NORTHWESTERN UNIVERSITY

Essays on Labor History

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Economics

By

Luca Bittarello

EVANSTON, ILLINOIS

September 2019

Abstract

This dissertation contains three empirical studies in economic history and labor economics.

The first chapter discusses two sources of historical data on work stoppages in the United States: the *Third Annual Report of the Commissioner of Labor* (1888) and the *Tenth Annual Report of the Commissioner of Labor* (1896). It describes a new transcription of the strike tables, which includes all rows for the first time, and provides instructions for users. Four replication exercises illustrate the advantages of the new file.

The second chapter uses the the data from the *Third Report* to test whether labor unions help workers win strikes. Unorganized workers were still responsible for two fifths of all strike activity in the United States in the early 1880s, which allows me to identify the effect of unions on strike outcomes. Because organized workers might attempt riskier confrontations than the unorganized, I construct an instrument for the involvement of a union in a strike from the location of the assemblies of the Knights of Labor. I estimate that unions raised strikers' success rate by 32 percentage points from a baseline of 38 percent; moreover, they decreased the incidence of job loss by 22 percentage points from a baseline of 56 percent. Although unions increased the probability that employers acceded to strikers' demands, I find no evidence of an impact on the size of those concessions.

The third chapter evaluates how an increase in the supply of skilled labor affects task assignment within and between occupations. Guided by a simple theoretical framework, Francis Kramarz, Alexis Maitre and I exploit detailed information about individual workers' tasks from multiple surveys to examine the impact of a twofold rise in the share of university graduates in the French workforce between 1991 and 2013. Our identification strategy uses variation in the

change in the graduate share across local labor markets. We find that higher average educational attainment is associated with more routine, fewer cognitive and fewer social tasks within occupations and with fewer routine, more cognitive and more social tasks across occupations.

Preface

Labor markets occupy a special place in the economy. Most active citizens put more hours into wage work than any other pursuit. Finding and holding a job is their foremost economic concern, though the task is simpler for some than others.

This dissertation explores two topics in economic history and labor economics. The first two chapters examine industrial conflict. In particular, the second chapter analyzes the impact of union intervention on strike outcomes in the U.S. in the early 1880s. It takes advantage of an experiment in radical mass unionism to isolate unions' true effect from changes in members' behavior. I find that unions increased the probability of success of a strike, which emboldened workers to undertake riskier confrontations. The third chapter considers the assignment of tasks across employees in France between 1991 and 2013. On the one hand, higher educational attainment increased the incidence of cognitive tasks by swelling the ranks of skilled occupations. On the other hand, oversupply reduced the pay for cognitive work, so workers spent less time on cognitive tasks within each occupation.

These chapters differ in subject and setting. Nonetheless, there are similarities. First, all are empirical exercises. They infer economic behavior from a mass of individual observations at particular points of a historical event – whether the growth of the Knights of Labor in the U.S. or the expansion of tertiary education in France. They required much attention to the details of data collection, data processing and institutional context. Second, both main chapters take the impact of policies on incentives into account. Workers respond strategically to changes in labor markets: for example, organized workers confronted stronger employers than the unorganized because they could count on the support of their union. Economic intuition was key in identifying such effects and establishing causality. Third, all chapters share a concern for workers' efforts to improve their own lot – whether through industrial action or career choices.

As I wrote this dissertation, a great many helped me with indispensable support and kind advice. I can not thank them enough.

I must first thank my family: my father, André Luis, my late mother, Cristina Maria, and my brother, Miguel. They prepared me for the challenges ahead and were infallible in their support at every turn.

I am grateful to Sciences Po for providing me with excellent instruction and broadening my horizons. I would not gone far without that initial scholarship. I also met my first advisor at Sciences Po, Francis Kramarz. He did not only guide me as I discovered research, but he is responsible for my abiding interest in labor economics. In addition, he is a coauthor of the third chapter in this dissertation. I owe him much and more.

