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

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

July 3, 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 the impacts of unions on overall wage levels. 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 reassess 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 described 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, which would presumably impose 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 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 percentage 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 (e.g., Freeman and Medoff [1981], Donsimoni [1981], Holzer [1982], Kahn [1980], Dickens and Katz [1986], Hirsch and Neufeld [1987], Podgursky [1986], Neumark and Wachter [1995], Fortin et al. [2021]). In an important analysis, Farber [2005] shows that spillover estimates are highly sensitive to the source of variation and included controls, helping explain the disparity in earlier results. When controlling for a wide range of potential omitted variables, he finds at most a small positive effect of union power on non-union wages. Most earlier studies also lack a clear identification strategy to address the endogeneity of unionisation rates.[2] Finally, none of the papers in the literature even mention the twin problem of potential selectivity bias: as unionisation declines, the composition of non-union workers and firms is likely to 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 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

[2] An exception is Farber [2005], who uses an event study of Right to Work (RTW) laws in Idaho and Oklahoma. An earlier version of Fortin et al. [2019] extends this approach, but the identifying variation is limited because few states changed RTW status during their sample period.

the probability that the worker can transition 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 national-level changes in industry growth, industry wage premia, and transition probabilities across job types defined by industry and union status.

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 impact 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$\times$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 from 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-unionisation from 1980 to 2010, even though this decade saw the largest decline in unionisation. 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.

Our work also relates to a large literature on declining unionisation, including studies of the union wage premium and the role of de-unionisation in rising wage inequality. Card et al. [2004] and Card et al. [2018] provide comprehensive overviews of this research, following the early contribution of Freeman [1980]. Farber et al. [2021] provides a comprehensive account of the relationship between union density and inequality in the U.S., using new survey data to extend their analysis back to the 1930s. They find that rising unionisation significantly reduced inequality after WWII, while the subsequent decline in unionisation had a more modest effect on rising inequality over the past 50 years. Their framework allows for spillover effects onto non-union wages, but these are not examined directly. Our results suggest that spillovers may have contributed significantly to the inequality-reducing effects of unions in their analysis and help explain why those effects appeared weaker during the major union decline of the 1980s.[3]

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 the impact of spillovers on wage structure movements and the role played by our two channels. Section 7 contains conclusions.

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. Throughout the paper, we refer to non-union firms' wage responses to resist unionisation as standard threat effects, reflecting their treatment in the existing 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

[3] See, e.g., DiNardo et al. [1996], who attribute 14% of the rise in male wage inequality (1979–1988) to declining unionisation, and extensions in DiNardo and Lemieux [1997], Card [2001], Gosling and Lemieux [2001], Fortin et al. [2021], and Firpo et al. [2018].

all workers—even workers in firms not directly threatened with unionisation and 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.

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.[4] 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. Unionisation is costly and involves a one-time fixed cost, λ_c^*, which varies across cities. 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.

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

[4] 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.

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 choose whether or not to respond. When the cost of unionisation exceeds the value of the amenity, unionisation reduces the total surplus generated by the match between the firm and its workforce. However, for some combinations of amenity, productivity, and unionisation cost, the enhanced bargaining power from unionisation allows workers to capture more surplus than under the Simple Non-union arrangement—even though the total surplus is smaller. In such cases, both parties would be better off avoiding unionisation, provided the workers receive a higher wage that leaves them at least as well off as under unionisation.[5] Recognizing this, the firm will offer the workers a wage that is just slightly above the wage that will make them indifferent about unionising and the firm will remain non-union.[6] That wage will be higher than the Simple Non-union wage, and will be increasing in the benefits of unionising (the union wage and amenities) and decreasing in the cost of unionising. 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 workers' threat to unionise. The workers will then proceed to unionise (bearing the one-time fixed cost of doing so) and wages will be set in bargaining between the firm and the union. 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.

In this model, an increase in the cost of unionisation reduces the union threat, lowering both the number of Emulating Non-union firms and the wages they pay. It also reduces workers' outside options in both Simple and Emulating Non-union firms, since there are fewer high-wage Union and Emulating Non-union firms to move to. This decline in outside options affects bargained wages across sectors.

[5] 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.

[6] We can think of this as dividing up the surplus gained by remaining non-union relative to the alternative of unionising. The firm captures the whole surplus because the workers are unwilling to bear the cost of forming a union and bargaining as a group.

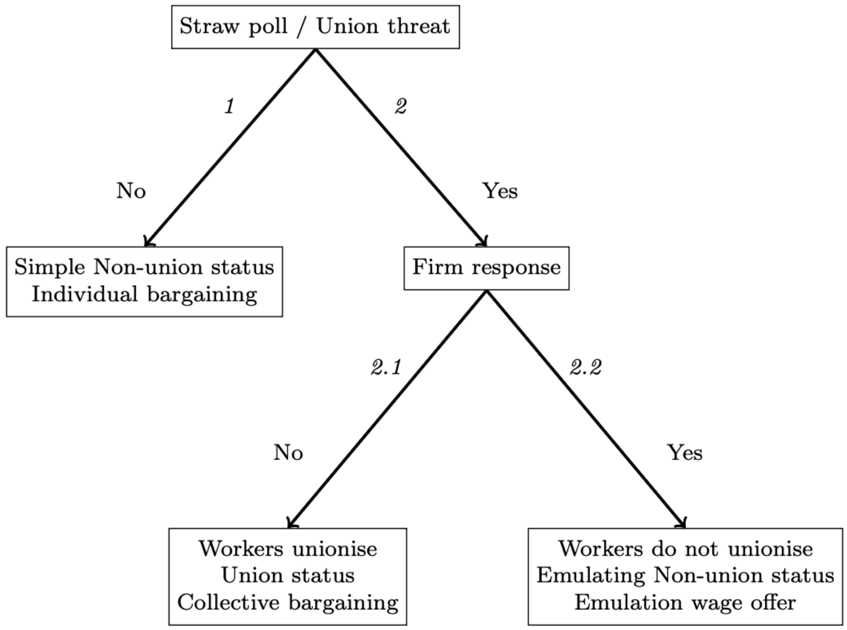

Figure 1: Sequence of events after hiring.

These mechanisms give rise to two distinct spillover effects of union shifts on non-union wages that we aim to estimate: a bargaining effect, which operates through non-union workers' outside options in Simple Non-union firms, and a threat effect, which operates through the incentives for non-union firms to emulate union wages. To connect the model to the data, note that we only observe whether a worker is employed at a union or non-union firm—not the type of non-union firm. As a result, as discussed in detail in Section 3, the observed non-union wage equation will be a weighted average of the Simple Non-union wage and the Emulation wage, with the latter being a function of Union wages.

With this structure in mind, we next fill in the details needed to derive our estimating non-union wage equation. 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.[7] 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. 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 the mobility friction, $\varphi_{k|j}$. 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^u_c or η_{kc} is low).

2.3 Firms

We assume the number of firms operating with type j jobs (defined as a combination of union status τ and industry i) in city c is fixed, leaving endogenizing firm formation for future work. For notational simplicity, we drop the (firm-job-city-specific) subscript on firm employment

[7] 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.

