psi_{fuic}, \tag{10}$$

where $\tilde{\gamma}_{0i}$ is a function of the price, p_i, and constant terms stemming from the expansion point values. $\tilde{\gamma}_1$, $\tilde{\gamma}_2$, $\tilde{\gamma}_3$, and $\tilde{\gamma}_4$ are all positive. Expressions for each, written as functions of underlying parameter values, are given in Appendix A.7. Thus, Union wages are a positive function of productivity (captured in $\tilde{\gamma}_{0i}$, u_{fic}, and ϵ_{ic}), the workers' outside option value ($\sum_k T_{kc|ui} w_{kc}$), the tightness of the labour market, as reflected in ER_c, and a negative function of union amenities (ψ_{fuic}). As pointed out in BGS, in a frictionless environment, the wage would only be a function of productivity and the union amenity. In particular, the value of a worker's outside option would not play a role in wage determination.

2.5.2 Simple Non-union Firms

In Simple Non-union firms, the firm bargains with each worker individually, yielding wages that satisfy the bargaining rule:

$$\beta S_{fnic}(n) = (1-\beta)\left[V^E_{fnic}(n) - V^U_{nic}\right] \tag{11}$$

As with Union workers, the worker outside option is the value of unemployment, though the size of that option can differ because Union and Non-union workers have potentially different probabilities of accessing jobs of various types.[13] For firms, the fact that they are bargaining with one worker at a time means the firm surplus equals profits at the current firm size minus the profits the firm would attain if it lost this one worker plus the cost of having to hire one additional worker the following period. In Appendix A.3 (following TD), we show that this implies that:

$$S_{fnic}(n) \;\; = \;\; \frac{\partial \pi_{fnic}(n)}{\partial n} + \rho^e(1-\delta^m)\frac{\kappa}{q^v_c}, \tag{12}$$

where $\frac{\partial \pi_{fnic}(n)}{\partial n} = p_i \frac{\partial y_{fnic}(n)}{\partial n} - w_{fnic}(n) - n\frac{\partial w_{fnic}(n)}{\partial n}$.

Solving for the Simple Non-union wage in the same way as for the Union wage and again taking a linearization leads to our Simple Non-union wage expression:

$$w_{fnic} = \gamma_{0i} + \gamma_1 \sum_k T_{kc|ni} w_{kc} + \gamma_2 ER_c + \gamma_3 \epsilon_{ic} + \gamma_3 u_{fic} \tag{13}$$

As with the Union equation, γ_{0i} is a function of the price, p_i, and constant terms stemming from the expansion point values, and the other coefficients are all positive. Expressions for each of these coefficients are given in Appendix A.7. Importantly, $\tilde{\gamma}_{0i} > \gamma_{0i}$ and, so,

[13] A referee correctly pointed out that Union and Non-union workers may also have different job separation rates. Appendix B.2 shows that allowing for this implies that the coefficients in our linearized wage equations should differ between the Union and Simple Non-union wage equations. Since we estimate the Non-union and Union wage equations separately, our specification allows for such differences.

Union wages within an industry are on average higher, reflecting the fact that Union wages are proportional to total product while Simple Non-union wages are proportional to the marginal product of a worker and the latter is smaller if there are span of control issues. Intuitively, unions can bargain higher wages because they can threaten to withdraw the whole labour force, while a non-union worker can only threaten to withdraw her labour. In addition, $\tilde{\gamma}_3 > \gamma_3$, i.e., unions can capture a greater share of productivity shocks than non-union workers.

2.5.3 Emulating Non-union Firms

As discussed earlier, if the workers initially signal a desire to unionise through a straw vote, the firm may respond with a wage offer designed to induce the workers to vote against a union in subsequent straw votes. They do so in circumstances where unionisation would reduce the total size of the surplus but is still in the interest of the workers on their own. What follows is a sequence of offers culminating in a solution that is represented as a Nash bargaining solution. The outside options in this bargaining situation (what each side would revert to if bargaining breaks down) is the value of being in a union for the workers and operating as a unionised firm for the firm. This is the case because the workers have the legal right to form a union and doing so is in their best interests if there is no wage response from the firm. In contrast, in the Simple Non-union and Union cases, the outside options are the value of unemployment for workers and of production with fewer workers for firms.

We assume that in steady state, a firm knows if it will be an Emulating Non-union firm and will choose its number of workers to maximize profits conditional on paying the Emulating Non-Union wage. However, they also understand that if bargaining were to break down, one consequence would be operating as a Union firm with sub-optimal (union) employment, as firms are legally prohibited from laying off workers for organizing a union (though, in practice they sometimes violate these legal constraints [Bronfenbrenner, 2009]). In the model, because firms cannot adjust employment, their outside option reflects the cost of operating as a Union firm with sub-optimal employment and is part of the threat that workers hold over the firm. Although this scenario never occurs because it lies off the equilibrium path, bargaining takes place with that threat in the background.

Given these assumptions, the Emulating Non-union wage solves the bargaining problem:

$$\beta S_{feic}(n) = (1-\beta)n\left[V^E_{feic}(n) - (V^E_{fuic} - \lambda^*_c)\right], \tag{14}$$

where, the worker surplus on the right-hand side equals the difference between the value of being in an Emulating Non-union job relative to the value of being in a Union job minus the cost of unionising, and where n denotes the number of workers in the Emulating firm. The firm surplus, $S_{feic}(n)$, is given by:

$$S_{feic}(n) = \left[\pi_{feic}(n) + \rho^e \Pi_{feic}(n)\right] - \left[\pi^*_{fuic}(n) + \rho^e \Pi^*_{fuic}(n)\right] \tag{15}$$

where, the first term in the brackets reflects the discounted profits following successful bargaining. The second term in the brackets captures the discounted profits if the negotiation breaks down, i.e., if the firm is unionised (but employs n workers). In Appendix A.3, we show that $S_{feic}(n) = \frac{1}{1-\rho^e}\left[\pi_{feic}(n) - \pi^*_{fuic}(n)\right]$, i.e., the firm surplus is given by the difference in profits when operating as an Emulating Non-union firm versus as a Union firm.

Solving the bargaining problem yields the wage expression:

$$w_{feic} = w_{fuic} + \bar{\xi}\left[(\mathbb{A}\psi_{fuic} - \lambda_c^*) - \Delta_{eic,uic}\right], \tag{16}$$

where $\bar{\xi}$ and $\mathbb{A}$ are positive functions of model parameters and $\Delta_{eic,uic}$ is a function of differences in transition rates to other firms between non-union workers and union workers that drops out in the linearization step, with the specific forms of each given in Appendix A.5. This expression says that the Emulating Non-union wage equals the Union wage plus an adjustment that is positively related to any amenities the union would create and negatively related to the cost of unionisation. Substituting in our expression for the Union wage, w_{fuic}, and linearizing, we arrive at:

$$w_{feic} = \tilde{\gamma}_{0i} + \tilde{\gamma}_1 \sum_k T_{kc|ui} w_{kc} + \tilde{\gamma}_2 ER_c + \tilde{\gamma}_3 \epsilon_{ic} + \tilde{\gamma}_3 u_{fic} - (\tilde{\gamma}_4 - \bar{\xi}\mathbb{A})\psi_{fuic} - \bar{\xi}\lambda_c^*, \tag{17}$$

where $\tilde{\gamma}_4 - \bar{\xi}\mathbb{A} > 0$. Notice that the Emulating Non-union wage depends on $\sum_k T_{kc|ui} w_{kc}$, the outside option value for union workers since when that is higher (when, for example, there is a lot of high wage jobs accessible to union workers), Emulating firms are forced to pay a higher wage to prevent workers from unionising.

2.6 Union Arrangement Determination and Selection

After the firm hires its optimal workforce, union status is determined through the decision process shown in Figure 1. This process implies that two amenity thresholds, ψ^*_{fuic} and ψ^{**}_{fuic}, and one productivity threshold, ϵ^*_{fic}, jointly determine union status. Figure 2 summarizes the various union statuses in the amenity-productivity space. A full derivation of union status determination is provided in Appendix A.8.[14]

The first amenity threshold, ψ^*_{fuic}, is the point at which the employment value of working in a Union firm, net of the unionisation cost, equals the employment value if the firm remains Simple Non-union. Below ψ^*_{fuic}, workers willingly remain non-union; above it, they threaten to unionise. The second amenity threshold, ψ^{**}_{fuic}, is the point where the total surplus from the Emulating Non-union status (relative to Union status) equals zero. We show that ψ^{**}_{fuic} is equal to the unionisation cost, with firms entering Emulating Non-union status if $\psi^*_{fuic} < \psi_{fuic} \leq \lambda_c$ and being Union if $\psi_{fuic} > \lambda_c$. Finally, below the productivity threshold ϵ^*_{ic}, there will be no range of amenity values where the emulation wage is bargained. Since Union wages fall faster with productivity than Simple Non-union wages, when ϵ_{fic} is low enough, the net value of unionised employment is less than the value of Simple Non-union employment, leaving no room for an intermediate, Emulating Non-union wage to forestall unionisation. Hence, union firms are more common in cities with lower unionisation costs and among more productive firms, where unions can capture a larger share of the productivity shock, reflecting a standard selection mechanism.

[14] Recent decades have seen a high and rising level of illegal interventions by firms in union certification drives [Bronfenbrenner, 2009]. To simplify, the baseline model assumes no intimidation, i.e., workers face an exogenous cost of unionising, λ_c^*. In Appendix B.1, we consider the case where firms can pay to increase the cost of unionisation.

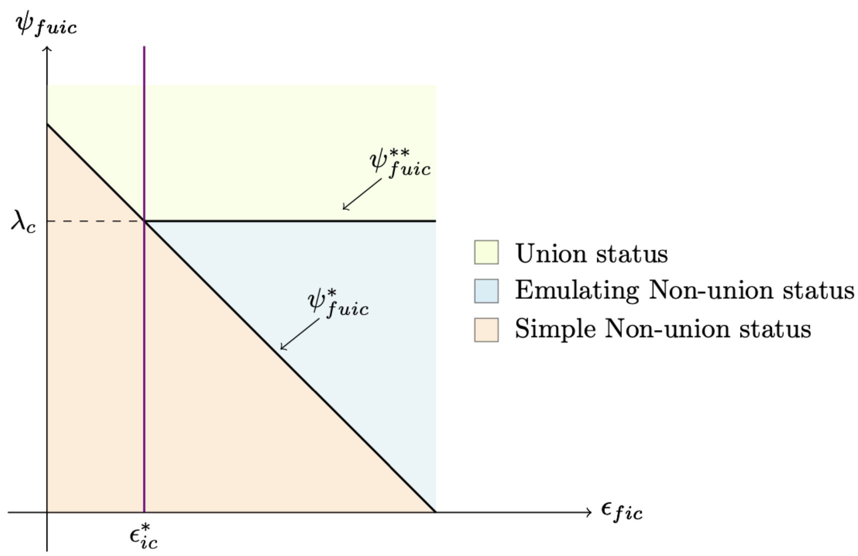

Figure 2: Statuses in the amenity-productivity space.

3 Empirical Specification

We are now in a position to present our empirical specification, which is set at the industry $\times$ city cell level. Importantly, in our data, we can only see whether a worker is union or non-union, not the type of non-union firm they work for. As a result, our observed dependent variable, the non-union wage, is a weighted average of the wages in Simple Non-union and Emulating Non-union firms:

$$w^n_{ict} = (1 - P^{ne}_{ict}) \cdot w_{nict} + P^{ne}_{ict} \cdot w_{eict}, \tag{18}$$

where: w^n_{ict} is the observed mean non-union wage in industry i in city c at time t; we have now introduced a time subscript; w_{nict} and w_{eict} are mean wages across Simple Non-union and Emulating Non-union firms, respectively; and, P^{ne}_{ict} is the probability a firm is an Emulating Non-union firm conditional on it being a non-union firm of either kind. P^{ne}_{ict} corresponds to the proportion of non-union firms that are under direct threat of unionisation, and hence, captures the union threat.