I am also indebted to my committee beyond measure. Joseph Ferrie introduced me to economic history. His anecdotes were a source of inspiration for my research; moreover, he was always ready with such practical advice as the location of a data set or speaking faster. Joel Mokyr pushed me to see the big picture and be more ambitious in my work. He was so attentive as to check in on me after he had not heard from me in a while (more than once at that). He was a model proofreader too. Matthew Notowidigdo taught me identification. He helped me interpret my results, motivate them and present them to a broader audience. He could often see the connections between my own projects more clearly than me.

The Center for Economic History provided me with generous funding and fostered an intellectually stimulating environment at Northwestern University. This dissertation would have poorer without it.

I must finally thank my other friends and relatives. I should not list them all here, but their support was essential in getting me through graduate school. I am forever grateful.

Contents

List of Tables

List of Figures

1 A Guidebook to the Strike Data from the Commissioner's Annual Reports

1.1 Introduction

Economists and historians have long borrowed historical strike data from the reports of the U.S. Bureau of Labor. Two sources are especially valuable: the *Third Annual Report of the Commissioner of Labor* (U.S. Bureau of Labor, 1888), which covers the period from 1881 to 1886, and the *Tenth Annual Report of the Commissioner of Labor* (U.S. Bureau of Labor, 1896), which covers the years from 1887 through the first semester of 1894. They contain detailed information on 16,694 observations between strikes and lockouts. Other reports extend the series through the 1950s, but they only provide aggregate statistics, while postwar microdata say little about the bargaining units or the resolution of each conflict.

This paper presents a new transcription of the *Third* and the *Tenth Report*. It includes all rows of the strike tables for the first time. The lockout table of the *Third Report* was also digitized. Furthermore, observations were geolocated. I describe each variable in the reports and discuss the limitations of the data. To illustrate the advantages of the new file, I replicate four important studies: Card and Olson (1995), Currie and Ferrie (2000), Friedman (1988) and Rosenbloom (1998). Qualitative results are generally robust, though point estimates differ.

I am indebted to Joseph Ferrie, Gerald Friedman, Joel Mokyr, Matthew Notowidigdo and Joshua L. Rosenbloom for advice. Joseph Ferrie and Joshua L. Rosenbloom also shared data for this project. Priyanka Panjwani provided invaluable research assistance. All mistakes are mine.

1.2 Context

Strike statistics were first gathered in the late 1870s (Weeks, 1886). Data collectors were partly motivated by perceived increase in the incidence of work stoppages, by their impact on production and by the search for alternative conflict resolution strategies. There was also broad public interest in the "labor question" in the aftermath of such insurrections as the Paris Commune and the Great Railroad Strike. For example, the *Third Report* begins: "The industrial disturbances which have been so frequent in this country since 1877 really establish the period as one of strikes and lockouts." Bevan (1880) writes that he was able to "make an aggregate of the number of labour disputes, which may perhaps startle those who have engaged in them, if they ever do happen to reflect upon the enormous hindrance to labour and trade that these quarrels represent".

Bevan (1880) secured information on ten years of strikes across the United Kingdom, obtaining the earliest sample for an entire country (Weeks, 1886). American data were first assembled by Ohio's Bureau of Labor Statistics (1878), Massachusetts's Bureau of Statistics of Labor (1880) and Pennsylvania's Bureau of Industrial Statistics (1882).[1] Weeks (1886) followed their efforts with a postal inquiry for the *Tenth Census* of 1880, which yielded the first nationwide statistics for the United States.

The U.S. Congress established the Bureau of Labor in 1884. Carroll D. Wright became its first commissioner, having previously led the labor bureau of Massachusetts. It began recording work stoppages in 1886, as unions prepared the May strike wave for the eight-hour day. There were four reports on strikes and lockouts: the *Third*, the *Tenth*, the *Sixteenth* and the *Twenty-first*.[2] The Bureau did not collect strike data between the *Twenty-first Report* in 1906 and the

[1] Chapter IV in the *Third Report* recapitulates these data.