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. It will prove useful to write $\epsilon_{fic} = \epsilon_{ic} + u_{fic}$, where ϵ_{ic} is a city-level, sector-wide productivity shock, and u_{fic} is a mean-zero, firm-specific shock. This specification implies that technology is common across cities within an industry, but comparative advantage in producing each good varies across cities, as captured by the ϵ_{ic} terms. 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.[8] We assume that σ_i is sufficiently below 1 so that, combined with the assumption of a fixed number of firms in each ic cell, 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 they know will be bargained with their workers or set 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 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.

[8] 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.

Hence, in the bargaining game, both parties treat the future wage as predetermined and independent of today's bargaining outcome. V_{jc}^U is the value of unemployment for a worker formerly employed in a job of type j and is given by:

$$V_{jc}^U = b + \rho \left[q_c^u \sum_k T_{kc|j} V_{kc}^E(w'_{kc}) + (1 - q_c^u) V_{jc}^U \right], \tag{5}$$

where, b is the flow value of being unemployed, $V_{kc}^E(w'_{kc})$ 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 $T_{kc|j}$.[9]

Equation (5) says that the value of unemployment is higher when b is higher, when the probability of making a match (q_c^u) is higher, and when the expected value of the match, $\sum_k T_{kc|j} V_{kc}^E$, is higher. In Appendix A.4, we show that in steady state, $\sum_k T_{kc|j} V_{kc}^E$ 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).

2.5 Wage Setting

Before turning to wage setting by union status, it is useful to recall that all three wage equations—Union, Simple Non-union, and Emulation—are relevant for estimating the observed non-union wage equation. The threat effect operates through the Emulation wage, which is itself a function of Union wages, while the bargaining effect operates through the Simple Non-union wage.

Also recall that the job subscript j combines union status $\tau = \{u, n, e\}$ and industry i. To make the notation more intuitive, 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.

[9] 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.

2.5.1 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 is:

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

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. In Appendix A.1, we 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 in 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:

$$w_{fuic} = \tilde{\gamma}_{0i} + \tilde{\gamma}_1 E_{uic} + \tilde{\gamma}_2 ER_c + \tilde{\gamma}_3 \epsilon_{ic} + \tilde{\gamma}_3 u_{fic} - \tilde{\gamma}_4 \psi_{fuic}, \tag{10}$$

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. E_{uic} is our shorthand notation for the outside option value for union workers, i.e., $E_{uic} = \sum_k T_{kc|ui} w_{kc}$.

Equation (10) states that union wages increase with productivity—captured by $\tilde{\gamma}_{0i}$, u_{fic}, and ϵ_{ic}—and with local labour market tightness, as reflected in ER_c, and decrease with union amenities (ψ_{fuic}). In addition, union wages rise with the workers' outside option value (E_{uic}). To understand how union wages depend on the determinants of this outside option term, it is useful to substitute in the expression for transition probabilities from equation (6). We obtain:

$$E_{uic} = \sum_k T_{kc|ui} w_{kc} = \sum_k \eta_{kc} \frac{\varphi_{k|ui}}{\sum_{k'} \eta_{k'c} \varphi_{k'|ui}} w_{kc}. \tag{11}$$

The value of the outside option varies across cities and depends on three key determinants. The first is η_{kc}, the share of city employment in each job type, defined by the combination of union status and industry. The second is the wage paid in that job type, w_{kc}. To this point, we have assumed that workers are homogeneous, so wage differences across job types reflect rents—that is, pay above what is required to attract the marginal worker. Those rents are maintained because of the frictions in the labour market. It is important that we consider rents since wage differences due to compensating differentials or skill requirements do not provide a basis for bargaining higher wages with a current employer. The third determinant is the ease with which a worker in job type ui can transition to job type k, captured by $\varphi_{k|ui}$. This includes transitions from union to non-union jobs. Thus, for workers in job type ui, the outside options—and hence their wage—is higher in cities with more jobs that pay high rents and are easily accessible to them.

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:

$$\beta S_{fnic}(n) = (1-\beta)\left[V^E_{fnic}(n) - V^U_{nic}\right]. \tag{12}$$

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.[10] 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:

$$S_{fnic}(n) \quad = \quad \frac{\partial \pi_{fnic}(n)}{\partial n} + \rho^e(1-\delta^m)\frac{\kappa}{q^v_c}, \tag{13}$$

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

[10] 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.

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:

$$w_{fnic} = \gamma_{0i} + \gamma_1 E_{nic} + \gamma_2 ER_c + \gamma_3 \epsilon_{ic} + \gamma_3 u_{fic}. \tag{14}$$

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. The term E_{nic} is shorthand for the outside option value for non-union workers, i.e.:

$$E_{nic} = \sum_k T_{kc|ni} w_{kc} = \sum_k \eta_{kc} \frac{\varphi_{k|ni}}{\sum_{k'} \eta_{k'c} \varphi_{k'|ni}} w_{kc}. \tag{15}$$

Its coefficient, γ_1, will capture the bargaining effect in the observed non-union wage equation, discussed below. When non-union firms are not under threat of unionisation, the observed non-union wage is given by equation (14), and unionisation affects it solely through its effect on outside options.

Importantly, $\tilde{\gamma}_{0i} > \gamma_{0i}$ and, so, 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 any formal vote. In particular, if the cost of unionising is greater than the value of the union amenity, unionising shrinks the size of the total pie. Workers are still interested in unionising because the slice of the smaller pie that they get through forming a union and bargaining together is bigger than the slice of the larger non-union pie they get when they bargain individually. In this circumstance, the firm recognizes that the lost surplus from moving to unionisation can be saved (and the firm can keep as much of it as possible for itself) if it offers the workers a wage just high enough to induce them to vote against a union but still lower than the union wage it would have to pay. It will make a unilateral offer and the workers will accept it without further bargaining, exactly because the workers vote against taking on the cost to form a bargaining unit.

In making its offer, the firm recognises that the workers will choose to unionise if the value of employment at the firm under unionisation, net of the unionisation cost λ_c^*, exceeds the value of employment when the firm is nonunion. Thus, the wage offer that would make them indifferent about forming a union (the emulation wage, w_{feic}) solves:

$$V^E_{feic}(w_{feic}) = V^E_{fuic}(w_{fuic}) - \lambda_c^*, \tag{16}$$

where $V^E_{feic}(w_{feic})$ and $V^E_{fuic}(w_{fuic})$ denote the value of employment for a worker when the firm is an Emulating Non-union firm and a Union firm, respectively. Plugging in the steady-state expressions for the value of employment, the emulation wage is given by:

$$w_{feic} = w_{fuic} + \psi_{fuic} - \lambda_c - \Delta_{eic,uic}, \tag{17}$$

where $\lambda_c = [1 - \rho(1 - \delta)] \lambda^*_c$, and $\Delta_{eic,uic}$ captures the difference in the value of unemployment between being a non-union and a union worker, and drops out in the linearization step (see Appendix A.5). Equation (17) implies that the Emulating Non-union wage equals the Union wage plus an adjustment that is equal to the difference between the amenity provided by the union and the cost of unionisation.