Differencing at the decadal level, thus eliminating any industry by city time-invariant characteristics, and dividing through by a base wage so that we are working with log wages, we obtain:[15]

$$\begin{aligned}\Delta \ln w^n_{ict} \quad = \quad & \underbrace{\Delta\gamma_{0it} + \gamma_1 \Delta[(1 - P^{ne}_{ict})E_{nict}] + \gamma_2 \Delta ER_{ct}}_{\text{A. Influencing factors in the absence of a firm response}} \\ & + \underbrace{\Delta(\gamma^*_{0it} P^{ne}_{ict}) + (\tilde{\gamma}_3 - \gamma_3)\epsilon_{ict} \Delta P^{ne}_{ict} + \tilde{\gamma}_1 \Delta(P^{ne}_{ict} E_{uict}) - \gamma_6 \Delta(P^{ne}_{ict}\lambda_{ct})}_{\text{B. Factors related to union threat changes}} + \bar{u}_{ict},\end{aligned} \tag{19}$$

[15]Differencing at the decadal level and dividing through by a constant base wage is equivalent to taking the (change in) a log-linear approximation of log wages, where the base wage corresponds to a constant term – specifically, the wage evaluated at x_0.

where, $\Delta x_t = x_t - x_{t-1}$ is a decadal difference, $\gamma^*_{0it} = (\tilde{\gamma}_{0it} - \gamma_{0it})$, and $\gamma_6 > 0$. We view the different time periods (which are a decade apart in our data) as corresponding to different steady states with different draws on productivity shocks, amenity values, and the cost of creating a union.

E_{nict} and E_{uict} are our shorthand notation for the outside option values for non-union and union workers, respectively. In particular,

$$E_{nict} = \sum_k \underbrace{\eta_{kct} \frac{\varphi_{kt|ni}}{\sum_{k'} \eta_{k'ct}\varphi_{k't|ni}}}_{=\eta_{kct}\chi_{kt|ni}=T_{kct|ni}} w_{kct} \tag{20}$$

with E_{uict} formed analogously.[16] The value of the outside option varies across cities and is a function of three determinants. The first is η_{kct}: the proportion of employment accounted for by a given job type, defined by the combination of union status and industry, in the city. The second is the wage rate paid in that job type in the city, w_{kct}. Importantly, to this point, we have assumed that workers are homogeneous, implying that wage differences across job types correspond to rents – differences in pay over and above what is required for the marginal worker to want to join that industry. Those rents are maintained because of the frictions in the labour market. It is important that we consider rents since wage differences across industries that correspond to compensating differentials or skill differentials cannot be the basis of bargaining a higher wage with your current employer. The third determinant driving the outside option value is the ease with which an employee from job type ni can transit to a job of type k, captured in $\varphi_{kt|ni}$. This includes the ease of moving from non-union to union jobs. Thus, for workers in job type ni, a city will have higher wages if it has more jobs that workers can actually access that pay high rents.

The last term in (19) is given by

$$\bar{u}_{ict} = \Delta\mu_{ict} + \gamma_3 \Delta\epsilon_{ict} + (\tilde{\gamma}_3 - \gamma_3) P^{ne}_{ict-1} \Delta\epsilon_{ict},$$

where μ_{ict} captures selection of firms into either type of Non-union status and the last two components reflect variations in sectoral productivity. One advantage of deriving our empirical specification from a model is that it allows us to see what is in the error term and what that implies for both endogeneity problems and solutions. Since productivity is unobserved, the second component in the second line of (19) is part of the error term, which is expressed as $\tilde{u}_{ict} = \bar{u}_{ict} + (\tilde{\gamma}_3 - \gamma_3)\epsilon_{ict}\Delta P^{ne}_{ict}$. The structure of this error term introduces a range of identification challenges, which we discuss in the next section. A complete derivation of this specification, including the form of μ_{ict}, is provided in Appendix A.9.

The two parts in equation (19) represent two sets of factors influencing the observed non-union wage:

Part A includes the factors that would determine non-union wages even in the absence of a firm response to potential unionisation (i.e., if $P^{ne}_{ict} = 0$). The first of these factors is movements in output prices, captured in a complete set of industry×time effects ($\Delta\gamma_{0it}$).

[16] We assume that hiring firms can distinguish between whether a worker's previous job was Union or Non-union but not whether it was Simple or Emulating Non-union. Given this, workers from both types of Non-union firms face the same outside option, E_{2ict}. We also show in Appendix A.9 that at reasonable values for the structural parameters, $\gamma_2 \approx \tilde{\gamma}_2$, so we do not include interactions of ER_{ct} with P^{ne}_{ict}.

Second, non-union wages are increasing in the value of the outside option for non-union workers, E_{nict}, fitting with results in, among others, Beaudry et al. [2012], Tschopp [2017], Jarosch et al. [2019], Caldwell and Danieli [2021], and Bassier [2022]. Third, the wage is predicted to increase with the employment rate, ER_c, since a tighter labour market implies that workers can access their alternative options more easily.

The inclusion of industry effects means that the relevant identifying variation for the estimated coefficients comes from across-city within-industry variation. Intuitively, we identify the impact of outside options by comparing wage changes in the same industry in two different cities that are experiencing different changes in the quality of outside employment prospects, holding the employment rate constant. For example, we could compare construction workers in Pittsburgh in the 1980s, when the decline of big steel meant a decline in the possibility of high-rent jobs, to construction workers in a city not substantially altering its sectoral composition, and would predict larger wage declines in Pittsburgh.

Part B contains factors related to changes in the union threat, as captured by movements in P^{ne}_{ict}. The impact of an increase in the threat probability will be higher when union wages (and, as a consequence, emulating firm wages) are higher, and the terms on the second line correspond to reasons why union wages might be higher than non-union wages. Union wages are higher, in part, because of the basic bargaining environment emphasized in the model – workers have more bargaining power when they organize. That means that union workers are able to capture a bigger proportion of rents, as reflected in the $\Delta(\gamma^*_{0it}P^{ne}_{ict})$ terms (which corresponds to union wages capturing a bigger proportion of price differences across industries, p_i) and $(\tilde{\gamma}_3 - \gamma_3)\epsilon_{ict}$ (which corresponds to unions capturing a larger share of industry×city specific rents, ϵ_{ict}). In addition, union wages are higher because union workers' outside options are better than those of non-union workers. In our data, union workers are much more likely to access union jobs than non-union workers, and, as a result, their outside option value is larger and moves differently from that of non-union workers. If the union workers' outside option value were a significant determinant of the non-union wage, we would view this as strong evidence in favour of the threat of unionisation affecting non-union wage setting. In contrast to these forces, the effects of increases in the threat probability are mitigated when unionising costs are higher (as reflected in the $\gamma_6\Delta(P^{ne}_{ict}\lambda_{ct})$ component).

To capture the $\Delta(\gamma^*_{0it}P^{ne}_{ict})$ and $\gamma_6\Delta(P^{ne}_{ict}\lambda_{ct})$ components of the specification, we include a complete set of interactions between P^{ne}_{ict-1} and industry×time effects, as well as between P^{ne}_{ict-1} and city×time effects. We use P^{ne}_{ict-1} rather than ΔP^{ne}_{ict} to avoid endogeneity issues.[17]

3.1 Implementation and Identification Challenges

We turn, next, to describing the set of challenges with taking our empirical specification to our data and how we solve them.

[17] We think of the interactions of the lagged proportion unionised with industry dummies as the equivalent of including a Bartik type variable that distributes national-level changes at the industry level to cities based on their initial levels of union activity at the local level.

3.1.1 Worker Heterogeneity

The first challenge comes from the fact that while workers are homogeneous in our model, they are heterogeneous in our data. Our response is to treat individuals as representing different bundles of efficiency units of work and to assume those bundles are perfect substitutes in production. We then interpret firm wages in the model as the cost per effective labour unit. Let effective labour units be $\exp(H_l'\beta_t + a_l)$, where H_l and a_l capture observable and unobservable skills of worker l, respectively. Adding industry, city and time subscripts, workers' observed non-union log wages, $\ln w_{lict}^n$, are given by:

$$\ln w_{lict}^n = H_{lt}'\beta_t + \ln w_{ict}^n + a_{lt}. \tag{21}$$

The $\ln w_{ict}^n$ values are our object of interest. To obtain a measure of these, we estimate (21), capturing $\ln w_{ict}^n$ as the coefficients on a complete set of job×city fixed effects under an assumption that a_l is orthogonal to job×city effects. Our specification of H_l includes a complete interaction of dummies for educational attainment, a quadratic in potential experience, and gender and race dummy variables. We estimate (21) using only non-union workers, separately by year. This allows for flexible changes in the returns to education and other observable characteristics over time. The estimated $\ln w_{ict}^n$ coefficients are regression adjusted mean wages and constitute the dependent variable in our regressions.

3.1.2 Endogeneity and Reflection Issues with the Outside Option Terms

To understand the nature of the identification problems we face with the outside option and threat terms in equation (19), it is useful to write the terms out in full. From Part A:

$$\Delta\left((1-P_{ict}^{ne})E_{nict}\right) \;=\; (1-P_{ict}^{ne})\cdot E_{nict} - \left(1-P_{ict-1}^{ne}\right)\cdot E_{nict-1} \tag{22}$$

Three identification issues arise from this variable. The first is a standard reflection problem: the wage in a given industry×city cell, ic, is determined by the wages in the other cells within the same city (via the outside option term E_{nict}), which, in turn, are determined by the wage in ic. The second is that the value of outside options is determined, in part, by the share of high or low-paying jobs in the employment structure in a city (the η_{kct}'s). We would expect the shares for given industry×city cells to increase when productivity in the cells increase, and since changes in those productivities (the $\Delta\epsilon_{ict}$'s) are in the error term, this creates a relationship between the error term and changes in the outside option terms. Third, we also expect P_{ict}^{ne} to be endogenous since the model predicts that unionisation probabilities are higher in more productive environments.

We approach these problems by constructing variables corresponding to E_{nict} and P_{ict}^{ne} which we argue are not endogenous and correct the reflection problem. These variables are then used to create our final instruments. We will argue, in doing this, that the E_{nict} and E_{nict-1} terms present different issues, and therefore, we construct distinct instrument components for each.

The wages at other jobs in the city (the w_{kct}'s in equation (20)) are the source of the reflection problem. Movements in those wages alter the value of the outside options and, hence, the wage for workers in job ni in city c. But w_{nict} is also part of the outside option for the other jobs in the city, making it impossible to tell if the movement in one wage has a causal effect on another. In a related model, BGS show that one can replace the w_{kct} values with national-level rents by job type (ν_{kt}) in the outside option expressions in a model-consistent manner, and we follow that path here. As described earlier, it is important that the outside option variables are constructed from rents. To meet that requirement with our data, we implement a similar exercise to constructing our dependent variable, regressing log wages on the same set of skill and demographic variables (H_{lt}) plus a complete set of job-type dummy variables. We refer to the coefficients on those dummy variables as job-specific (industry-union status) rents, ν_{kt}, and use them to form altered outside options that do not depend on local wage variation. Using ν_{kt} alters the outside option slightly to say that we predict wages will be higher for workers in job ni in city c if there are relatively high proportions of jobs in high-rent sectors (as defined at the national level) to which the worker has effective access.

In considering endogeneity issues, it is important to recall that our identifying variation is within industries and across cities. As a result, identification problems arise from cross-city elements of our right-hand side variables being correlated with cross-city differences in productivity. Referring back to equation (20), which defines E_{nict}, we can see that this means the job-to-job transition probabilities (the $\varphi_{kt|ni}$ terms) are not a source of concern, as they are defined at the national level and do not vary across cities. Instead, the endogeneity problem is located in the η_{kct} employment shares in the outside option expressions.