[2] Many other countries published official microdata in varying detail before the First World War. For sources, see the *Twenty-first Report*. French and Swedish data have been digitized: see Enflo and Karlsson (2018), Karlsson (2019) and Tilly and Jordan (2012). Canadian postwar microdata are available in electronic form (work stoppages since 1946 and collective bargaining agreements since 1977).

first editions of the *Monthly Labor Review* in 1915 (Peterson, 1938).[3] It has thereafter published stoppage data without interruption, including limited microdata;[4] however, the series have often changed in scope and methodology.[5]

1.3 Methodology

The Bureau's initial data collection was retrospective. Since Weeks (1886) had already obtained data for 1880, it spanned the period from 1881 to 1886. After the publication of the *Third Report*, the Bureau gathered information on a rolling basis.

Agents collected data in two stages. They first compiled a list of work stoppages from newspapers, trade journals, etc. They then visited every locality on the list for canvassing. They sought accounts from both sides of each dispute; as the case warranted, they also interviewed journalists and other third sources. Moreover, the Commissioner instructed agents to ask interviewees about any additional conflicts in the area. This procedure had three goals: securing as many details about each confrontation as possible, reconciling contradictory information and improving coverage. Strikes of less than a day's duration were later discarded.

Each report contains separate tables for strikes and lockouts. The strike table of the *Third Report* encompasses 5407 complete rows; the lockout table, 358. The strike table of the *Tenth Report* encompasses 10,487 complete rows; the lockout table, 442. (Subsection 1.4.2 discusses the unit of observation.) Table 1.1 presents the contents of the reports at length.

[3] So far as I was able to ascertain, the Bureau published its last annual report in 1911.

[4] Postwar samples include the *Work Stoppages Historical File* (1953 to 1981) and the *Work Stoppages Program* (since 1993). The Federal Mediation and Conciliation Service provides data as well (since 1984). These files are available in digital format. Section 1.2 discusses their scope and methodology. See McConnell (1990) for other sources.

[5] The annual reports included any stoppage of a day or more. The *Review* initially included all strikes and lockouts, but it limited the sample in 1927 to a minimum of six workers and a day's duration. The Bureau raised the threshold to a thousand workers in 1982. Whereas it undertook field interviews for the annual reports (q.v. Section 1.3), it used postal inquiries for later data.

Table 1.1: Contents of the Commissioner's reports on strikes and lockouts

Column	Availability	Notes
Year	Both	
State	Both	The *Third Report* does not distinguish between the Dakotas, but the transcription does. The Bureau split general stoppages by states as far as possible and noted exceptions (mostly in the rail industry in the *Tenth Report*).
Locality	Both	Localities are not listed in any obvious order, which suggests that the most affected locality comes first and so forth.
Industry	Both	The reports use different classifications. The transcription classifies the miscellaneous category in greater detail than the reports. Affected observations are flagged, so one can recover the original classification. The Bureau split general stoppages by industry (except for the general strike in New Orleans in 1892).
Occupation	Both	This column was only partly transcribed in the case of the *Third Report*. Rows are flagged if workers were described as "employees", "helping hands" or "laborers".
Cause or object	Both	Only an aggregate classification is available for most rows from the *Third Report*. Observations are flagged if the stoppage was defensive (according to the first demand) or if multiple demands were made.
Ordered by organization	Both	It is unclear if the Bureau flagged only stoppages ordered by organizations or if it flagged all stoppages involving organizations. This variable refers to labor unions in the case of strikes and to manufacturers' associations in the case of lockouts.
Number of establishments	*Third Report*	The *Third Report* splits stoppages into multiple rows if affected establishments shut down for different lengths.
Number of closed establishments	*Tenth Report*	This column can be replicated for the *Third Report* by combining the number of establishments and the number of days closed (vide supra).
Number of open establishments	*Tenth Report*	This column can be replicated for the *Third Report* by combining the number of establishments and the number of days closed (vide supra).
Days closed	*Third Report*	
Beginning date	Both	
End date	Both	According to the *Tenth Report*, it is the date by which most workers had either returned to work or been replaced.

Continues...