Substituting in our expression for the Union wage, w_{fuic}, and linearizing, we arrive at:

$$w_{feic} = \tilde{\gamma}_{0i} + \tilde{\gamma}_1 E_{uic} + \tilde{\gamma}_2 ER_c + \tilde{\gamma}_3 \epsilon_{ic} + \tilde{\gamma}_3 u_{fic} + (1 - \tilde{\gamma}_4)\psi_{fuic} - \lambda_c, \tag{18}$$

where $1 - \tilde{\gamma}_4 > 0$. Importantly, equation (18) depends on E_{uic}, the outside option value for union workers. When this value is high—for example, when many high-wage jobs are accessible to union workers—Emulating Non-union firms must offer higher wages to deter unionisation. For this reason, $\tilde{\gamma}_1$ will capture the threat effect in the estimated non-union wage equation.

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 gives rise to two amenity thresholds. 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 $\psi_{fuic} > \lambda_c$. For technical reasons, there is also a productivity threshold below which there are no Emulating Non-union firms. Overall, 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 surplus—reflecting a standard selection mechanism. A full derivation is provided in Appendix A.8, and Appendix Figure 2 illustrates the resulting union statuses in the amenity-productivity space.

3 Empirical Specification

We are now in a position to present our empirical specification, which is defined at the industry×city cell level. Our objective is to estimate both types of spillovers—threat and bargaining effects—arising from shifts in unionisation on non-union wages. Importantly, in our data, we observe only whether a worker is employed at a union or non-union firm, not the type of non-union firm. As a result, our dependent variable—the observed non-union wage—is a weighted average of wages in Simple Non-union (equation (14)) and Emulating Non-union firms (equation (18)):

$$w^n_{ict} = (1 - P^{ne}_{ict}) \cdot w_{nict} + P^{ne}_{ict} \cdot w_{eict}, \tag{19}$$

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. We refer to P^{ne}_{ict} as the threat probability, as it represents the share of non-union firms that are under direct threat of unionisation. Because the Emulating Non-union wage is a function of the Union wage, the Union wage equation (10) is also relevant for our specification.

We use equations (14) and (18), difference at the decadal level—thereby removing any industry-by-city time-invariant characteristics—and divide through by a base wage to express the equation in log terms. This yields our main estimating equation for non-union wages:[11]

$$\begin{aligned} \Delta \ln w^n_{ict} &= \underbrace{\gamma_1 \Delta[(1 - P^{ne}_{ict}) E_{nict}]}_{\text{Bargaining effect}} + \underbrace{\tilde{\gamma}_1 \Delta(P^{ne}_{ict} E_{uict})}_{\text{Threat effect}} \\ &+ \gamma_2 \Delta ER_{ct} + \Delta\gamma_{0it} + \Delta(\gamma^*_{0it} P^{ne}_{ict}) - \Delta(P^{ne}_{ict}\lambda_{ct}) + \tilde{u}_{ict}, \end{aligned} \tag{20}$$

where $\Delta x_t = x_t - x_{t-1}$ denotes a decadal difference, $\gamma^*_{0it} = (\tilde{\gamma}_{0it} - \gamma_{0it})$, and $\tilde{u}_{ict}$ is the error term, described in detail below. We view the different time periods (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.[12]

The first two terms in equation (20) capture the bargaining and threat effects. The bargaining effect arises because unionisation impacts non-union wages through non-union workers' outside options: when the threat probability, P^{ne}_{ict} is zero, observed non-union wages are equivalent to Simple Non-union wages and unionisation affects them only via its impact on those outside options. The threat effect reflects firms raising wages to deter unionisation. More wage deterrence is needed the greater is P^{ne}_{ict} and the higher are union wages. Union wages, in turn, are higher when union worker outside options, E_{uict}, are higher. In our data, union workers are much more likely than non-union workers to access union jobs. As a result, their outside option value is larger and moves differently from that of non-union workers. If the outside option value of union workers were a significant determinant of non-union wages, this would provide strong evidence that the threat of unionisation plays a role in non-union wage setting.

The non-union wage is also increasing in changes in the employment rate, ΔER_{ct}, because tighter labour markets makes it easier for workers to access alternative job opportunities. Increases in output prices also imply higher wages and are captured in a full set of industry×time

[11]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. As a result, all the coefficients are scaled by $w(x_0)$, though, for simplicity, we have not changed the coefficient labels.

[12]We 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}.

fixed effects ($\Delta\gamma_{0it}$). The inclusion of those effects implies that we identify the bargaining and threat effects by comparing wage changes within the same industry across cities with different movements in outside option values. 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.

In addition to differences in outside option values, union wages are higher than non-union wages because greater union bargaining power implies that they are able to capture a larger share of the p_i terms (i.e., $\gamma^*_{0it} > 0$). On the other hand, if the costs of unionisation are higher in a city then the emulating wage is lower. The first of these forces amplifies the effects of movements in P^{ne}_{ict}, while the second reduces it. In equation (20), these are represented by the $\Delta(\gamma^*_{0it}P^{ne}_{ict})$ and $\Delta(P^{ne}_{ict}\lambda_{ct})$ terms, respectively. We capture them in our empirical specification by including 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. These interactions are meant to proxy for heterogeneity in rent-sharing (across industries) and union costs (across cities) that influence how changes in the threat of unionisation affect wages. We use P^{ne}_{ict-1} rather than ΔP^{ne}_{ict} to avoid endogeneity issues.[13]

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. The error term in equation (20) is given by:

$$\tilde{u}_{ict} = \underbrace{\Delta\mu_{ict}}_{\text{Selection}} + \underbrace{\gamma_3\Delta\epsilon_{ict} + (\tilde{\gamma}_3 - \gamma_3)P^{ne}_{ict-1}\Delta\epsilon_{ict}}_{\text{Variation in sectoral productivity}} + \underbrace{(\tilde{\gamma}_3 - \gamma_3)\epsilon_{ict}\Delta P^{ne}_{ict}}_{\text{Rent-capture term}},$$

where μ_{ict} captures selection of firms into either type of Non-union status, the next two terms reflect components related to variation in sectoral productivity $\Delta\epsilon_{ict}$, and the last term reflects another reason why unions obtain higher wages—that they capture a larger share of industry-city specific rents ($(\tilde{\gamma}_3 - \gamma_3)\epsilon_{ict}$).[14] The structure of this error term introduces a range of identification challenges, which we discuss in the next section. A complete derivation of our empirical specification, equation (20), including the form of the selection term μ_{ict}, is provided in Appendix A.9.

3.1 Implementation and Identification Challenges

We turn, next, to describing the set of challenges with taking equation (20) to our data and how we solve them.

[13] 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.

[14] While we can address selectivity directly and, under our identifying assumptions, the productivity variation terms are independent of our instruments, the rent capture term in the error potentially introduces omitted variables bias. As a result, our estimated coefficients on the outside option terms reflect both outside option and rent capture effects. Since our goal is to estimate the total effect of de-unionisation on non-union wage changes, this is not necessarily a concern, though it does complicate our decomposition. In Appendix F, we show via simulation that the rent capture component is negligible for the primary coefficients of interest.

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.[15] 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$\times$city fixed effects under an assumption that a_l is orthogonal to job$\times$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 identification problems posed by the two interaction terms involving the outside options and the threat probability in equation (20), it is useful to write them out in full. Starting with the first term in the equation:

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

We address the associated identification challenges by constructing an instrument for this term, using exogenous variables corresponding to each of the endogenous components: E_{nict-1}, E_{nict}, and P_{ict}^{ne}. Note that under the identifying assumption that validates our proposed exogenous variables, the start-of-period threat probability P_{ict-1}^{ne} is likewise exogenous and not subject to endogeneity concerns. The resulting instrument is:

$$IV_{nict} = \left(1 - \hat{P}_{ict}^{ne}\right) \cdot \hat{E}_{nict} - \left(1 - P_{ict-1}^{ne}\right) \cdot \tilde{E}_{nict-1}, \tag{23}$$

where $\tilde{E}_{nict-1}$, $\hat{E}_{nict}$, and $\hat{P}_{ict}^{ne}$ are the exogenous counterparts to the endogenous components E_{nict-1}, E_{nict}, and P_{ict}^{ne}, respectively.