We respond to this issue using a Bartik-type instrument approach, leveraging start-of-period employment and national-level growth rates to predict the η_{kct} terms. In constructing our instruments, we also take a "leave-one-out" approach, dropping the job type defined by the combination n and industry i to ensure that the instrument does not derive its power from the very sector we are analysing. Combining the responses to the reflection and the endogeneity problem, we form:

$$\hat{E}_{nict} = \sum_{k \neq ni} \hat{\eta}_{kct} \frac{\varphi_{kt|ni}}{\sum_{k'} \hat{\eta}_{k'ct} \varphi_{k't|ni}} \nu_{kt} \tag{23}$$

where $\hat{\eta}_{kct}$ is the predicted end-of-period share of employment in city c that is in job type k. $\hat{\eta}_{kct}$ is constructed by multiplying employment in city c in $t-1$ in each job by the national level growth rate between $t-1$ and t for that job to obtain predicted end-of-period values, and then using those values to form employment shares. Note that $\hat{E}_{nict}$ gets its cross-city variation from variation in the η_{kct-1}s. Thus, the relevant identifying assumption is that changes in productivity (the key component in the error term) are independent of the city's industrial composition at the start of the period. BGS further show that this is equivalent to a random walk assumption for the ϵ_{ict} process, such that changes in ϵ_{ict} are independent of their values at the start of a period. The validity of this assumption may not be immediately obvious. In Section 5, we present results from a test of an over-identifying restriction arising from the model that supports the assumption.

Our expression for the $\Delta\left((1 - P^{ne}_{ict})E_{nict}\right)$ term also depends on the lagged value of the outside option for non-union workers, E_{nict-1}. This term is naturally a function of the employment shares in period $t-1$, η_{kct-1}, as well as the national level job-to-job transition rates and local wages across different job types in $t-1$. If our identifying assumption that the η_{kct-1}s are independent of the error term holds, then E_{nict-1} does not face an endogeneity problem. The reflection problem still exists, however, and we again address it by replacing the local job wages with national-level wage rents. Thus, we work with a variable:

$$\tilde{E}_{nict-1} = \sum_{k \neq ni} \eta_{kct-1} \frac{\varphi_{kt-1|ni}}{\sum_{k'} \eta_{k'ct-1} \varphi_{k't-1|ni}} \nu_{kt-1} \tag{24}$$

The last element of $\Delta\left((1-P^{ne}_{ict})E_{nict}\right)$ that we need to consider is P^{ne}_{ict} – the proportion of non-union firms that are under direct threat of unionisation. This variable also faces endogeneity issues. We view the unionisation threat facing a firm as dependent on four factors: the level of rents that unions are expected to capture (with, as shown in the model, higher rents implying a higher threat); the level of organizing activity by unions in its sector or locality; the demographic makeup of its workforce, as some groups may be more easily organized [Card et al., 2018]; and the regulatory environment. The first of these factors is the basis of the endogeneity issue, since changes in rents are positively associated with the productivity shocks in the error term. We address that issue by working with variables associated with the second and fourth factors to form the predicted probability, $\hat{P}^{ne}_{ict}$.

We form a proxy for the second of our factors determining the threat, the level of union organising, by working with changes in union organisation at the national-sectoral and city level.[18] Specifically, we capture these changes by constructing a leave-one-out measure of the decadal growth rate in elections per establishment at the city level (UA_{ct}) and at the national industry level (UA_{it}).[19] For the latter, the premise is that if national-level unions shift toward more activist leadership then election drives in the industries in which they operate will increase.[20] At the same time, organizational spillovers could create different recruiting environments across cities (Holmes [2006]).

We proxy for the regulatory environment using two variables: one for Right To Work states, RTW_{ct}, and another for whether the Republican Party controlled all three branches of the state legislature, R_{ct}, which we assign to cities based on the state in which they reside and average based on population shares for cities that cross state borders.

We approximate P^{ne}_{ict} as the number of firms successfully unionised divided by the number of non-union firms in the industry-city cell in the previous four years. This is based on the idea is that when that proportion is higher, the threat of unionisation is more present, and a larger proportion of firms have to pay a higher wage to emulate union wages and dissuade their workers from unionising.[21] We then create our predicted value of the threat probability by first regressing growth in the industry $\times$ city rate of union organizing success on UA_{ct},

[18] We cannot use changes in union organizing at the city$\times$industry level because we expect those to be related to $\Delta\epsilon_{ict}$ (the productivity shocks).

[19] Note that since we leave out the specific city, for UA_{ct}, and the specific industry, for UA_{it}, of an ic observation when constructing the instrument values, these variables actually take different values for each location and industry.

[20] For instance, when John Sweeney became president of the AFL-CIO in 1995, he pledged to increase unionisation drives, allocating \$20 million to 'organize at a pace and scale that is unprecedented' (cited in Bronfenbrenner [1997], p. 196).

[21] This follows directly from the model. As stated earlier in Section 2.6, for $\epsilon_{ic} > \epsilon^*_{ic}$, ψ^{**}_{fuic} – the threshold value of amenities above which firms are unionised – equals the cost of unionisation, λ_c. The threshold determining Emulating Non-union status, ψ^*_{fuic}, is also a monotonically increasing function of λ_c (see Appendix A.8). Thus, when more firms are unionised (i.e., ψ^{**}_{fuic} falls), the emulation threshold also falls, and the proportion of firms that are Emulating Non-union firms among all non-union firms rises.

UA_{it}, RTW_{ct}, R_{ct}, and the interaction of UA_{ct} and UA_{it}. The predicted values from this regression, $\hat{g}(P^{ne}_{ict})$, represent the estimated growth rate in the local sector of the threat of unionisation, driven by trends in union organizing efforts outside that local sector and modulated by local regulatory conditions. Using this, we form estimates: $\hat{P}^{ne}_{ict} = \hat{g}(P^{ne}_{ict}) \cdot P^{ne}_{ict-1}$

Substituting $\hat{E}_{nict}$, $\tilde{E}_{nict-1}$, and $\hat{P}^{ne}_{ict}$ for their relevant counterparts in equation (22), we arrive at our instrument for the key $\Delta\left((1 - P^{ne}_{ict})E_{nict}\right)$ variable:

$$IV_{nict} = \left(1 - \hat{P}^{ne}_{ict}\right) \cdot \hat{E}_{nict} - \left(1 - P^{ne}_{ict-1}\right) \cdot \tilde{E}_{nict-1}, \tag{25}$$

where our identifying assumption implies that P^{ne}_{ict-1} does not face endogeneity issues. Using the same logic, we form an analogous instrument related to $\Delta(P^{ne}_{ict}E_{uict})$ as:

$$IV_{uict} = \hat{P}^{ne}_{ict} \cdot \hat{E}_{uict} - P^{ne}_{ict-1} \cdot \tilde{E}_{uict-1}. \tag{26}$$

Under our identifying assumption, these instruments allow for consistent estimation of the coefficients in our main specification.

3.1.3 Endogeneity of the Employment Rate

We also expect productivity changes in the error term to be related to labour market tightness, which we capture with the ER_{ct} variable in our regression. However, we do not instrument for ΔER_{ct}. We follow Stock and Watson [2011] in interpreting the employment rate as a control variable – a variable that is not of direct interest in its own right but is useful for picking up its own effect and those of correlated omitted variables. In our case, we view the employment rate as capturing its own effect plus the impact of general, local demand shifts. This allows us to isolate the outside option effects we care about from general demand effects, thus strengthening our claims for identifying the former.[22]

3.1.4 Selection into Non-union Status

As mentioned in Section 2.6 and detailed in Appendix A.8, the endogenous determination of union status involves a classic selection problem. Specifically, as can be seen in Figure 2, a change in the cost of unionisation, λ_c (which alters the values of ψ^*_{fuic} and ψ^{**}_{fuic}), will change the conditional distribution of productivity and firm amenities for firms observed to be non-union. The conditional means of those variables are captured in the μ_{ict} term in the error term in equation (19). Since a change in λ_c will also change P^{ne}_{ict} and, with it, the outside option term values, the change in firm composition will be picked up in the coefficients on the outside option terms – a classic selection problem.

[22] The control variable argument, in our case, implies that the required identifying assumption is that $\tilde{u}_{ict}$ is conditionally mean independent of our IVs, i.e., that the instruments are independent of the error term once we condition on the control variable ($E(\tilde{u}_{ict}|IV, \Delta ER_{ct}) = E(\tilde{u}_{ict}|\Delta ER_{ct})$). Stock and Watson [2011] show that if this condition is met, then the coefficients on $\Delta((1 - P_{ict})^{ne}E_{nict})$ and $\Delta(P^{ne}_{ict}E_{uict})$ are consistent for the causal effects of the outside option terms on $\Delta \ln w^n_{ict}$ while the coefficient on ΔER_{ct} does not have a causal interpretation. They also show that standard IV inference results, such as weak instrument tests, are valid under the conditional mean independence assumption.

To address this, we apply a generalized Heckman two-step approach [Heckman, 1979, Snoddy, 2019], which corrects for omitted variable bias by including a control function that corresponds to μ_{ict}. Since μ_{ict} can be expressed as a non-linear function of the probability of selection (i.e., the probability of being non-union), the control function can be a polynomial in this probability or in the exogenous variables that determine it.

We examine potential selection effects using two sets of variables. First, we include a quadratic in ΔP_{ict}, the change in the proportion of unionised workers in the industry×city×time cell. By doing so, we take the model very seriously in the sense that it says that access to union jobs affects non-union wages only through the transition rates in the outside option terms. The union proportion does not directly affect wage setting and, so, its inclusion can be interpreted as capturing selection effects. Note that this uses industry×city×time variation to identify the control function effect. Following Fortin et al. [2019], we also estimate specifications in which we proxy for costs of unionisation using NLRB data on certification elections as a robustness check. Details of this approach, including the construction of the selection variables, are provided in Appendix E.

3.1.5 Rent Capture Term

Because we do not have direct observations of the sectoral rents, we cannot form the variable corresponding to the extra rent capture term by unions, $(\tilde{\gamma}_3 - \gamma_3)\epsilon_{ict}\Delta P^{ne}_{ict}$. That variable, then, becomes part of the error term, and its effect will be reflected in the estimated coefficients on the right-hand side variables according to a standard omitted variables bias argument. More specifically, the estimated coefficients will capture not only the outside option channel but also the influence of the rent capture channel. Given that our interest lies in estimating the total effect of de-unionisation on changes in non-union wages, the inclusion of at least part of the rent capture term in our estimated effects is not necessarily a problem. However, it also complicates the decomposition of the effect into its components, as each estimated coefficient reflects some influence from the rent capture channel. In Appendix F, using simulations, we show that the rent capture component is negligibly small for the primary coefficients of interest (e.g., $\hat{\gamma}_1$), amounting to less than 0.001 in magnitude. Consequently, our estimates largely reflect the outside option effects.

Our empirical specification, as set out to this point, rests on three different unionisation variables, with the theory indicating a different role for each. The first is the probability that workers who are switching jobs can move into a union job (captured in the $\varphi_{kt|ni}$s). The second is the proportion of workers who are unionised, and the third is the probability a firm will face a successful union campaign. Each represents a specific way de-unionisation affects observed non-union wages (through outside option values for the workers, selection effects, and the threat of unionisation for the firm, respectively) and our empirical specification uses them in different ways. All three are, naturally, related, but we show in Appendix D that the transition probabilities and unionisation success have separate identifying variation relative to each other (i.e., that the transition probabilities are correlated with non-union wage movements even conditioning on unionisation success rates and vice versa.)

4 Data and Descriptive Patterns

Our analysis uses data from the Current Population Survey Merged Outgoing Rotation Groups for 1983-2019 and the CPS May extracts for 1978-1982. We are interested in comparisons across steady states over a medium-long time horizon, and, as such, we consider 10 year differences. We pool observations across 3 years for each period to reduce statistical noise. We consider variation across 1980, 1990, 2000, 2010, and 2020 using the years 1978-1980, 1988-1990, 1998-2000, 2008-2010, and 2018-2019.

From this data, we keep all workers between the ages of 20-65 who do not report being in school either full-time or part-time. We follow Lemieux [2006] in constructing our wage data, working with weekly wages. We use an aggregated grouping of industry codes based on the 1980 industrial classification from the Census Bureau. We obtain a consistent industry classification using crosswalks provided by IPUMS and the Census Bureau that map the 1970, 1990, and 2000 industry codes to the 1980 classification. The result is a consistent classification system with 51 industries. Appendix C contains additional processing details.

We construct a set of cities with as consistent geographic boundaries as possible, given data limitations in the CPS. We are constrained by the number of SMSAs available in the May extract data and end up with 43 cities. Making use of the limited number of counties identified in the CPS, we can create a set of cities which are reasonably, though not always perfectly, consistent over time.[23] The final geographic definition we use pools data for these 43 cities and the remaining population. Specifically, we create additional regions comprising the remaining state population absent the population living in these 43 cities. In the end, our core geographic measure is composed of 93 areas that are fairly consistently defined over the course of the sample period.