Contents of the Commissioner's reports on strikes and lockouts (continued)

Column	Availability	Notes
Duration	Both	The *Third Report* splits stoppages if the conflict ended at different dates across establishments. The *Tenth Report* gives the average duration across establishments. Unlike the *Third Report*, the *Tenth Report* does not give the duration if affected establishments closed permanently.
Succeeded	Both	This column reads "succeeded" if all demands were met, "partly" if only some were met and "failed" if none were met. The transcription separates these outcomes, since outcomes sometimes differed across establishments. Flags were also added for incomplete stoppages and permanently closed establishments.
Employees' loss	*Third Report*	The report does not define this variable. It apparently subtracts compensatory overtime and forgiveness for lost time by employers from workers' wage loss. It is not comparable to the related variable in the *Tenth Report*.
Employers' wage loss	*Tenth Report*	The report does not define this variable. It is not comparable to the related variable in the *Third Report*.
Employees' assistance	Both	This column includes the payment of defense funds by labor unions, donations, etc.
Employers' loss	Both	This column includes losses as a consequence of lost work and boycotts.
Employees	Both	By gender. Before and after stoppage.
Average daily wages	*Third Report*	By gender. Before and after stoppage.
Employees for whom strike was undertaken	*Tenth Report*	By gender.
Strikers or employees locked out	Both	The *Third Report* gives the total number. The *Tenth Report* splits it by gender.
Strikers' or locked-out employees' daily pay	*Third Report*	Before and after stoppage.
Idled employees	Both	By gender. This column has different labels in each report ("employés striking and involved", "employés locked out and involved" and "employees thrown out of employment"). It includes voluntarily and involuntarily idled workers.
New employees after strike	Both	By gender. This column does not include temporary strikebreakers.
New employees brought from other places	Both	This column does not include temporary strikebreakers.
Weekly working hours	Both	Before and after stoppage.

1.4 Limitations

The Commissioner's reports exhibit two deficiencies: incomplete coverage (Subsection 1.4.1) and inconsistencies in the unit of observation (Subsection 1.4.2).

1.4.1 Exhaustiveness

The Commissioner believed that the Bureau achieved a near census of strikes and lockouts (U.S. Bureau of Labor, 1888). However, Bailey (1991) shows that it undercounted work stoppages in fact. Following the Bureau's methodology, he found that local newspapers record at least twice as many strikes in Terre Haute as the reports.[6] He could not identify a clear pattern in excluded disputes, though short and wildcat strikes seem overrepresented. The two reports exhibit similar omission rates, though the *Third* used a retrospective survey. As Card and Olson (1995) observe, missing data do not jeopardize statistical analysis so long as they are random.

For additional evidence, I repeated the exercise for Chicago, Decatur and Milwaukee. Unlike Bailey (1991), I restricted the search to the first six months of 1881. I found six missing strikes in Chicago (against 46 in the report), none in Decatur (vs. 3) and three in Milwaukee (vs. 3). It appears that fifty percent is an upper bound on the omission rate.

1.4.2 Consistency

It is often difficult to present a general strike in tabular form (Peterson, 1938). Even if strikers coordinate across establishments, each action might still differ in its demands, dates, resolution, etc.

In planning the *Third Report*, the Bureau therefore chose the establishment as the unit of observation. This design traded the aggregation challenge for a printing one: the line count quadrupled. The Bureau compromised by grouping related standoffs to the extent that they

[6] Bailey (1991) commits mistakes too: e.g., the *Third Report* did record the telegraphers' strike of 1883.

coincided in dates and outcomes. As a result, the *Third Report* contains 5407 rows for 3902 strikes across 22,304 establishments (U.S. Bureau of Labor, 1888). A row may thus represent either a strike against a single establishment, a general strike or part of one. Although this strategy earned praise at the time (Smith, 1888), the consequent ambiguity complicates the interpretation of the data.

The Bureau had a change of heart as it prepared the *Tenth Report*. The unit of observation became the strike or lockout. Where dates differed, the Bureau recorded the first starting date, the last ending date and the average duration across establishments; where outcomes differed, the Bureau added footnotes. This redesign eliminated ambiguity. However, it is less informative and reduced comparability with the *Third Report*.[7]

Empirical studies have generally ignored this change. Weighting is a possible fix. If one weights observations by the number of affected establishments or their prestrike workforce, both samples yield results in terms of a consistent unit. On the other hand, estimates may be sensitive to outliers. Winsorized weights are a sensible compromise, but the choice of cutoff is not obvious and even winsorized weights yield less precise estimates than their unweighted counterparts (cf. Section 1.6). Resourceful researchers might also be able to aggregate all observations to the strike level, especially as the *Third Report* gives the number of strikes per year.