By the same logic, $\Delta(P_{ict}^{ne}E_{uict})$—the second term of equation (20)—faces the same identification challenges and is constructed analogously as $IV_{uict} = \hat{P}_{ict}^{ne} \cdot \hat{E}_{uict} - P_{ict-1}^{ne} \cdot \tilde{E}_{uict-1}$. Under our identifying assumption, these instruments allow for consistent estimation of the coefficients in our main specification. We now turn to the construction of each exogenous component that makes up equation (23).

[15] An alternative interpretation is that labour markets are segmented by observable skill and demographic characteristics, such as education and gender, and that our model applies to homogeneous workers within each segment. We also conduct our analysis separately by education group as a specification check below.

Constructing $\tilde{\mathbf{E}}_{\mathbf{nict-1}}$ The first issue is that E_{nict-1} creates a reflection problem in relation to the $\ln w^n_{ict-1}$ part of the differenced wage that is our dependent variable. This arises because the value of outside options depends on wages in other jobs within the same city—the w_{kct} terms in equation (15) (w_{kct-1} terms in the lagged version). Movements in those wages alter the value of the outside options and, hence, the wage for workers in job ni in city c. However, w_{nict-1} itself enters the outside option for other jobs in the city in $t-1$, making it difficult to determine whether changes in one wage causally affect another. In a related model, BGS show that one can replace the local wage terms w_{kct} with national-level rents by job type, ν_{kt}, in a model-consistent way. This yields an outside option expression that is independent of local wage variation. We adopt the same approach here.

To construct ν_{kt}, we follow a procedure similar to that used for our dependent variable: we regress log wages on the same set of skill and demographic controls H_{lt}, along with a full set of job-type fixed effects. The coefficients on these job-type indicators are interpreted as job-specific (industry–union status) rents, ν_{kt}, and are used to form $\tilde{E}_{nict-1}$:

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

where, in constructing $\tilde{E}_{nict-1}$, we 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.

Constructing $\hat{\mathbf{E}}_{\mathbf{nict}}$ The second issue arises from the potential endogeneity of E_{nict}. Recall that our identifying variation is within industries across cities. As such, concerns arise when the cross-city components of our right-hand-side variables correlate with cross-city differences in productivity. Looking at equation (15), which defines E_{nict}, it is clear that the job-to-job transition probabilities ($\varphi_{kt|ni}$) are not a source of concern since they are defined at the national level and do not vary across cities. The endogeneity problem is thus confined to the city-level employment shares η_{kct} that enter the outside option terms. It seems plausible that these shares are positively correlated with $\Delta\epsilon_{kct}$, the change in productivity in an industry-city cell that is in the error term

We address this issue using a Bartik-style instrument, combining start-of-period employment with national-level growth rates to predict the η_{kct} terms. In addition, E_{nict} suffers from the same reflection problem as E_{nict-1}, which we handle in the same way. This yields our exogenous counterpart to E_{nict}:

$$\hat{E}_{nict} = \sum_{k \neq ni} \hat{\eta}_{kct} \frac{\varphi_{kt|ni}}{\sum_{k'} \hat{\eta}_{k'ct}\varphi_{k't|ni}} \nu_{kt}. \tag{25}$$

Here, $\hat{\eta}_{kct}$ denotes the predicted end-of-period share of employment in city c in job type k. It is constructed by multiplying start-of-period $(t-1)$ employment in each job in city c by the corresponding national-level growth rate between $t-1$ and t, yielding predicted end-of-period job counts. These predicted values are then used to compute employment shares.

Note that $\hat{E}_{nict}$ gets its cross-city variation from variation in the η_{kct-1} terms. 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. While this assumption may not be immediately obvious, in Section 5 we present evidence from a test of an over-identifying restriction implied by the model that supports its validity.

Constructing $\hat{\mathbf{P}}^{\mathbf{ne}}_{\mathbf{ict}}$ The third issue concerns the endogeneity of P^{ne}_{ict}, the proportion of non-union firms under direct threat of unionisation. As highlighted by the model, this threat probability tends to be higher in more productive environments—those with higher ϵ_{ict}—because the rents that unions are expected to capture are larger (see Section 2.6).

To address this endogeneity, we construct a predicted value of the threat probability that is exogenous under the same identifying assumption described above—namely, the random walk assumption for the ϵ_{ict} process. This predicted value is given by:

$$\hat{P}^{ne}_{ict} = \hat{g}(Z_{ict}) \cdot P^{ne}_{ict-1},$$

where $\hat{g}(Z_{ict})$ denotes the predicted growth rate in the probability that a firm is an Emulating Non-union firm, conditional on being non-union, and Z_{ict} is a vector of exogenous factors that influence the probability of a union threat. Specifically, $\hat{g}(Z_{ict})$ is obtained from a regression of the growth in P^{ne}_{ict} on the vector $Z_{ict} = [UA_{it}, UA_{ct}, UA_{it} \times UA_{ct}, RTW_{ct}, R_{ct}]$. UA_{it} and UA_{ct} capture changes in union organizing activity at the industry and city levels, respectively, and are constructed as leave-one-out measures of the decadal growth rate in union elections per establishment at the national industry and city levels, respectively.[16] The premise for UA_{ct} is that if national unions adopt more activist leadership, election drives will rise in the industries where they operate.[17] Meanwhile, UA_{ct} captures potential organizational spillovers across industries in the same city. RTW_{ct} and R_{ct} capture the regulatory environment, with RTW_{ct} indicating whether the state has Right-to-Work laws and R_{ct} indicating whether the Republican Party controlled all three branches of the state legislature. These variables are assigned to cities based on their state.

We approximate P^{ne}_{ict} as the number of firms successfully unionised divided by the number of non-union firms in the same industry-city cell in the previous four years. The underlying idea is that a higher proportion of recent successful unionisations signals a greater threat of unionisation, increasing the share of non-union firms that are Emulators.[18]

[16] We cannot use changes in union organizing at the city×industry level because we expect those to be related to productivity shocks, $\Delta\epsilon_{ict}$. Note, however, 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.

[17] 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).

[18] This follows directly from the model. As discussed in Appendix A.8, 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. 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.

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.[19]

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 of the Appendix, 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 (20). 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.

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.

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

[19] 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 $\Delta((1-P^{ne}_{ict})E_{nict})$ 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.

workers who are switching jobs can move into a union job (captured in the $\varphi_{kt|ni}$ terms). 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, following Lemieux [2006], we restrict the sample to workers aged 20–65 and focus on weekly wages. Industries are grouped using the 1980 Census industrial classification. To create a consistent time series, we use crosswalks from IPUMS and the Census Bureau to map the 1970, 1990, and 2000 industry codes to the 1980 system, resulting in 51 aggregated industry categories. Additional data processing details, including details on the construction of weekly wages, are provided in Appendix C.