Additionally, we use NLRB case data for the sets of three years associated with each of our decadal points to construct probabilities of firms facing successful union certification drives.[24] We use the county of the unit involved in the election to construct our geographic measures, aggregating counties to our city definition discussed above. In particular, we calculate the proportion of firms in an ic cell that experienced a successful unionisation drive in the previous 4 years. We view those probabilities as proxies for the proportion of non-union firms that are emulating non-union firms based on the idea that when more firms are being unionised, the threat of unionisation for the remaining non-union firms is greater. Unfortunately, the election data ends before 2020 and, because there are no establishments in the public sector, we cannot generate the unionisation drive variable. As a result, we estimate the full model over the years 1980, 1990, 2000 and 2010 and only for the private sector. We do, however, use public sector wages as part of the outside options.

Central to our empirical work are the outside option terms characterising alternative job prospects in either the union or non-union sectors. As defined above, these terms are composed of the rents a worker would get in expectation when searching for a new job

[23] The metropolitan area definition used by the IPUMS identifies a general pattern of expanding metropolitan area definitions over time that we overcome to some extent, but not perfectly: `https://usa.ipums.org/usa/volii/county_comp2b.shtml`. Estimation using states as the geographic unit yields very similar results, suggesting that issues related to geographic definitions are not driving our results.

[24] We are grateful to Hank Farber for providing this data. We use data on certification elections in which a conclusive decision on certification was reached.

and are functions of the average wage rent paid in each possible job by city cell (w_{kct}), the proportion of workers in each cell in the city (η_{kct}), and the term that captures the difficulty with which a worker in a job of type j can move to a job of any other type k ($\varphi_{k|j}$). For the rent component, we use our regression-adjusted wages in order to get as close as possible to rents rather than skill differentials since wage differences that reflect skill differentials cannot be used as an outside option in bargaining (a janitor cannot use the opening of new jobs for lawyers in town to bargain a better wage).

We compute the η_{kct}s (the proportion of employment in city c that is in job k) directly from the CPS data. We proxy the $\varphi_{k|j}$ (the probability a worker in job type j moves to k) terms with transition probabilities at the national level, estimated using the matched CPS. In particular, we calculate the proportion of workers in a given cell, j, in year t observed in each possible other cell in year $t+1$.[25] We do this for each of the three CPS years at each decade point (e.g., initial years 1988 – 1990 for the 1990 observation) and average over those three years. This is done separately at each decade point, allowing for changes in transitions over time.[26]

Before turning to estimation, it is important to highlight key patterns in unionisation over the 1980–2019 period. Nationally, unionisation rates declined from 25% in 1980 to 13% in 2019. The decline was particularly steep in cities like Detroit, Gary, and Pittsburgh, where unions historically played a major role, and smaller in cities like Dallas, with lower initial unionisation rates. This variation, both across cities and over time, provides valuable identifying variation for disentangling the effects of union declines from broader trends. Additional details and figures, including city-level changes and their implications for nonunion wages through bargaining and threat channels, are available in Appendix D.

5 Estimation Results

Non-Union Wage Results: Table 1 presents the results from specification (19). The dependent variable—the decadal change in non-union wages—is adjusted for education, age, gender, and race, with standard errors clustered at the city×year level.[27] To ensure

[25] We calculate the proportions using only workers who were employed at both survey points. We discuss the implications of this sample restriction in Appendix B.3. Also, our framework assumes that bargaining effects operate only through the unemployment channel, that is, workers must first transition through unemployment to access other jobs. However, due to data limitations, our transition measures use transitions between sectors, which may or may not have included an intervening unemployment spell. Thus, the union outside option term may reflect on-the-job search dynamics. Formally modeling on-the-job search, or job laddering, is beyond the scope of this paper. As noted by Beaudry et al. [2012], it is not straightforward and is sensitive to the modeling of the search process and its relationship to wage determination.

[26] One complication in this is workers observed in the same cell in years t and $t+1$ since we cannot observe whether they have moved to a different firm in the same cell. To the extent they do, the wage in that cell is part of their outside option with their current firm. We estimate the proportion of workers making such a transition by calculating the proportion of workers who are observed in the same cell in both years but have different values of a set of job characteristics, including how they are paid (hourly versus not hourly), worker class (private versus public), and sub-industry.

[27] In Appendix I, we discuss recent papers on clustering and standard errors using Bartik Instruments. The proposed approach in Borusyak et al. [2022], in which data is aggregated to the level of shocks, is not possible in our case where the Bartik instruments vary with each jct cell. Adao et al. [2020] argue that standard

robustness, we exclude industry×union status×city cells with fewer than 10 observations and weight observations by the square root of the cell size to avoid small cells disproportionately influencing the results.

Table 1: Non-Union Wages and Outside Options

	Non-Union				Union	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nict}\right)$	0.65***	0.66***		0.65***		
	(0.11)	(0.095)		(0.10)		
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nct\|ni}\right)$			0.64***			
			(0.12)			
$\Delta\left((1-P^{ne}_{ict})\cdot E_{uct\|ni}\right)$			0.65***			
			(0.11)			
$\Delta\left(P^{ne}_{ict}\cdot E_{uict}\right)$	0.77***	0.74***	0.77***	0.77***		
	(0.26)	(0.27)	(0.26)	(0.26)		
ΔE_{ucit}					0.41***	0.44***
					(0.15)	(0.15)
ΔER_c	0.36		0.38	0.36*	0.25	0.22
	(0.22)		(0.24)	(0.21)	(0.30)	(0.30)
Obs.	5960	5960	5960	5960	1661	1661
Year × Ind.	Yes	Yes	Yes	Yes	Yes	Yes
P^{ne}_{ict-1}× Ind. × Year	Yes	Yes	Yes	Yes		
P^{ne}_{ict-1}× City × Year	Yes	Yes	Yes	Yes		
Selection controls						
ΔP_{ict} Quadratic	Yes	Yes	Yes		Yes	
First-Stage p-Stat.:						
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nict}\right)$	0.000	0.000		0.000		
$\Delta\left(P^{ne}_{ict}\cdot E_{uict}\right)$	0.000	0.000		0.000		
$\Delta\left((1-P^{ne}_{ict})\cdot E_{uct\|ni}\right)$			0.000			
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nct\|ni}\right)$			0.000			
ΔE_{ucit}					0.000	0.000
Over-id. p-val			0.824			

Notes: This table displays results from the estimation of equation (19) via 2SLS. The dependent variable is the decadal change in the regression adjusted average hourly wage of non-union workers in an industry-city cell, using CPS data from 1980-2010 across 50 industries and 93 cities. Standard errors, in parentheses, are clustered at the city-year level.

Column (1) contains the coefficients from our full specification. Based on equation (19) and the arguments that follow it, our specification has, as its key covariates, the bargaining and threat channel variables, $\Delta\left((1-P^{ne}_{ict})E_{nict}\right)$ and $\Delta(P^{ne}_{ict}E_{uict})$, as well as the change in the employment rate at the city level. Based on our argument that selection is a potential problem, we also include, as a control function, a quadratic in the change in unionisation

errors with Bartik instruments face a clustering problem because of correlations across observations with a similar base period composition of the shock exposure shares. Since that would correspond to industry shares in our case and we already control for time-varying industry effects, we do not face this issue.

proportions in ic cells. Following our theory, our specification includes a complete set of industry×time effects and interactions of P^{ne}_{ict-1} with both a complete set of industry×time effects and a complete set of city×time effects. We instrument using our IV_{nict} and IV_{uict} variables. We report the First-stage p-values for the Sanderson-Windmeijer test statistics for weak instruments [Sanderson and Windmeijer, 2016] at the bottom of the table. They are less than 0.001 in all cases, indicating that we do not face weak instrument problems.

The key first takeaway from our estimates is that unionisation in a local economy does affect non-union wage setting, and it does so through both the bargaining and the threat channels. The coefficients on the variables corresponding to both channels are statistically significant at the 1% level of significance. If the coefficients on each variable were equal (a restriction we cannot reject at standard significance levels), our estimates imply that a \$1 increase in the outside option value of both non-union and union workers would generate a \$2.45 increase in the mean non-union wage. The fact that the effect is greater than 1 fits with arguments in BGS that such spillover effects can be large. For example, the closure of even one large union firm reduces the outside options of all non-union workers in the city simultaneously. The resulting reduction in non-union wages then serves to further reduce outside options, leading to further wage losses.[28] Given this, the widespread decline in the union sector has the potential to reduce non-union wages substantially – a point we return to in our decomposition exercise in Section 6.

The significant effect of non-union worker outside options on non-union wages aligns with results in earlier papers. The significant effect of the union worker outside option is, perhaps, less expected, and we view it as evidence in support of our model. Its effect is identified relative to the non-union worker outside option because of differential changes in transition rates for union and non-union workers. In Figure 3, we present the mean probabilities, separately, that union and non-union workers transit to a union job by the following year for each of our sample years using the national level data (i.e., mean values of $\varphi_{ui't|ui}$ and $\varphi_{ui't|ni}$). These show a strong decline in the probability of accessing a union job for non-union workers (from 0.24 in 1980, to 0.091 in 2000, and 0.07 in 2020) but higher levels that do not decline as fast for union workers (where the probability is 0.275 in 1980, 0.197 in 2000, and 0.16 in 2020). The impact of these differences on local outside option values is mediated through their interactions with changes in local industrial composition (the η_{kct}s) and changes in wage premia for different job types, or, for our instruments, changes in national level job rents (the ν_{kt}s). It is the variation in our instruments, IV_{nict} and IV_{uict}, that is most relevant for our identification, and the differences in transition rates underlying each translate into a correlation between the instruments of only 0.16 across industry×city cells.[29] In Appendix G, we present results from a quasi-reduced form specification in which we regress $\Delta \ln w^n_{ict}$ on ΔE_{uict}, ΔE_{nict}, and ΔP^{ne}_{ict} separately, including all the same controls

[28] More specifically, γ_1 is the initial impact of a \$1 change in the outside option value on the nonunion wage. However, the change in the wage will lead to a change in the outside option value for others, altering their wages, and so on. The sum of the effects across all rounds is $\frac{\gamma_1}{1-\gamma_1}$.

[29] In comparison, the correlation between changes in the outside option variables, ΔE_{uict} and ΔE_{nict}, is 0.86. This is much larger than the correlation between the instruments because the outside option values use the local wages, w_{uict} and w_{nict} and the local changes in the job type shares. Since these tend to move together at the local level in a way that the national level ν_{kt}s and the start of period η_{kct}s do not, the outside options are much more correlated than the instruments that actually generate our estimated effects.

and using the same instruments as for equation (19). We find that all three elements enter significantly, supporting the argument that our instruments for the two outside option values have identifying variation relative to each other.

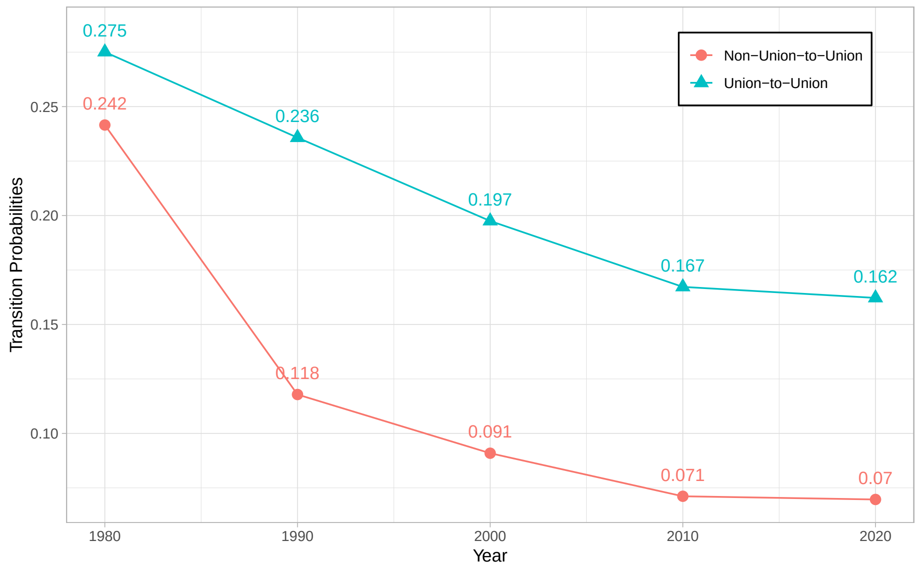

Figure 3: Transitions to Union Jobs

Notes: This figure reports transition probabilities for non-union and union workers into union jobs. The transition probabilities exclude same-job transitions for union workers. The data comes from matched CPS data, described in the main text.