The design of the *Third Report* has consequences for inference as well. Whenever several rows pertain to the same strike, they are likely to share outcomes and characteristics. Inference should therefore take cross-correlation into account. It is reasonable to assume that cross-correlation decreases over time and across space. Multiple inference strategies are valid in this context. One could cluster at the intersection of a time unit (e.g., the year) and a geographic unit (e.g., the county), though clustering requires an implausibly sharp discontinuity between

[7] Note that neither report gives the name of affected establishments. As a consequence, it is not possible to link the data with other sources. Moreover, it is not possible to identify repeated establishments.

units. One could also base inference on Conley (1999) via either a composite distance measure or a multidimensional generalization.

Table 1.2 shows summary statistics for both samples in an attempt to assess their comparability. Because strike characteristics evolve for substantial reasons too, it presents each measure for two time frames: all years and six months near the change in the unit of observation.[8] The first four rows show the percentages of observations across multiple establishments, across multiple localities, with multiple causes and without strikers' specific trade. One expects larger shares in the *Tenth Report*, yet the differences are small and no pattern emerges. The next two rows show the percentages of observations in which some establishments remained open while others closed or in which outcomes differed across establishments. They are a small fraction of all rows in the *Tenth Report*, but they represent an important fraction of all affected establishments. (The *Third Report* does not group such cases.) The last six rows show medians and interquartile ranges for the number of strikers, the number of idled workers and the duration. Contrary to expectations, the unweighted statistics are higher in the *Third Report*. Weighting mostly reduces the differences in proportional terms, reversing it in some instances.

1.5 Transcription

So far as I am aware, Gerald Friedman prepared the first digital transcription of the Commissioner's reports on work stoppages in the 1980s. He sampled one in five observations in the strike table of the *Third Report* and one in ten in the *Tenth* (Rosenbloom, 1998). This file was the basis for Friedman (1988) and Rosenbloom (1998). Card and Olson (1995) began a parallel digitization effort in the meanwhile. The initial sample comprised all strike rows in the *Third Report* from Illinois, Massachusetts and New York.[9] Currie and Ferrie (2000) added

[8] I do not use the year 1886 because the May strike wave distorts the estimates.
[9] Card and Olson (1995) exclude 77 rows from their sample because of "clerical errors". I could not identify these errors. They may have been transcription mistakes.

Table 1.2: Comparability of the Commissioner's reports on strikes and lockouts

	Unweighted				Weighted			
	Third Rpt. All years	*Tenth Rpt.* All years	*Third Rpt.* 7–12, 1885	*Tenth Rpt.* 1–6, 1887	*Third Rpt.* All years	*Tenth Rpt.* All years	*Third Rpt.* 7–12, 1885	*Tenth Rpt.* 1–6, 1887
Binary variables								
Multiple establishments (%)	21	21	19	23	81	82	72	86
Multiple localities (%)	2	3	7	2	3	10	5	8
Multiple causes (%)	9	8	3	9	18	18	8	26
No specific trade (%)	27	20	28	24	19	24	22	20
Partial shutdown (%)	—	4	—	5	—	21	—	26
Multiple outcomes (%)	—	2	—	3	—	15	—	13
Continuous variables								
Strikers (median)	50	30	57	44	519	471	325	450
Strikers (IQR)	110	83	126	107	1900	1839	1427	3385
Idled workers (median)	68	38	95	54	650	580	364	575
Idled workers (IQR)	167	116	180	155	1874	2013	2863	3861
Duration (median)	10	6	11	7	12	12	11	11
Duration (IQR)	25	12	27	13	22	28	33	15

Notes: Lockouts are excluded. The last four columns weight rows by the number of establishments (truncated at 75).

observations from the *Tenth Report* and eleven other states,[10] which increased coverage from 14 to 82 percent of all rows. This file was later used by Geraghty and Wiseman (2008), Naidu and Yuchtman (2018) and Schmick (2018).