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 which are reasonably, though not always perfectly, consistent over time. 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, ending up with 93 geographic areas.

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.[20] We use the county of the unit involved in the election to construct our geographic measures, aggregating counties to our location 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. As described above, we use these proportions as our measure of the proportion of non-union firms that are Emulating Non-union firms, P^{ne}_{ict}, under the theoretically consistent assumption that a greater union drive success rate implies a proportionally larger value for P^{ne}_{ict}. Unfortunately, the election data ends before 2020 and, because there are no establishments in the public sector, we cannot generate the unionisation drive variable for the public sector. As a result, we estimate the full model over the years 1980, 1990, 2000

[20] 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 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 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 η_{kct} (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 $t+1$.[21] 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.[22]

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 are available in Appendix D.

5 Estimation Results

Non-Union Wage Results: Table 1 presents 2SLS results based on specification (20). The dependent variable—the decadal change in non-union wages—is adjusted for education, age, gender, and race. Standard errors are clustered at the city×year level.[23] To ensure robustness, we exclude industry×union status×city cells with fewer than 10 observations and

[21] 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 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 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.

[22] A complication arises for workers observed in the same cell in both years t and $t{+}1$, since we cannot tell whether they changed firms within the cell. If they did, the wage in that cell would be part of their outside option. To estimate the frequency of such moves, we calculate the share of workers with the same cell in both years but different values on job characteristics—such as pay type (hourly vs. non-hourly), worker class (private vs. public), or sub-industry.

[23] See Appendix I for a discussion of standard errors with Bartik instruments, including potential concerns raised in Borusyak et al. [2022] and Adao et al. [2020], which do not arise in our setting.

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)
Bargaining Effects (γ_1):						
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nict}\right)$	0.65***	0.66***		0.65***		
	(0.11)	(0.095)		(0.10)		
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nct\mid ni}\right)$			0.64***			
			(0.12)			
$\Delta\left((1-P^{ne}_{ict})\cdot E_{uct\mid ni}\right)$			0.65***			
			(0.11)			
ΔE_{ucit}					0.41***	0.44***
					(0.15)	(0.15)
Threat Effects ($\tilde{\gamma}_1$):						
$\Delta\left(P^{ne}_{ict}\cdot E_{uict}\right)$	0.77***	0.74***	0.77***	0.77***		
	(0.26)	(0.27)	(0.26)	(0.26)		
Other Effects:						
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
$P^{ne}_{ict-1}\times$ Ind. × Year	Yes	Yes	Yes	Yes		
$P^{ne}_{ict-1}\times$ City × Year	Yes	Yes	Yes	Yes		
Selection controls						
ΔP_{ict} Quadratic	Yes	Yes	Yes		Yes	
First-Stage p-values:						
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nict}\right)$	0.000	0.000		0.000		
$\Delta\left(P^{ne}_{ict}\cdot E_{uict}\right)$	0.000	0.000		0.000		
$\Delta\left((1-P^{ne}_{ict})\cdot E_{uct\mid ni}\right)$			0.000			
$\Delta\left((1-P^{ne}_{ict})\cdot E_{nct\mid ni}\right)$			0.000			
ΔE_{ucit}					0.000	0.000
Over-id. p-values			0.824			

Notes: This table reports 2SLS estimates based on specification (20). Columns (1)–(4) use the decadal change in the regression-adjusted average hourly wage of non-union workers in an industry–city cell as the dependent variable; columns (5)–(6) use the corresponding wage change for union workers. Wage data are constructed from CPS microdata from 1980–2010 across 50 industries and 93 cities. Standard errors, in parentheses, are clustered at the city-year level. Reported first-stage p-values are from Sanderson–Windmeijer tests for instrument relevance; the overidentification p-value is from the Hansen J test.

Column (1) contains the coefficients from our full specification. Based on equation (20) 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. Given our argument that selection is a potential problem, we also include, as a control function, a quadratic in the change in unionisation 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.[24] 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. 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 (e.g., BGS, Caldwell and Danieli [2021]). 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 2, 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 (via η_{kct}) and changes in wage premia for different job types, or, for our instruments, changes in national level job rents (via ν_{kt}). 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.[25] 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

[24] 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 $\frac{\bar{\gamma}_1}{1-\bar{\gamma}_1}$, where $\bar{\gamma} = \frac{1}{2} \cdot (0.65 + 0.77)$.

[25] 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 ν_{kt} and start-of-period η_{kct} 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 (20). 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.

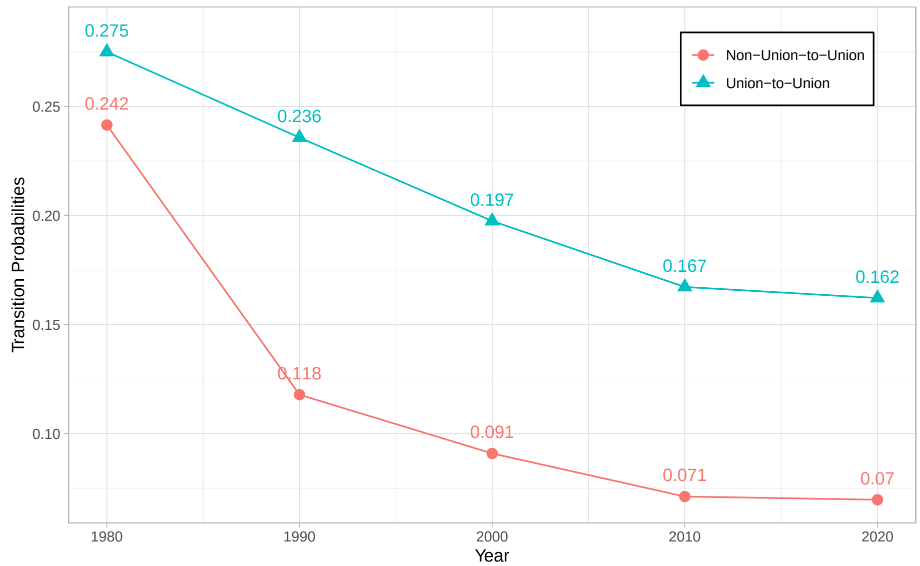

Figure 2: Transitions to Union Jobs

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,

$$E_{nict} = \sum_k T_{kct|ni} w_{kct} = \underbrace{\sum_{i'} T_{ui'ct|ni} w_{ui'ct}}_{\text{Union component}} + \underbrace{\sum_{i'} T_{ni'ct|ni} w_{ni'ct}}_{\text{Non-union component}}, \tag{26}$$

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 (26) 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 $\eta_{\tau i'ct-1}$ terms 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.[26]

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. The estimation challenges and instruments are similar to those for the non-union wage equation and are discussed in Appendix H.1. 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.[27] 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 (20)) for a set of sub-populations defined by gender, age, and

[26] 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.

[27] 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	**Bargaining Effect** (γ_1)	**Threat Effect** ($\tilde{\gamma}_1$)	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 reports 2SLS estimates from specification (20) on separate subsamples. Columns (1) and (2) report coefficient estimates and standard errors for each group; all first-stage p-values for instrument relevance are below 0.03. Column (3) reports the number of observations. Columns (4) and (5) show each group's 1980 unionization rate and union wage premium, respectively, with the latter adjusted for worker characteristics.

education, following on evidence that there is considerable heterogeneity in experiences with 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.