In the remaining columns of the left panel of Table 1, we present variations on our main specification. In column (2), we present results in which we drop the ΔER_{ct} control. As described earlier, we derived our model under partial equilibrium assumptions, including treating labour market tightness as fixed. However, de-unionisation could affect labour market tightness if, for example, firms that de-unionise face lower wage costs and, as a result, post more vacancies. By not controlling for changes in the employment rate, we allow any such effects to show up in the estimated outside option coefficients – though at the cost of using a specification that is not strictly interpretable under our theory. The estimated coefficient on the outside option value is very similar to what we obtain in the previous column, where we control for ΔER_{ct}. This suggests that the indirect effects of de-unionisation through labour market tightness are unlikely to be large.

In Section 3, we described the identifying assumption underlying our instruments as being that changes in industry by city-level productivity are unrelated to the initial industrial composition of the city. The model provides a natural over-identification test of this assumption. In particular, we can write E_{nict} as,

$$E_{nict} = \sum_k T_{kct|ni} w_{kct} = \underbrace{\sum_{i'} T_{ui'ct|ni} w_{ui'ct}}_{\text{Union component}} + \underbrace{\sum_{i'} T_{ni'ct|ni} w_{ni'ct}}_{\text{Non-union component}}, \tag{27}$$

where the first component of the outside option is associated with potential union jobs and the second term with potential non-union jobs. We will refer to the first component in equation (27) as $E_{uct|ni}$ and the second component as $E_{nct|ni}$. In column (3), we present IV results from a specification in which we include them separately (constructing instruments related to each analogously to IV_{nict}). Our theory says that the two components should have an equal effect on bargained wages since it does not matter to the employer in what specific sector a worker's improved outside options arise. That is what we find – with the two estimated coefficients being almost identical. Importantly, the associated instruments, which we call $IV_{uct|ni}$ and $IV_{nct|ni}$ contain different employment share values $\eta_{\tau i'ct-1}$ (for union versus non-union jobs) and multiply them by different transition rates. Thus, if the identifying assumption that the $\eta_{\tau i'ct-1}$ values are independent of the error terms is incorrect, then the two instruments should weight the offending correlation of the $\eta_{\tau i'ct-1}$s with the error term differently and we should obtain different estimates of the outside option effect depending on which instrument we use. We report the test for the restriction that this is not the case (a standard Hansen's J test performed on the column (1) specification using $IV_{uct|ni}$ and $IV_{nct|ni}$) at the bottom of column (3). The test statistic does not reject the null hypothesis (with a p-value of 0.824). Together, we interpret these results as being supportive of our identifying assumption.

In column (4), we drop the control function, with little impact on our outside option channel coefficients. Thus, while the existence of selection effects is plausible, we find little evidence that they actually affect our estimates.[30]

Union Wage Results: In columns (5) and (6), we present the results from estimating the union wage specification (10), with and without the control function.[31] Due to the significant decline in unionisation over time, we lose approximately two-thirds of our industry×city cells when focusing on union wages. Additionally, we exclude public sector jobs, as our model may not apply well to wage setting in this sector, which further reduces our sample size. We do, however, continue to use public sector jobs as part of the outside options.[32] We find significant bargaining effects related to outside options for union workers in both columns (5) and (6). Interestingly, a referee pointed out that the theory predicts that the coefficient on the outside option term should be smaller for the union than the non-union wage equation, which is what we find (see Appendix A.5).

5.1 Heterogeneity in Spillover Effects

In Table 2, we present estimates of the bargaining and standard threat effect coefficients (γ_1 and $\tilde{\gamma}_1$ in equation (19)) for a set of sub-populations defined by gender, age, and education, following on evidence that there is considerable heterogeneity in experiences with

[30] A test of the hypothesis that the parameters in the quadratic equal zero is not rejected at any standard significance level, and the estimates for the key covariates change very little from column (2). As a robustness check, Appendix H reports results using NLRB certification election data to proxy for unionisation costs, which similarly has minimal impact on our estimates.

[31] We discuss estimation challenges and instruments in Appendix H.1.

[32] Including public sector jobs in the union specification produces estimates with the expected sign but with low precision. We have also estimated both our full, non-union specification and the union specification dropping the public sector from the outside option computation with little impact on the estimated coefficients.

Table 2: Subsample Analysis - Coefficient Estimates on Outside Options

	(1)	(2)	(3)	(4)	(5)
	Coefficient			1980	
Sample	γ_1	$\tilde{\gamma}_1$	N	Union Prop.	Union Prem.
Men	0.54** (0.15)	0.71** (0.22)	4551	0.32	0.14
Women	0.66** (0.08)	0.75** (0.10)	3702	0.18	0.17
Age 20–35	0.59** (0.11)	0.64** (0.16)	4032	0.23	0.17
Age 36–55	0.49** (0.18)	0.61** (0.20)	4133	0.29	0.13
$\leq$ HS	0.45** (0.10)	0.61** (0.14)	4081	0.30	0.18
$>$ HS	0.39** (0.19)	0.42* (0.22)	4293	0.21	0.11
Men Young/Low skill	0.62** (0.13)	0.55* (0.29)	2648	0.35	0.20
Men Young/High Skill	–0.06 (0.70)	0.03 (0.74)	2449	0.20	0.08
Men Old/Low Skill	0.56** (0.12)	0.98** (0.30)	2919	0.41	0.14
Men Old/High Skill	0.06 (0.47)	0.23 (0.44)	2720	0.21	0.00

Notes: This table displays results from the estimation of equation (19) via 2SLS on separate subsamples. All first-stage p-values for tests of instrument relevance are below 0.03. The last two columns show each group's 1980 unionization rate and union wage premium with the latter adjusted for worker characteristics.

unionisation (Farber et al. [2021], Card et al. [2018]). Each row corresponds to estimates for a different sub-sample. We calculate the transition rates from any job type to any other job type, $\varphi_{kt|j}$, for the specific population being examined and, based on those transition rates, calculate outside option values for each sub-sample. In all cases, the p-values from SW weak instrument tests for the instruments corresponding to the two outside option terms are 0.03 or less, implying the absence of weak instrument problems.

The first two rows contain separate results for men and women. These indicate that both the bargaining and standard threat effects are similar in size for men and women, though the bargaining effect is slightly larger for women. The following rows show that both types of effects are stronger for younger workers and for people whose highest level of education is high school graduation or less. In the last four rows, we delve deeper into skill-related differences for males, using an approach from Card [2009] for creating skill groups. In this method, weights are generated for each person that correspond to their contribution to four groups: young, low educated; young, high educated; old, low educated; and old, high educated.[33] We focus on men since they suffered the largest declines in unionisation. The last two columns of these rows show that the young/low-skilled and old/low-skilled men had

[33] In particular, people are assigned an age weight for each of two categories – young (with the weight generated from a quadratic kernel centred on age 27.5 with a 20-year bandwidth) and old (using a quadratic kernel centred on age 50 with a 20-year bandwidth). They are also assigned a weight for the low-educated group and for the high-educated group using Card [2009]'s efficiency weights. The low-educated group puts a weight of 1 on high school graduates and smaller weights on adjacent education categories, while the high-educated group puts a weight of 1 on those with a BA. The four skill groups are formed by multiplying the weights for the age groups with the weights for the education groups.

particularly large values of the union wage premium and the proportion unionised. Thus, these are groups where we would expect both the union threat and bargaining spillover effects to be particularly large and, indeed, the estimated effects are large relative to other groups – particularly more skilled workers.

6 Counterfactual Exercise

Our results thus far indicate a significant relationship between the quality of job opportunities in both the non-union and union sectors and non-union wage setting. However, the exact magnitude of the estimated effects remains unclear. In this section, we pursue a counterfactual exercise, asking what path mean wages in a typical city would have followed if unionisation rates and union wage premia had remained at their 1980 levels. This both provides a way of characterizing the size of our estimated effects and some insight into whether de-unionisation played an important role in wage changes over the last four decades.

6.1 Loss of Union Power and Movements in the Average Wage

Our focus is on changes in total mean wages at the city level, expressed as the weighted average of non-union and union mean wages, with the weight being the proportion unionised at the city level, P^u_{ct}:

$$w_{ct} = P^u_{ct} \cdot w^u_{ct} + (1 - P^u_{ct}) \cdot w^n_{ct}, \tag{28}$$

where w^u_{ct} is the mean log union wage and w^n_{ct} is the mean log non-union wage in city c at time t. We use residualized industry-city wages from our regressions (to abstract from the confounding effects of changes in education, age, and other factors), combined with local industrial shares, to create city-level wages.[34]

Changes in union strength affect average city wages through four channels:[35]

1. Union Proportion (P^u_{ct}): This is the most direct effect, representing the shift from higher-paid union jobs to lower-paid non-union jobs, holding sector wages constant. This is the "between" component in standard decompositions.

2. Probability of a Non-union Firm Being Unionised (P^{ne}_{ict}): This captures part of the classic threat effect, representing the changes in the likelihood of a non-union firm becoming unionised.

3. Probability of Finding a Union Job: Changes in transition rates, $T_{kct|j}$ –which combine how changes in mobility frictions, $\varphi_{kt|j}$, and job shares, η_{kct} impact outside options, affecting wages through the bargaining and emulation channels. When linking changes

[34] We set the wage level to correspond to the mean wage across all worker types. In particular, mean wages correspond to the wages of white workers, holding the proportion of education×gender groups at their 1980 levels.

[35] A fifth channel, selection effects, could theoretically increase observed non-union wages by changing the productivity composition of non-union firms. However, we find no substantial evidence of this effect, so it is not included in our decompositions.

in transition rates to the decline in union power, we do not want to attribute all of the changes in job shares to union effects. Instead, we assume that shifts in the industrial distribution for non-union workers capture changes in the overall economy, while a change in the industrial distribution for union workers relative to what happens for non-union workers is a union decline effect. We denote these relative job shares as η^*_{uict} (the difference between the actual growth in the industrial share in industry i in the union sector and what would have happened if it had grown at the non-union rate) and transition rates using these shares as $T^*_{kct|j}$.

4. Union Wage Premium: Declines in union bargaining power could reduce the union wage premium $(w^u_{ct} - w^n_{ct})$, lowering the value of the outside option of finding a union job. This could arise because unions become less effective at unifying worker resistance during bargaining or become afraid to threaten the withdrawal of the whole workforce in a new policy environment.

Figure 4 plots the percentage change in these key drivers relative to their 1980 values, aggregated across cities using city populations as weights. Thus, the trends shown depict the movements of each component for an average city. The trend in the probability of unionisation (P^u_{ct}) is labelled as 'Proportion Union' in the figure. In the line labelled 'Transitions' we present the movement in the national level probabilities of a non-union worker in any industry transiting to a union job in any industry $(\varphi_{ui't|ni})$, averaged across industries. We present this series rather than the local transition probabilities $(T^*_{kct|j})$ to provide an unadulterated look at the main driving force in the transition rates. This force is obviously related to changes in P^u_{ct}, though one could imagine that it could decline faster than the overall union proportion (if older union workers keep their jobs but new job searchers have difficulty getting into a union job) or slower (if the proportion declines quickly because union workers suddenly start taking early retirement). In fact, the figure shows that the two proportions fell since the 1980s, but the probability of entering a union job declines faster. Notably, the probability a non-union firm is successfully unionised (P^{ne}_{ict}), labelled as 'Threat' in the figure, fell the fastest of any of the unionisation measures, particularly in the 1980s when the policy environment was strongly against unionisation.[36]

Perhaps the most interesting line in Figure 4 corresponds to the union wage premium, labelled as 'Union Premium'. The premium actually increases in the 1980s before showing a sizeable decline in the 1990s and a smaller one thereafter. Both Card [2001] and Farber et al. [2021] have highlighted the seemingly odd result: the union wage premium did not decline during the 1980s when union power fell substantially.[37] Our model (echoing an argument in Farber [2005]) provides an explanation for the increase in the premium in the 1980s in our data and, potentially, the longer-term stability in the premium demonstrated in Farber et al. [2021] based on the threat channel. Recall that the observed mean non-union wage equals a weighted average of the Simple Non-union wage (w_{nict}) and the Emulating wage

[36] It is worth noting that the probability of a firm facing a union election was small even in 1980 (on the order of 4%).