This project expands the data set of Currie and Ferrie (2000). It involved four steps.

First, I added all remaining strike rows. Researchers now have access to the strike entire microdata from the reports in electronic form for the first time. The table of lockouts in the *Third Report* was also digitized.

Second, I fixed transcription errors. To locate them, I took advantage of the summary tables at the end of each report. These tables show such aggregates as the number of affected establishments for each state by years, for each state by industries, etc. The new file allows their full reproduction, so only such mistakes remain as cancel each other within narrow cells. Note however that I could not resolve the following discrepancies:

- *Third Report*: the number of closed establishments is off by one and the number of days closed is off by five for strikes in the clothing industry in Georgia in 1885.
- *Third Report*: aggregate duration is off by two days for strikes in the glass industry in Pennsylvania in 1886.
- *Third Report*: aggregate duration is off by 170 days for lockouts in the printing industry in Texas in 1885.
- *Tenth Report*: aggregate duration is off for strikes in the construction industry in New York by two days in 1891, by eight days in 1892 and by three days in 1894.
- *Tenth Report*: aggregate duration is off by thirty days for strikes in the construction industry in Pennsylvania in 1890.
- *Tenth Report*: the number of idled employees is off by a hundred for strikes in the metals and transportation industries in Illinois.

[10] To wit: Connecticut, Delaware, Indiana, Maine, Maryland, Michigan, New Hampshire, New Jersey, Ohio and Pennsylvania.

These discrepancies may reflect tabulation errors, printing errors or overlooked transcription errors.

This procedure was not applicable to dates, hours or wages, which are not part of the summary tables. I performed basic checks instead: the duration should equal the difference between the starting and ending dates, wages should not change after an unsuccessful strike for a wage raise, etc.

Third, I fixed a few obvious mistakes in the reports. For instance, the *Third Report* gives June 31 as a starting date. I flagged affected observations. Furthermore, I split the loss variables among rows of a general stoppage when the report failed to. I apportioned losses according to the number of affected employees and duration.

Fourth, I geolocated all observations. Geolocation is a work in progress: some coordinates are approximate and others are surely wrong. Four difficulties obtain: some localities became ghost towns and disappeared; others became part of expanding cities and lost their identity; others changed names; and others shared names with one or more different places in the same state. I resolved such cases through contemporaneous maps and narrative evidence from newspapers, government publications, etc.

1.6 Replications

This section presents four replication exercises. They illustrate two advantages of the new file: the inclusion of all strike rows and the correction of transcription mistakes (cf. Section 1.5). They follow the methodology of each paper as closely as possible: e.g., they use the original inference strategies and industrial classifications. I present both unweighted and weighted estimates (cf. Subsection 1.4.2).

TABLE 1.3: DETERMINANTS OF THE SUCCESS RATE (FRIEDMAN, 1988)

	Original coefficients		Unweighted replication		Weighted replication	
Union strike (1881–1894)	0.14	(0.91)	0.33	(5.51)***	0.64	(4.27)***
Union strike (1887–1894)	0.57	(2.93)***	0.14	(1.91)*	−0.24	(1.33)
Log of establishment size	−0.14	(4.23)***	−0.13	(10.15)***	−0.16	(4.57)***
Big city	0.44	(4.39)***	0.30	(8.25)***	0.33	(3.38)***
Striker rate by industry and state	0.03	(2.11)**				
Issue effects	3		3		3	
Industry effects	2		5		5	
Region effects			3		3	
Year effects	14		14		14	
Observations	2,052		15,880		15,880	
Parameters	24		29		29	

Notes: The table shows coefficients from logistic regressions and *t*-statistics (in parentheses in absolute value). The second column shows coefficients from Table 4 in Friedman (1988). The fourth column weights estimates by the number of establishments (truncated at 75).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

1.6.1 Friedman (1988)

Friedman (1988) conducts a comparative analysis of the impact of unionization on strike outcomes at the end of 19th century in France and the United States. He argues that industrial unions had little effect on the probability of success in the U.S. To do so, he estimates the effect of union sponsorship on the success rate in two subperiods: from 1881 to 1886, when industrial unions dominated, and from 1887 to 1894, when craft unions became dominant again.