The first two rows present separate results for men and women. These estimates indicate that the bargaining and standard threat effects are broadly similar in magnitude, though the bargaining effect is slightly larger for women. Notably, the gender differences are at least as large as those observed across age or education groups. The following rows show that both types of effects are stronger for younger workers and for those whose highest level of education is a high school diploma or less. In the last four rows, we examine skill-related differences for males, using an approach from Card [2009] to construct skill groups. In this method, each person is assigned weights that reflect their contribution to four groups: young/low-educated, young/high-educated, old/low-educated, and old/high-educated.[28] We focus on men, since they suffered the largest declines in unionisation. The final two columns show that young/low-skilled and old/low-skilled men had particularly high unionisation rates and

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

union wage premiums in 1980. These are therefore the groups where we would expect both union threat and bargaining spillover effects to be especially large, and indeed, the estimated effects are large relative to other groups—particularly more skilled workers, for whom the effects are not statistically significant.

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. 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 provides a way of characterizing the economic significance of our estimated bargaining and threat effects, and offers 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, defined as the weighted average of mean union and non-union wages, where the weight is the proportion unionised in each city, P^u_{ct}:

$$w_{ct} = P^u_{ct} \cdot w^u_{ct} + (1 - P^u_{ct}) \cdot w^n_{ct}, \tag{27}$$

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. To construct city-level wages, we use residualized industry–city wages from our regressions—abstracting from the confounding effects of changes in education, age, and other factors—and combine them with local industry employment shares.[29]

Changes in union strength affect city-level wages through four channels.[30] Figure 3 plots the percentage change in these four components relative to 1980, averaged across cities using population weights. The trends shown, therefore, reflect the evolution of each component for an average city.

The most direct channel is the proportion of unionised workers (P^u_{ct}), shown as 'Proportion Union' in the figure. This captures the reallocation of workers from higher-paid union jobs to lower-paid non-union jobs, holding sector wages fixed—corresponding to the "between" component in standard wage decompositions. The second channel operates through transitions into union jobs. Transition rates, $T_{kct|j}$, reflect changes in mobility frictions ($\varphi_{kt|j}$) and job shares (η_{kct}), both of which affect outside options and thus bargaining and threat effects. In Figure 3, we plot the national average probability that a non-union worker enters a union job ($\varphi_{ui't|ni}$), which we use in place of $T_{kct|j}$ to isolate the primary driving force behind mobility-related wage effects. While obviously related to P^u_{ct}, this rate may evolve differently—for

[29]The wage level equals the mean wage of white workers, holding the proportion of education×gender groups at their 1980 levels.

[30]A fifth channel—selection effects—could theoretically increase observed non-union wages by shifting the productivity mix of non-union firms. However, we find no meaningful evidence of such effects and exclude them from our decomposition.

instance, it may fall faster if older union workers keep their jobs while new job seekers face increasing barriers, or more slowly if union workers retire earlier. In practice, both series decline post-1980, but the transition rate falls more steeply.

The third mechanism is the classic threat effect. In the figure, the line labelled 'Threat' shows the probability that a non-union firm is successfully unionised (P^{ne}_{ict}), capturing one dimension of this channel. This probability fell the fastest of any of the unionisation measures, especially during the 1980s, when the policy environment was strongly against unionisation.

The fourth channel, and perhaps the most interesting line in Figure 3, corresponds to the union wage 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. Our model (echoing an argument in Farber [2005]) provides an explanation for the increase in the premium in the 1980s and, potentially, for 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 (w_{eict}), as in (19). Suppose that larger forces—such as trade or technological change—drive down both w_{nict} and the union wage, w_{uict}, to the same extent. If, at the same time, the threat of unionisation declines, then the observed non-union wage will fall even further, as fewer firms emulate union contracts and the emulation wage itself is lower. This pattern of faster decline in mean observed wages in the non-union sector is what we observe in the 1980s—the decade in which the union threat fell fastest relative to other unionisation measures.[31]

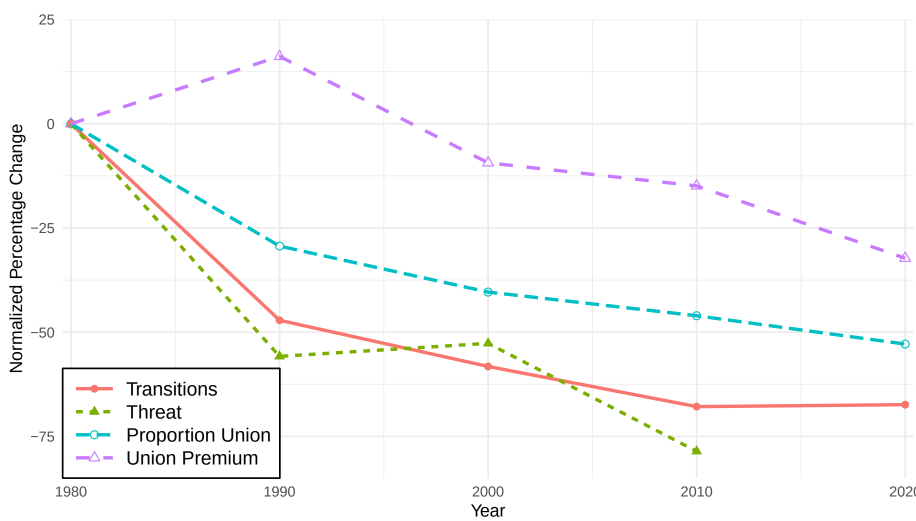

Figure 3: Components of Decomposition

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.

[31] 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.

6.2 Overall Decomposition

We present our decomposition of the overall trend in average city wages in Figure 4. The bottom line in Figure 4 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.

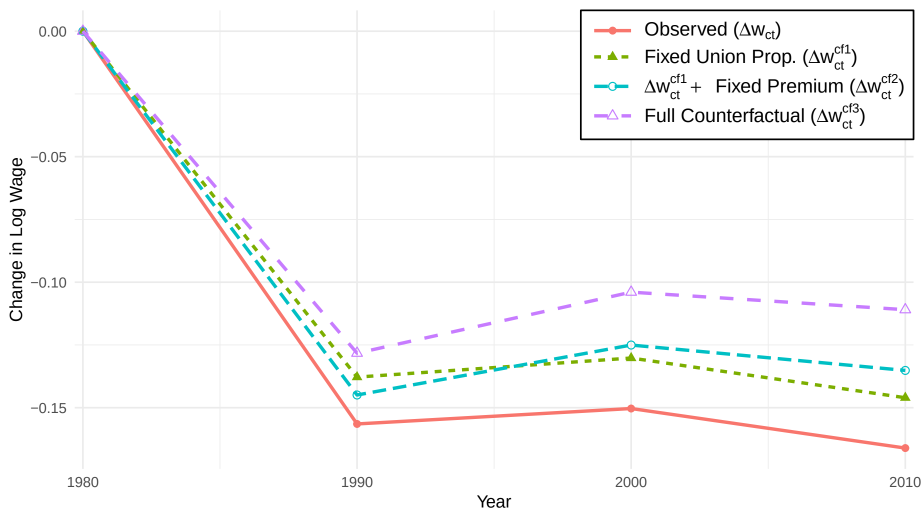

Figure 4: Average Wage Decomposition

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. The figure ends in 2010 because we do not have data on union election outcomes beyond that year. See the main text for further details.