[37] Farber et al. [2021] plot union wage premiums over an extended time period. Their plot differs from ours in showing a flat premium over the 1980s but is similar in showing a decline after 1990. Their estimates are based on family income and do not include controls for education that are part of our estimation.

(w_{eict}). The weights are $(1 - P^{ne}_{ict})$ and P^{ne}_{ict}, respectively. Suppose that larger forces (e.g., trade, technological change) drive down both w_{nict} and the union wage, w_{uict}, to the same extent. If the threat of unionisation declines simultaneously, the observed non-union wage will fall further because there will be fewer emulating firms, and the emulation wage they have to pay will not be as high. This pattern of faster decline in mean observed wages in the non-union sector is what we observe in the 1980s. It is striking that this is the decade in which the union threat fell fastest relative to other unionisation probabilities.[38]

Figure 4: Components of Decomposition

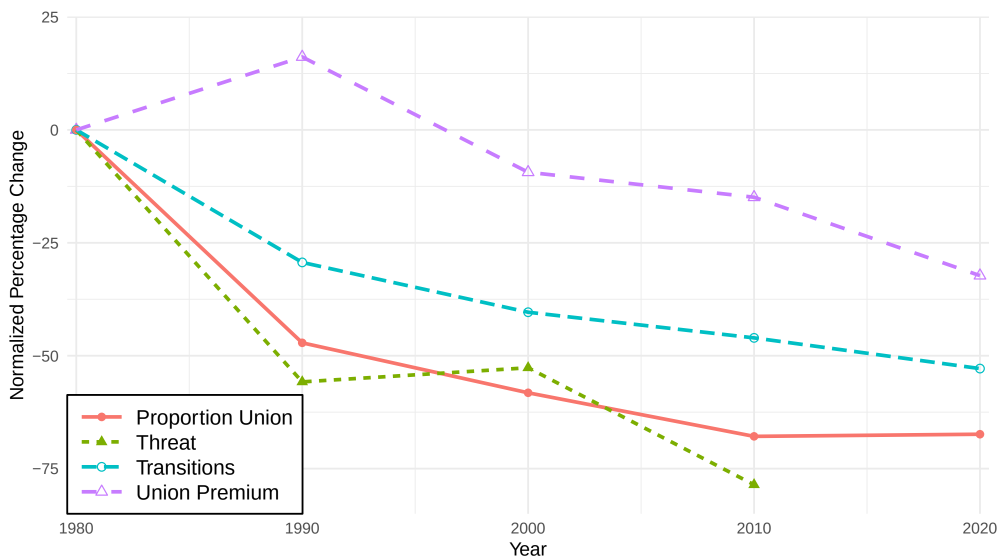

Notes: Data from the CPS and NLRB are shown as percentage changes relative to 1980 levels. Proportion union, union premium, and transitions are constructed from CBS data, as detailed in Appendix C. The threat of union election is derived from NLRB data and also described in Appendix C.

6.2 Overall Decomposition

We present our decomposition of the overall trend in average city wages in Figure 5. The bottom line in Figure 5 is the actual trend in an average city's (residualized) mean wage. It depicts an overall real wage trend that is strongly decreasing between 1980 and 1990 – falling by 15.6% in that decade – followed by a see-saw pattern of mild increases in the 1990s and declines in the 2000s.[39]

To understand the components in our decomposition, we use (28) to write the change in the city-level mean wage between period t and 1980 as follows:

[38]In Appendix H, we report on a rough check on this argument in which we regress changes in the union wage premia in industry×city cells on changes in our union threat variable, P^{ne}_{ict}. As our theory predicts, the union threat effect is negative and statistically significant.

[39]We end our figure in 2010 because we only have data on one element of our decomposition – the part related to union elections – up to that year.

Figure 5: Average Wage Decomposition

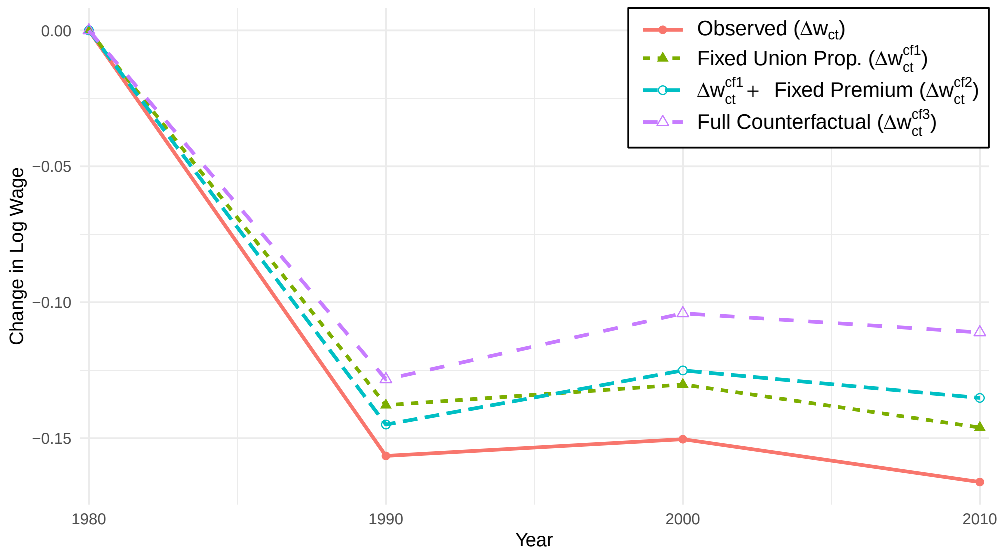

Notes: Data from the CPS and NLRB are presented as changes in log relative to 1980 levels. Wages for union and non-union workers are adjusted for worker characteristics and averaged across city-industries using fixed 1980 weights. See the main text for further details.

$$\Delta w_{ct} = \underbrace{\Delta P^u_{ct} \cdot (w^u_{ct} - w^n_{ct})}_{\text{Change in Union Proportion}} + \overbrace{\underbrace{P^u_{c80} \cdot \Delta(w^u_{ct} - w^n_{ct})}_{\text{Change in Union Wage Premium}} + \underbrace{\underbrace{\Delta w^n_{ct}}_{\text{Change in Non-Union Wages}}}_{\Delta w^{\text{cf2}}_{ct}}}^{\Delta w^{\text{cf1}}_{ct}}. \tag{29}$$

The first component of our decomposition is formed by setting $\Delta P^u_{ct} = 0$ (i.e., holding the union proportion at its 1980 value while allowing other factors that determine wage changes to vary). We denote this counterfactual wage series as $\Delta w^{\text{cf1}}_{ct}$ in Figure 5. This line shows that the decline in unionisation contributed to a 0.019 log-point drop in the mean wage in the 1980s, accounting for about 12% of the overall drop in the mean wage during that decade, with a similar effect on the drop from 1980 to 2010.

In examinations of the impact of unions on mean wage movements, authors often combine this first 'shifting weights' component with changes in the union wage premium. Thus, we form a second component by additionally setting $\Delta(w^u_{ct} - w^n_{ct}) = 0$, i.e., holding the union wage premium at its 1980 level. We refer to this counterfactual as $\Delta w^{\text{cf2}}_{ct}$ in the figure. As highlighted in Figure 4, the wage premium increased in the 1980s but declined thereafter. As a result, the impact of the union premium offsets the effect of de-unionisation in the 1980s but reinforces it in later decades. These two forces together account for a 3.1% drop in mean wages between 1980 and 2010. In a similar vein, Card et al. [2004] calculate that a standard shift-share analysis incorporating both declines in the unionisation rate and the union wage premium implies a drop in the mean US wage by 2.6% between 1984 and 2001.

A standard decomposition stops at this point. However, our estimates imply that de-unionisation affected the remaining component (the change in mean non-union wages) through both the bargaining and threat channels. To account for these effects, we return to our non-union wage specification, equation (19), which indicates that changes in the mean non-union wage are driven by changes in (1) industry wage premia in the non-union sector (γ_{0it}), (2) changes in outside option values (E_{uict} and E_{nict}), and (3) changes in the threat probability (P^{ne}_{ict}) (as well as ΔER_{ct} and other factors captured in the error term). De-unionisation affects the mean non-union wage through changes in (2) (due to underlying changes in the probability of finding a union job (transition rates, $T_{kct|j}$), and the relative value of union work (wage premium, $\nu_{uit} - \nu_{nit}$)) and (3).[40] We denote a counterfactual non-union wage as if changes in these factors did not occur as $w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}}$. Thus, non-union wage trends can be decomposed as:

$$\Delta w^{\mathrm{cf2}}_{ct} \equiv \Delta w^n_{ct} = \underbrace{\left[w^n_{ct} - w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}}\right]}_{\text{Non-Union Spillover Effect}} + \underbrace{\left[w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}} - w^n_{c80}\right]}_{\Delta w^{\mathrm{cf3}}_{ct}}. \quad (30)$$

To estimate $w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}}$, we use our estimated non-union wage equation and plug in 1980 values for the indicated components. However, these initially estimated wages are only first-round effects of de-unionisation. If the counterfactual wages in a particular *ic* cell are higher than observed, outside options for other workers would also be higher. Thus, we create a second round of counterfactual outside option values using the first round of counterfactual wages and then form a second round of counterfactual wages using updated outside options. We iterate this process until the predicted wages change by less than 0.1 percent. At the second and subsequent rounds, we update both the non-union and union wages using our estimated equations for each. This estimates the complete feedback loop inherent in bargaining schemes, ensuring that the union premia used in the outside option terms are consistent with the premia calculated from the set of counterfactual wages.

We refer to w^{cf3}_{ct} as the 'full counterfactual' in Figure 5 since it incorporates all the paths through which de-unionisation could affect the non-union mean wage. The last spillover component adds a further 2.4% over the full period, approximately doubling the estimated effect from the standard decomposition alone. Previous estimates of spillover effects of de-unionisation on non-union wages based on a regression of mean non-union wages on the union proportion range from large (Holzer [1982] and Denice and Rosenfeld [2018] – with estimated effects that would over-explain the decline in real non-union wages between 1980 and 2010 – to near-zero effects (Farber [2005]) to negative effects (Neumark and Wachter [1995]). Our estimates are closer to (though somewhat larger than) those in Farber [2005]. A related literature focuses on the impacts of inequality rather than wage levels. Fortin et al. (2021) find that taking account of spillovers roughly doubles the estimated 'shift-share' impact of de-unionisation on wage inequality over the 1979-2017 period.

Over the full 1980-2010 period, the three components together imply that de-unionisation can account for 33.2% of the total decline in the mean wage.[41] To provide further context

[40] We use national-level industrial premia differences ($\nu_{uict} - \nu_{nict}$) as drivers of outside option changes due to de-unionisation, not local wage premia ($w_{uict} - w_{nict}$). The former corresponds to ($\tilde{\gamma}_{0it} - \gamma_{0it}$) in our wage specifications and are treated as exogenous factors, while the latter are determined endogenously through spillovers within our model.

[41] In Appendix J, we present an alternate version of the decomposition based on separate estimates of the

for the size of our counterfactual effects, Autor et al. [2013]'s estimates of the impact of the China trade shock on the wages of non-manufacturing workers (their estimated effect on manufacturing wages is zero) amounts to a 0.009 decline between 1990 and 2000, and a 0.014 decline between 2000 and 2007. Together, these are approximately the same size as our estimated effect of de-unionisation on non-union wages alone from 1980 to 2010 and about 40% of our estimate of the total de-unionisation effect. Over this same period, the US federal minimum wage fell by 13% in real terms. If we take a relatively extreme estimate of minimum wage spillovers and assume that the wages of workers up to 1.2 times the minimum wage shift when the minimum wage shifts then using the fact that 18% of workers earned under 1.2 times the minimum wage in 1980 (Hardy et al. [2023]), the decline in the minimum wage would account for a 2.3% decline in the mean wage – again, about 40% of our total de-unionisation effect.