This paper is difficult to replicate because it offers few details on covariates. It does not define the "striker rate by industry and state", nor does it explain its categorization of cities, industries or strike causes. I group causes and cities according to Rosenbloom (1998), whose data are due to Friedman (1988). Rosenbloom uses a finer industrial classification than Friedman, so I group industries into five sectors ad hoc instead.[11] I use region effects in lieu of the striker rate.[12] My results are robust to the choices of classification and to the use of interaction terms.

[11] To wit: construction, mining, heavy manufacturing, light manufacturing and services.

[12] My regions are: New England, the Mideast, the Great Lakes, the Plains, the South and the West. They correspond to the BEA regions, but I split the Southwest between the South (Texas) and the West (others).

Table 1.3 shows the original coefficients and mine. Friedman (1988) estimates that unions increased the odds ratio by 0.14 log points from 1881 to 1886 and by 0.71 from 1887 to 1894. The effect is insignificant in the first period. I find instead that unions increased the odds ratio by 0.33 log points from 1881 to 1886 and by 0.47 from 1887 to 1894. The coefficient is significant at the one-percent level in the first period, while the change between periods is only marginally significant.[13] Weighting exacerbates these differences: the effect becomes 0.64 in the first period and 0.40 in the second. Other coefficients are more robust: e.g., Friedman (1988) puts the impact of establishment size at –0.14 log points, whereas I find –0.13 without weights and –0.16 with them.

1.6.2 Rosenbloom (1998)

Rosenbloom (1998) studies the use of strikebreakers in the late 19th century.[14] Unlike Friedman (1988), he does not test a particular hypothesis.

The classification of occupations by skill level is the only noteworthy difference between this exercise and the original. Rosenbloom borrows his from Edwards (1933). I approximate it by treating strikers as skilled if they worked in the manufacturing sector and the reports specified their trade.[15] Skilled strikers constitute 37 percent of his observations and 32 percent of mine (by our respective definitions).

Table 1.4 shows the original coefficients and mine. Note that Rosenbloom (1998) does not report a fixed effect for 1892, which is a typographical error in all likelihood, since it appears in other regressions. The estimates are quite different in magnitude, yet all significant coefficients

---

[13] I am able to replicate the essence of Friedman's results with Rosenbloom's data file, which is an updated version of Friedman's. These discrepancies with respect to the new file may thus be due to random sampling or transcription mistakes. Note that my results are robust to doubling the weights on the data from the *Third Report*, which Friedman samples at twice the rate of the *Tenth*.

[14] Note that the reports give the number of new employees at the end of the strike, so we do not know whether firms hired temporary strikebreakers.

[15] Workers without a specific trade are "employees", "laborers" or "helping hands".