To understand the components in our decomposition, we use (27) to write the change in the city-level mean wage between period t and 1980 as follows:

$$\Delta w_{ct} = \underbrace{\Delta P^u_{ct} \cdot (w^u_{ct} - w^n_{ct})}_{\text{Change in Union Proportion}} + \overbrace{\underbrace{P^u_{c80} \cdot \Delta(w^u_{ct} - w^n_{ct})}_{\text{Change in Union Wage Premium}} + \underbrace{\underbrace{\Delta w^n_{ct}}_{\text{Change in Non-Union Wages}}}_{\Delta w^{cf2}_{ct}}}^{\Delta w^{cf1}_{ct}} . \qquad (28)$$

The first component of our decomposition holds the union proportion fixed at its 1980 level (i.e., sets $\Delta P^u_{ct} = 0$) while allowing other wage-determining factors to vary. We label this counterfactual series Δw^{cf1}_{ct} in Figure 4. It shows that declining unionisation led to a 0.019 log-point drop in mean wages during the 1980s—about 12% of the total decline that decade—with a similar contribution to the overall decline from 1980 to 2010.

Studies of union effects on mean wages often combine the 'shifting weights' component with changes in the union wage premium. Following this approach, we construct a second counterfactual by also holding the union wage premium fixed at its 1980 level—that is, setting $\Delta(w^u_{ct} - w^n_{ct}) = 0$. We label this series Δw^{cf2}_{ct} in Figure 4. As shown in Figure 3, the union premium increased in the 1980s but declined thereafter. As a result, it offset the impact of de-unionisation in the 1980s and reinforced it in later decades. Together, the decline in unionisation and the change in the premium account for a 3.1% drop in mean wages from 1980 to 2010. For comparison, Card et al. [2004] calculate that a standard shift-share analysis implies a 2.6% decline in U.S. mean wages between 1984 and 2001.

A standard decomposition stops at this point. However, our estimates suggest that de-unionisation also 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 in equation (20), which implies that changes in the mean non-union wage are driven by changes in industry wage premia within the non-union sector (γ_{0it}), changes in outside option values (E_{uict} and E_{nict}), and changes in the threat probability (P^{ne}_{ict}), along with other factors such as ΔER_{ct} and an error term. De-unionisation affects the second and third of these channels: outside options shift due to changes in the probability of finding a union job (transition rates, $T_{kct|j}$) and the relative value of union work (the wage premium, $\nu_{uit} - \nu_{nit}$), while changes in P^{ne}_{ict} reflect the decline in the threat of unionisation.[32]

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:

$$\Delta w^{\text{cf2}}_{ct} \equiv \Delta w^n_{ct} = \underbrace{\left[w^n_{ct} - w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}} \right]}_{\text{Non-Union Spillover Effect}} + \underbrace{\left[w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}} - w^n_{c80} \right]}_{\Delta w^{\text{cf3}}_{ct}}. \tag{29}$$

To estimate $w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}}$, we use our estimated wage equations 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 4 because it removes all the channels through which de-unionisation could influence the non-union mean wage. That is, it estimates what non-union wages would have been had threat probabilities, transition rates, and the relative union wage premium remained fixed at their 1980 levels. The final spillover effect—measured as the difference between the observed non-union wage and this full counterfactual—adds an additional 2.4% decline over the period, roughly doubling the total effect relative to the standard decomposition alone.

[32] We use national-level differences in industry wage premia ($\nu_{uict} - \nu_{nict}$) to capture changes in outside options caused by de-unionisation, rather than local differences ($w_{uict} - w_{nict}$). The national premia correspond to ($\tilde{\gamma}_{0it} - \gamma_{0it}$) in our specifications and are treated as exogenous, while local premia are endogenously determined within our model. When attributing changes in transition rates to union decline, we assume that shifts in the industrial composition of non-union workers reflect general economic conditions, whereas deviations in union-sector job growth relative to non-union trends capture union-specific effects. We denote these relative job shares as η^*_{uict}, and compute adjusted transition rates using these shares as $T^*_{kct|j}$.

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.[33] To provide further context 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, 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.[34]

It is worth noting that, in our model, changes in the union wage premium—the second decomposition component—are driven in part by changes in the difference in industrial wage premia between the union and non-union sectors ($\tilde{\gamma}_{0it}$ - γ_{0it}). This channel, reflecting unions' changing share of industry price movements, is treated as exogenous. In contrast, the remaining two factors—relative changes in outside option values and the union threat probability—are forces internal to the model. To assess their role, we construct counterfactual union and non-union wages using 1980 values of the threat and transition variables and the estimated regression coefficients. These model-based forces explain roughly 25% of the decline in the union wage premium, implying that threat and bargaining effects account for half of the overall impact of de-unionisation on average wage declines 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 wage decline more than twice as large as that of women, alongside a much steeper fall in unionisation. As a result, spillover effects are considerably larger for men. Nonetheless, the share of the total wage decline explained by de-unionisation is similar for both genders. Among less-educated men, both younger and older groups experienced sizeable effects from de-unionisation, accounting for 28.5% and 29.7% of their respective wage declines. In contrast, young workers with higher education saw virtually no impact on their mean wage trends. For older, highly educated workers, spillover effects actually raised mean wages. This arises because union jobs for this education group became more concentrated in higher-paying (public sector) jobs, implying

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

[34] This assumes that wages up to 1.2 times the minimum wage are affected by changes in the minimum wage, and uses the fact that 18% of workers earned less than 1.2 times the minimum wage in 1980 (Hardy et al. [2023]).

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) Bargaining Effect							
(4b.1) Fixed Transitions	–0.018	–0.032	–0.008	–0.041	–0.003	–0.061	0.000
(4b.2) 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.

increased average union wages that more than offset declines in the probability of getting a union job in their outside option term.

6.3 Implications for Wage Inequality

Given heterogeneous exposure to de-unionisation across groups, the decline in union power has implications for inequality. Figure 5 shows decompositions for four key log wage differentials: the gender gap, the overall post-secondary premium, and the post-secondary premium separately for men and women. Each series is shown for "Observed" trends (actual national averages) and a "Counterfactual" in which union-related factors remain at 1980 levels.

The gender gap narrowed by about 18 log points over this period, with most of the convergence occurring in the 1980s. As Bidner and Sand [2024] note, the literature on gender wage differentials has overwhelmingly focused on explanations that emphasize women's gains.[35] However, our counterfactual analysis reveals that this narrowing would have been less pronounced without de-unionisation. Since the union decline disproportionately reduced male wages, it indirectly narrowed the gender gap by just under 4%. On the other hand, the post-secondary premium—defined as the wage gap between workers with and without post-secondary education—would have been 6.4% lower had union power remained at 1980

[35] 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 render standard decompositions misleading.