It is worth noting that, in our model, changes in the union wage premium (the second decomposition component) are driven by three factors. The first is changes in the difference in industrial wage premia between the union and non-union sectors ($\tilde{\gamma}_{0it}$ - γ_{0it}). This arises in our model because unions capture a different share of industry price movements over time, and that share may change as unions become weaker. We view that mechanism as an exogenous force which our model does not explain. On the other hand, the second and third factors (relative changes in outside option values between the sectors and the union threat probability) are forces within our model. In that sense, we can estimate how much of the second decomposition component is determined by the threat and bargaining effects from our model by forming counterfactual versions of both the non-Union and union wages formed using the 1980 values of the threat and transition variables combined with the estimated coefficients from our non-union and union wage regressions. The result implies that the latter forces account for 25% of the second (wage premium) component. Thus, changes in the threat and bargaining stemming from de-unionisation drive about half of the total impact of de-unionisation on the decline in the overall mean wage from 1980 to 2010.

In the top panel of Table 3, we present our counterfactual analysis for various sub-groups. The first row shows the total mean wage decline between 1980 and 2010 for each subgroup. The second row displays the standard shift-share effect, the third shows changes in union wage premia, and the fourth shows non-union spillover effects. The second panel further decomposes the spillover component, which we discuss in Section 6.4. Row (5) sums the spillover, wage premium, and shift-share effects, while the last row shows this total as a proportion of the total wage decline, indicating how much the decline would have been reduced if union-related factors had remained at 1980 levels.

Columns (2) and (3) show that men experienced a decline in the mean real wage between 1980 and 2010 which was over double that experienced by women, and also had a much larger loss in unionisation. As a result, spillover effects are larger for men. In the end, though, the proportion of the overall wage decline explained by de-unionisation is quite similar for men and women. Low-educated male workers in both our young and old age groups experienced similar impacts from de-unionisation, and for both, the effects are sizeable: 28.5% for the

full specification for each decade in order to check on whether violations of the local linearization in our specification are problematic. When we chain the estimated changes together, the three components account for 40% of the total decline.

former group and 29.7% for the latter. For more educated, young workers, de-unionisation had essentially no impact on their mean wage movements. For the highly educated, older workers, the spillover effects actually imply increases in mean wages. This arises because union jobs for this education group became more concentrated in higher-paying (public sector) jobs, implying increased average union wages that more than offset declines in the probability of getting a union job in their outside option term.

Table 3: Outside Options Contribution to Changing Wages - Subsample Analysis

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	All	Men	Women	Men			
				Young Low Skill	Young High Skill	Old Low Skill	Old High Skill
	1980-2010						
(1) Observed	–0.166	–0.223	–0.093	–0.320	–0.112	–0.288	–0.087
(2) Union Prop.	–0.020	–0.032	–0.007	–0.074	–0.003	–0.061	0.002
(3) Union Premium	–0.011	–0.002	–0.012	0.030	–0.004	0.035	–0.000
(4) Non-Union Spillovers	–0.024	–0.038	–0.015	–0.046	–0.005	–0.059	–0.001
(4a) Fixed Threat	–0.005	–0.004	–0.006	–0.003	–0.004	–0.004	–0.006
(4b) Fixed Transitions	–0.018	–0.032	–0.008	–0.041	–0.003	–0.061	0.000
(4c) Fixed Union Prem.	–0.001	–0.001	–0.001	–0.003	0.002	0.006	0.005
(5) Total	–0.055	–0.071	–0.034	–0.091	–0.011	–0.086	0.001
(6) Total/Observed	0.332	0.320	0.365	0.285	0.103	0.297	–0.012

Notes: This table displays results from the decomposition for union and nonunion workers from 1980-2010. Each column contains the decomposition results for a different subsample. All figures are log changes from 1980 levels. Details described in main text.

6.3 Implications for Wage Inequality

Given the heterogeneous exposure to de-unionisation across groups, the decline in union power has implications for inequality. Figure 6 presents decomposition results for four log wage differentials that are key to overall inequality movements: the gender gap, the post-secondary premium, and the post-secondary premium separately for men and women. Each series is plotted for "Observed" trends (capturing the actual national average of city-industry wages) and "Counterfactual" trends (representing a scenario in which all union-related factors are held constant at their 1980 levels).

The gender gap narrowed substantially over this period, declining by approximately 18 log points, with most of the improvement occurring in the 1980s. As Bidner and Sand [2024] note, the literature concerned with gender wage differentials has overwhelmingly focused on explanations for this decline that emphasize women's gains. However, our counterfactual analysis reveals that this narrowing would have been less pronounced if union-related factors had remained at their 1980 levels. De-unionisation disproportionately reduced male wages, indirectly reducing the gender gap by just under 4 percent.[42]

The post-secondary premium, which measures wage gaps between those with more than a high school education and those with a high school education or less, was also influenced by the decline in union power. Our estimates suggest that de-unionisation amplified inequality between education groups, with the post-secondary premium being about 6.3% lower in the counterfactual scenario where union power remained at 1980 levels. Gender-specific trends highlight differences in how de-unionisation affected this premium: for women, the effect of union decline on the education premium was minimal, while for men, the post-secondary premium would have been about 11% lower in the counterfactual scenario. This reflects the disproportionate impact of declining union power on lower-skilled men, exacerbating educational wage disparities.

Figure 6: Implications for Wage Inequality

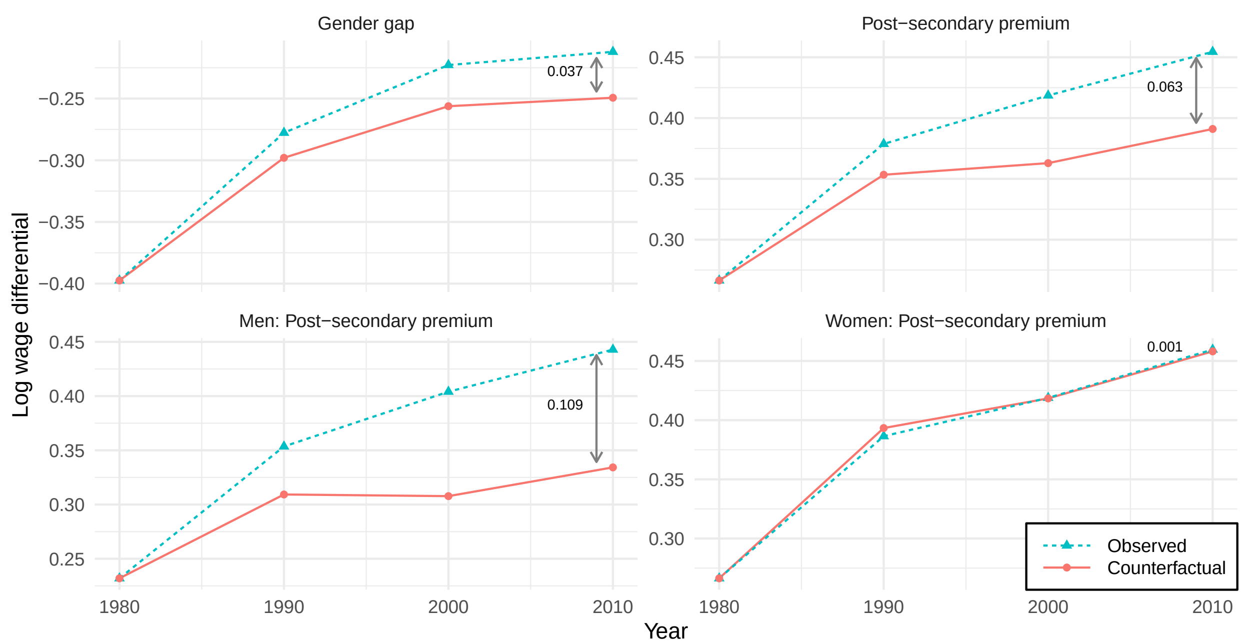

Notes: Data from the CPS and NLRB are presented as log changes relative to 1980 levels. Wages for union and non-union workers are adjusted for worker characteristics and averaged across city-industries using fixed 1980 weights. See the main text for further details.

6.4 Decomposing Non-union Wages

We next turn to decomposing the effect of de-unionisation on non-union wages, which, of course, is the focus of our estimation. To do so, we start with (30) and further decompose

[42]Bidner and Sand [2024] discuss how declines in men's employment opportunities have influenced the gender wage gap more broadly and highlight how general equilibrium forces can make standard decompositions misleading.

the Non-union Spillover Effect into its sub-components.

$$\text{Non-Union Spillover Effect} = \underbrace{\left[w^n_{ct} - w^n_{ct|P^{ne}_{ic80},T_{kct|j},\nu_{uit}-\nu_{nit}}\right]}_{\text{Union Threat (4a)}} + \underbrace{\left[w^n_{ct|P^{ne}_{ic80},T_{kct|j},\nu_{uit}-\nu_{nit}} - w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{uit}-\nu_{nit}}\right]}_{\text{Transitions (4b)}} + \underbrace{\left[w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{uit}-\nu_{nit}} - w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}}\right]}_{\text{Union Wage Premia (4c)}} \tag{31}$$

In Figure 7, we present each of the elements of our counterfactual non-union spillover effect. The line with diamonds corresponds to the total counterfactual effect of holding all the de-unionisation components constant at 1980 levels on the non-union wage. It says that all the factors combined resulted in a decline in the non-union wage of about 1.7 percentage point in the 1980s, rising to 2.4 percentage points by 2010. The remaining lines on the figure show the contribution to the full counterfactual of its constituent parts, and the sum of the points on those lines in a given year equals the total counterfactual effect. Holding the union wage premia to their 1980 values (shown in the line labelled 'Union Wage Premia Effect') would have resulted in a decrease in the non-union mean wage in the 1980s because increased premia in that decade increased the value of outside options. As we described earlier, our model provides an explanation for this seemingly odd trend in the union premium. In contrast, the large decline in the probability a non-union worker could find a union job in the 1980s (shown in the 'Transitions Effect' line) implied a substantial decline in the non-union wage. In fact, because the threat and wage premium effects happen to offset each other in that decade, the reduction in transition rates almost equals the size of the total spillover effect. In subsequent decades, the union wage premia decline and the transition effect stabilizes somewhat so that over the 1980 to 2010 period, the decline in the transition rates accounts for about 75% of the total spillover effect.

The last de-unionisation factor is the threat probability, which is captured by the line labelled 'Union Threat Effect'. What is most noteworthy about this effect is its size. While our estimates show clear evidence of the standard threat effect, its actual impact on non-union wage movements was small. This occurs mainly because the threat probabilities themselves are small, even in 1980. A small threat effect means that the sizeable spillover effect that emerges by 2010 in Figure 5 is almost completely accounted for by the bargaining channel. This has potentially important implications for policymaking aimed at raising wages since the threat effect can only be harnessed by increasing unionisation. The bargaining channel, in contrast, is not unique to unions – any policy that pushes up the outside option value for workers (such as eliminating non-compete clauses (Johnson et al. [2020] or expanding commuting options Hafner [2022]) can have this effect, and our results imply that this channel can be powerful. This is reminiscent of the results in Caldwell and Danieli [2021], who show that wages are increasing in their index of the value of outside options. Their index increases when workers have greater probabilities of transferring to other occupations and job opportunities. Our result is driven by decreases in the probability a worker can transfer to a union job.

Figure 7: Decomposition components: Non-union workers

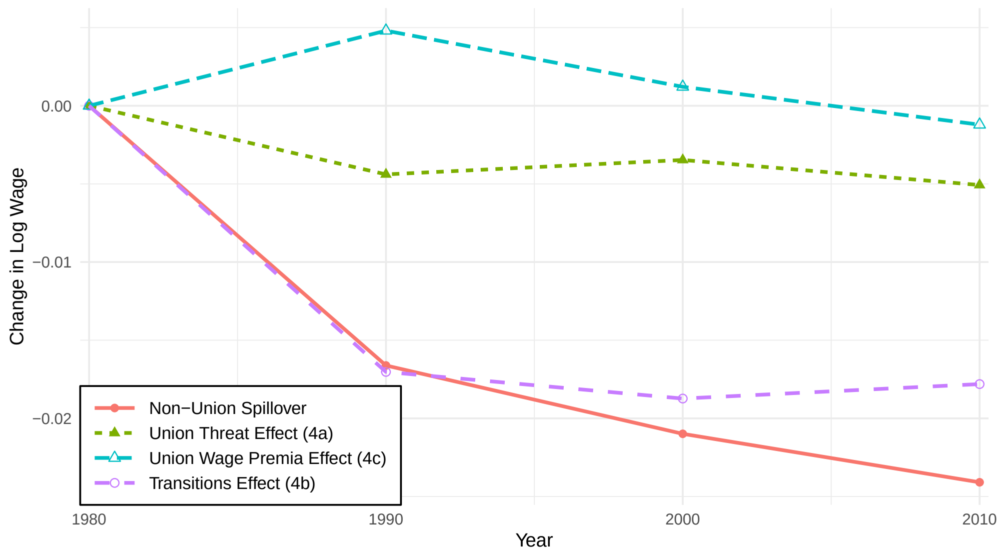

Notes: Data from the CPS and NLRB are shown as log changes relative to 1980 levels. Wages for non-union workers are adjusted for worker characteristics. Each series represents a decomposition component detailed in the main text.

The second panel in Table 3 shows the components of the non-union spillover effect for the full 1980-2010 sample period for different sub-groups. From this, we can see that the spillover effect for men is over double that for women, with most of that accounted for by differences in the transition effects. Similarly, the main reason that low-educated non-union workers were more affected by de-unionisation was because of a reduced chance of individual workers finding a union job rather than because of the reduced probability that their firm would be unionised.

7 Conclusion

In this paper, we provide new estimates of the impact of unions on non-union wage setting. We allow the presence of unions to affect non-union wages both through the typically discussed channel of non-union firms emulating union wages in order to fend off the threat of unionisation and through a bargaining channel in which non-union workers use the presence of union jobs as part of their outside option. We specify these channels in a search and bargaining model that includes union formation and the possibility of non-union firms responding to the threat of unionisation. By formalising wage setting and union formation, we derive a specification grounded in theory that provides guidance on what to control for, how to interpret our coefficients, and what is in the error term. Based on that, we derive a set of instruments and a model-based over-identification test, the values for which imply that our identification strategy is appropriate for this data.

Our estimates indicate that de-unionisation in the US after 1980 substantially affected non-union wages, particularly, and the wage structure in general. In a decomposition exercise, holding the probability a worker can find a union job, the probability a firm faces a unionisation drive, and union wage premia constant at their 1980 levels would have undone 33% of the 16% decline in the mean (composition constant) real wage in a typical city in the US between 1980 and 2010. While we find evidence for the spillover effects of unions on non-union wage setting through the traditional threat and bargaining channels, the latter dominates. That is important for policymakers looking for tools to help in raising wages. The union threat channel can only be implemented by increasing union power. However, the bargaining channel is not specific to unions. Any policy that raises worker outside option values will raise wages for a wide set of workers (Beaudry et al. [2012], Caldwell and Danieli [2021]). Unions are just one mechanism for doing that – though our estimates indicate it is a powerful and direct one. Finally, it is worth noting that what we have examined in this paper is only one path through which unions can affect labour market outcomes. When unions are stronger, there is also the possibility that they can affect investment (Alder et al. [2023]) and impact elections and policy-making, shifting policy on labour market regulation and minimum wages that would have their own effects on the wage structure (Feigenbaum et al. [2018]).

References

J. M. Abowd and H. S. Farber. Job queues and the union status of workers. *ILR Review*, 35(3):354–367, 1982.

K. G. Abraham and H. S. Farber. Returns to seniority in union and nonunion jobs: a new look at the evidence. *ILR Review*, 42(1):3–19, 1988.

K. G. Abraham and J. L. Medoff. Length of service and layoffs in union and nonunion work groups. *ILR Review*, 38(1):87–97, 1984.

K. G. Abraham and J. L. Medoff. Length of service and promotions in union and nonunion work groups. *ILR Review*, 38(3):408–420, 1985.

Ö. T. Açıkgöz and B. Kaymak. The rising skill premium and deunionization. *Journal of Monetary Economics*, 63:37–50, 2014.

R. Adao, M. Kolesar, and E. Morales. Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, 134(4):1949–2010, 2020.

J. T. Addison and B. T. Hirsch. Union effects on productivity, profits, and growth: Has the long run arrived? *Journal of Labor Economics*, 7(1):72–105, 1989.

S. Alder, D. Lagakos, and L. Ohanian. Labor market conflict and the decline of the rust belt. *Journal of Political Economy*, 131(10):296–315, 2023.

D. H. Autor, D. Dorn, and G. H. Hanson. The china syndrome: Local labor market effects of import competition in the united states. *The American Economic Review*, 103(6): 2121–2168, 2013.

I. Bassier. Collective bargaining and spillovers in local labour markets. Discussion Paper 1895, Centre for Economic Performance, 2022.

P. Beaudry, D. A. Green, and B. Sand. Does industrial composition matter for wages? a test of search and bargaining theory. *Econometrica*, 80(3):1063–1104, 2012.

C. Bidner and B. Sand. Pay gaps and outside wages: The u.s. gender wage gap 1980-2010. Technical report, Canadian Labour Economics Forum, 2024.

K. Borusyak, P. Hull, and X. Jaravel. Quasi-experimental, shift-share research designs. *The Review of Economic Studies*, 89:181–213, 2022.

K. Bronfenbrenner. The role of union strategies in nlrb certification elections. *Industrial and Labor Relations Review*, 50(2):195–212, 1997.

K. Bronfenbrenner. No holds barred: The intensification of employer opposition to organizing. Briefing Paper 235, Economic Policy Briefing Paper, 2009.

S. Caldwell and O. Danieli. Outside options in the labor market. Technical report, Tel Aviv University, 2021.

S. Caldwell and N. Harmon. Outside options, bargaining, and wages: Evidence from coworker networks. Technical report, University of California at Berkeley, 2019.

D. Card. The effect of unions on wage inequality in the us labor market. *ILR Review*, 54 (2):296–315, 2001.

D. Card. Immigration and inequality. *American Economic Review*, 99(2):1–21, 2009.

D. Card, T. Lemieux, and W. C. Riddell. Unions and wage inequality. *Journal of Labor Research*, 25(4):519–559, 2004.

D. Card, T. Lemieux, and W. C. Riddell. Unions and wage inequality: The roles of gender, skill and public sector employment, 2018. NBER working paper w25313.

P. Denice and J. Rosenfeld. Unions and nonunion pay in the united states, 1977-2015. *Sociological Science*, 5:541–561, 2018.

W. Dickens and L. F. Katz. Interindustry wage differences and industry characteristics. Technical report, National Bureau of Economic Research Cambridge, Mass., USA, 1986.

J. DiNardo and T. Lemieux. Diverging male wage inequality in the united states and canada, 1981–1988: Do institutions explain the difference? *ILR Review*, 50(4):629–651, 1997.

J. DiNardo, N. M. Fortin, and T. Lemieux. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5):1001–1044, 1996.

M.-P. Donsimoni. Union power and the american labour movement. *Applied Economics*, 13 (4):449–464, 1981.

H. S. Farber. Nonunion wage rates and the threat of unionization. *ILR Review*, 58(3): 335–352, 2005.

H. S. Farber, D. Herbst, I. Kuziemko, and S. Naidu. Unions and inequality over the twentieth century: New evidence from survey data. *The Quarterly Journal of Economics*, 136(3): 1325–1385, 2021.

J. Feigenbaum, A. Hertel-Fernandez, and V. Williamson. From the bargaining table to the ballot box: Political effects of right to work laws, 2018. NBER working paper w24259.

S. Firpo, N. Fortin, and T. Lemieux. Decomposing wage distributions using recentered influence function regressions. *Econometrics*, 6(2):28, 2018.

N. M. Fortin, T. Lemieux, and N. Lloyd. Labor market institutions and the distribution of wages: The role of spillover effects. Technical report, University of British Columbia, 2019. *Working Paper*.

N. M. Fortin, T. Lemieux, and N. Lloyd. Labor market institutions and the distribution of wages: The role of spillover effects. *Journal of Labor Economics*, 39(S2):S369–S412, 2021. doi: 10.1086/712923. URL https://doi.org/10.1086/712923.

R. B. Freeman. Unionism and the dispersion of wages. *ILR Review*, 34(1):3–23, 1980.

R. B. Freeman and J. L. Medoff. The impact of the percentage organized on union and nonunion wages. *The Review of Economics and Statistics*, pages 561–572, 1981.

A. Gosling and T. Lemieux. Labour market reforms and changes in wage inequality in the united kingdom and the united states, 2001. NBER working paper w8413.

F. Hafner. The equilibrium effects of workers' outside employment options: Evidence from a labor market integration. Technical report, Alto University, 2022.

B. Hardy, S. Holla, E. Krause, and J. Ziliak. Inequality in the united states: 1975-2022. Technical report, Institute for Fiscal Studies, 2023.

J. J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1):153–161, 1979.

B. T. Hirsch and A. N. Link. Unions, productivity, and productivity growth. *Journal of Labor Research*, 5(1):29–37, 1984.

B. T. Hirsch and J. L. Neufeld. Nominal and real union wage differentials and the effects of industry and smsa density: 1973-83. *The Journal of Human Resources*, 22(1):138–148, 1987.

T. J. Holmes. Geographic spillover of unionism. *NBER Working Paper Series*, 2006. URL https://api.semanticscholar.org/CorpusID:3105553.

H. J. Holzer. Unions and the labor market status of white and minority youth. *ILR Review*, 35(3):392–405, 1982.

G. Jarosch, J. S. Nimczik, and I. Sorkin. Granular search, market structure, and wages, 2019. NBER working paper w26239.

G. Jarosch, J. S. Nimczik, and I. Sorkin. Granular search, market structure, and wages. *Review of Economic Studies*, 2024.

M. Johnson, K. Lavetti, and M. Lipsitz. The labor market effects of legal restrictions on worker mobility, 2020. Available at SSRN: https://ssrn.com/abstract=3455381 or http://dx.doi.org/10.2139/ssrn.3455381.

L. M. Kahn. Union spillover effects on organized labor markets. *The Journal of Human Resources*, 15(1):87–98, 1980.

P. Krusell and L. Rudanko. Unions in a frictional labor market. *Journal of Monetary Economics*, 80:35–50, 2016.

T. Lemieux. Increasing residual wage inequality: Composition effects, noisy data, or rising demand for skill? *American Economic Review*, 96(3):461–498, 2006.

H. G. Lewis. *Unionism and relative wages in the United States: an empirical inquiry.* University of Chicago press, 1963.

D. Neumark and M. L. Wachter. Union effects on nonunion wages: Evidence from panel data on industries and cities. *ILR Review*, 49(1):20–38, 1995.

C. A. Pissarides. Trade unions and the efficiency of the natural rate of unemployment. *Journal of Labor Economics*, 4(4):582–595, 1986.

M. Podgursky. Unions, establishment size, and intra-industry threat effects. *ILR Review*, 39 (2):277–284, 1986.

S. Rosen. Trade union power, threat effects and the extent of organization. *The Review of Economic Studies*, 36(2):185–196, 1969.

E. Sanderson and F. Windmeijer. A weak instrument f-test in linear iv models with multiple endogenous variables. *Journal of Econometrics*, 190(2):212–221, 2016.

I. G. Snoddy. Learning about selection: An improved correction procedure. Technical report, University of British Columbia, 2019. *Working Paper*.

J. H. Stock and M. M. Watson. Introduction to econometrics, 3rd international edition, 2011.

M. Taschereau-Dumouchel. The Union Threat. *The Review of Economic Studies*, 87(6): 2859–2892, 05 2020. ISSN 0034-6527.

J. Tschopp. Wage formation: Towards isolating search and bargaining effects from the marginal product. *The Economic Journal*, 127(603):1693–1729, 2017.