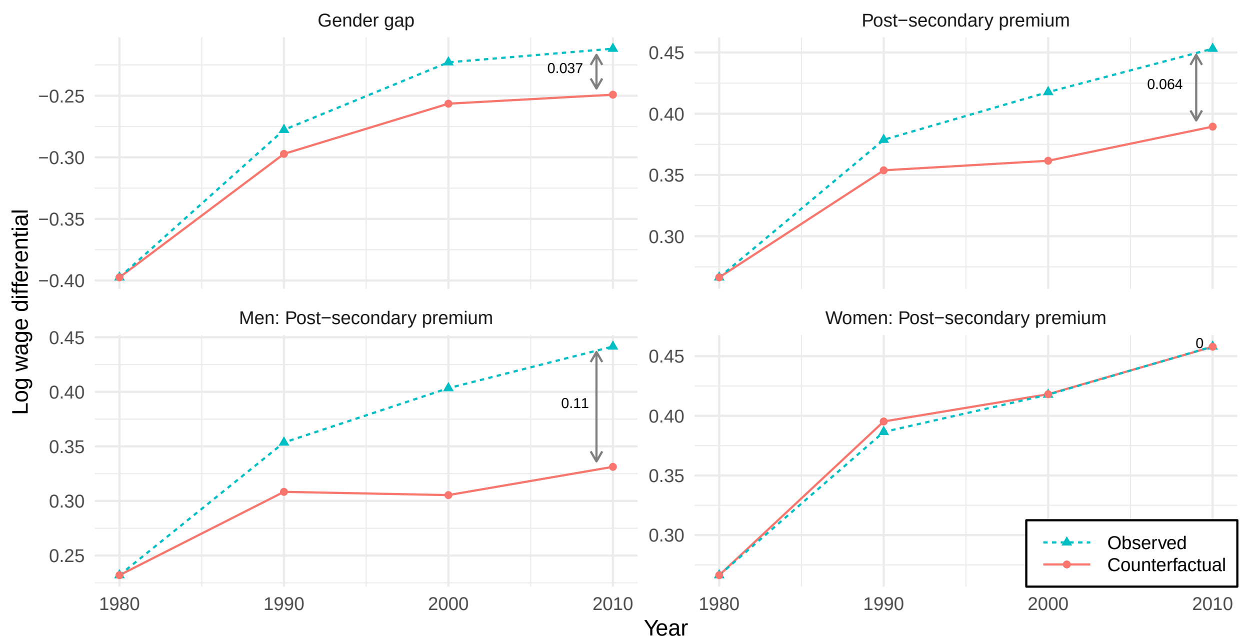

Figure 5: Implications for Wage Inequality

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.

levels. This was driven entirely by men, for whom the gap would have been 11% lower and reflects the disproportionate impact of declining union power on lower-skilled men, amplifying educational wage disparities.

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 (29) and further decompose the Non-union Spillover Effect into its sub-components:

$$\begin{aligned} \text{Non-Union Spillover Effect} = & \underbrace{\left[w^n_{ct} - w^n_{ct|P^{ne}_{ic80},T_{kct|j},\nu_{uit}-\nu_{nit}}\right]}_{\text{Union Threat (4a)}} \\ & \left.\begin{aligned} & + \underbrace{\left[w^n_{ct|P^{ne}_{ic80},T_{kct|j},\nu_{uit}-\nu_{nit}} - w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{uit}-\nu_{nit}}\right]}_{\text{Transitions (4b.1)}} \\ & + \underbrace{\left[w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{uit}-\nu_{nit}} - w^n_{ct|P^{ne}_{ic80},T^*_{kc80|j},\nu_{ui80}-\nu_{ni80}}\right]}_{\text{Union Wage Premia (4b.2)}} \end{aligned}\right\} \text{Bargaining Effect (4b)}. \end{aligned} \tag{30}$$

In Figure 6, we plot the components of our counterfactual non-union spillover effect. The solid line shows the total effect of holding all de-unionisation channels at their 1980 levels. It implies that these factors lowered the non-union wage by about 1.7 percentage points in the 1980s, rising to 2.4 percentage points by 2010. The other lines decompose this total into contributions from each channel.

Figure 6: Decomposition components: Non-union workers

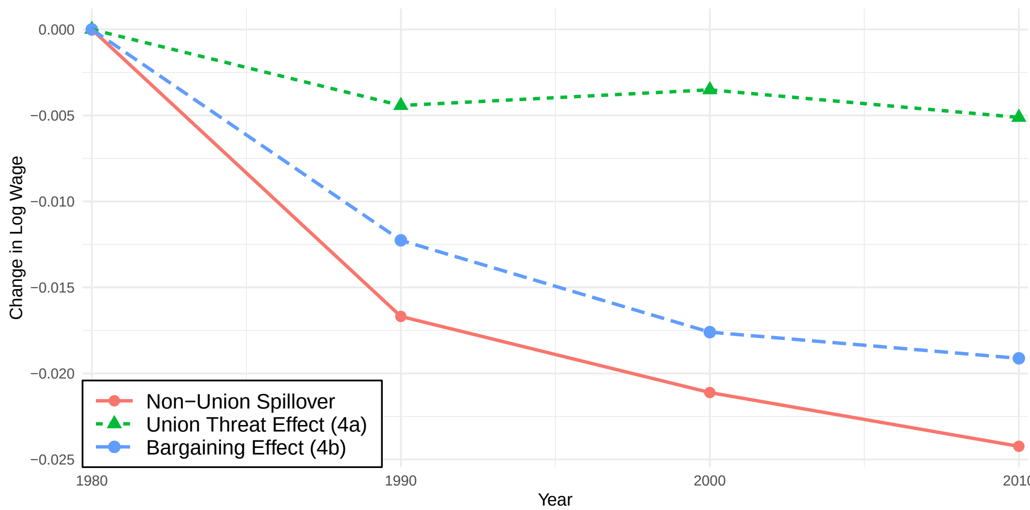

Notes: Data from the CPS and NLRB are shown as log changes relative to 1980 levels. Wages for non-union workers are adjusted for worker characteristics. Each series represents a decomposition component detailed in the main text, in which all union-related factors remain at 1980 levels.

The key takeaway from the decomposition is that the bargaining channel accounts for most of the spillover of de-unionisation onto non-union wages, making up approximately two-thirds in the 1980s and 80% over the entire period. The bargaining component is somewhat smaller in the 1980s because, as we saw earlier, the union premium actually rose during that decade, implying that non-union wages would have been higher at the end of the decade if the union wage premium were held at its 1980 level. Over the entire period, the bargaining effect is nearly entirely due to reductions in the probability of transitioning to union jobs (component 4b.1 in equation (30)), as seen in the lower panel of Table 3.

The union threat component contributes relatively little to the total effect, even though our estimates show clear evidence of a standard threat effect, largely because the threat probabilities were small even in 1980. That the bargaining channel is much more important has important policy implications: unlike the threat effect, which depends on expanding unionisation, the bargaining channel can be activated by policies that improve workers' outside options more generally—such as banning non-compete clauses [Johnson et al., 2020] or expanding commuting options [Hafner, 2022].[36]

The second panel of Table 3 presents the components of the non-union spillover effect over the full 1980–2010 period for different subgroups. The spillover effect for men is more than twice as large as that for women, with most of the difference driven by variation in the transition effect. Similarly, the greater impact of de-unionisation on low-educated non-union workers reflects their lower chances of transitioning into union jobs, rather than differences in the likelihood that their firm would become unionised.

[36] This echoes findings by Caldwell and Danieli [2021], who show that wages rise with an index of outside options, including the probability of switching jobs or occupations. In our setting, this outside option is the ability to transition into a union job.

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 these 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 raise 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], Tschopp [2017], 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.

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.

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.

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. *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.