TABLE 1.4: DETERMINANTS OF THE USE OF STRIKEBREAKERS

	Original coefficients		Unweighted replication		Weighted replication	
Strike characteristics						
Number of employees (log)	0.0435	(3.597)***	0.1054	(8.793)***	0.1192	(3.402)***
Number of strikers (log)	–0.0299	(2.412)**	–0.0709	(5.697)***	–0.0263	(0.744)
Union strike (1881–1894)	0.0176	(0.468)	–0.0473	(1.249)	–0.1830	(1.844)*
Union strike (1887–1894)	–0.0872	(1.826)*	0.0886	(1.928)*	0.2111	(1.850)*
Skilled occupation	0.0413	(1.354)	–0.0135	(0.422)	0.0505	(0.562)
Cause						
Hours reduction	0.0180	(0.383)	–0.1694	(6.718)***	–0.0963	(1.331)
Wage increase	–0.0417	(1.574)	–0.0722	(1.706)*	–0.0184	(0.175)
Defense	–0.0434	(1.219)	–0.1636	(5.216)***	0.0253	(0.285)
Industry						
Boots and shoes	–0.0383	(0.656)	–0.1433	(2.578)***	–0.0702	(0.626)
Clothing	0.0449	(0.889)	–0.1153	(2.496)**	0.0445	(0.341)
Construction	–0.1191	(2.756)***	–0.3509	(8.610)***	–0.1251	(1.259)
Food preparation	0.1311	(1.605)*	0.3319	(4.297)***	0.4373	(2.285)**
Furniture	0.0286	(0.415)	0.0877	(1.315)	0.2993	(1.696)*
Glass	–0.1088	(1.357)	–0.2007	(2.304)**	–0.0699	(0.412)
Machinery	–0.0613	(1.367)	0.0168	(0.212)	–0.0312	(0.169)
Metals and metallic goods	–0.0652	(0.817)	–0.0521	(1.299)	–0.0426	(0.385)
Mining	–0.1784	(3.827)***	–0.4104	(8.602)***	–0.1486	(1.239)
Printing and publishing	0.1836	(2.523)**	0.6740	(9.715)***	0.9351	(6.923)***
Stone quarrying and cutting	–0.0442	(0.710)	–0.1751	(3.077)***	0.0790	(0.571)
Tobacco	0.0198	(0.395)	–0.0356	(0.735)	0.0821	(0.595)
Transportation	0.1036	(1.890)*	0.3153	(6.308)***	0.4381	(3.510)***
Wooden goods	–0.0091	(0.117)	0.1945	(2.315)**	0.5196	(2.709)***

Continues...

Determinants of the use of strikebreakers (continued)

	Original coefficients		Unweighted replication		Weighted replication	
City size						
25,000–49,999	–0.0042	(0.087)	0.0571	(1.288)	0.0598	(0.596)
50,000–99,999	–0.0077	(0.169)	0.0744	(1.773)*	0.3697	(3.827)***
100,000–249,999	0.0712	(1.615)*	0.1395	(3.405)***	0.1530	(1.521)
250,000 or more	–0.0666	(2.365)**	–0.0354	(1.257)	0.1383	(1.868)*
Region						
Midwest	–0.0234	(0.891)	–0.0188	(0.750)	–0.2775	(4.067)***
South	0.0004	(0.009)	0.0137	(0.322)	–0.3135	(2.925)***
West	–0.0180	(0.160)	0.0508	(0.794)	0.0955	(0.636)
Year						
1882	0.0917	(1.209)	0.0639	(0.848)	–0.1497	(0.738)
1883	0.1769	(2.430)**	0.1961	(2.686)***	–0.2189	(1.140)
1884	0.0484	(0.649)	0.1321	(1.759)*	0.0624	(0.303)
1885	0.0124	(0.182)	0.0399	(0.580)	–0.1641	(0.962)
1886	0.0688	(1.160)	0.0670	(1.134)	–0.0367	(0.235)
1887	0.1001	(1.392)	0.0093	(0.138)	–0.0278	(0.165)
1888	0.1435	(1.807)*	0.0132	(0.181)	–0.0870	(0.468)
1889	0.0723	(0.953)	–0.0136	(0.192)	–0.2380	(1.264)
1890	0.0381	(0.540)	–0.1431	(2.149)**	–0.1328	(0.800)
1891	0.0725	(0.994)	–0.0534	(0.794)	0.1078	(0.643)
1892			–0.0414	(0.599)	0.0121	(0.070)
1893	0.0712	(0.970)	0.0440	(0.637)	–0.0199	(0.113)
1894	0.0926	(1.170)	–0.0249	(0.342)	–0.2814	(1.487)
Observations	2045		15,880		15,880	

Notes: The table shows marginal effects from probit regressions, evaluated at the mean of each regressor, and t-statistics (in parentheses in absolute value, robust to heteroscedasticity). The second column shows coefficients from Table 3 in Rosenbloom (1998). The fourth column weights estimates by the number of establishments (truncated at 75).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

at the five-percent level but one are also significant in the replication and have the same sign. The lone exception is the coefficient on city with 250,000 residents or more, which is only significant in the original, though its sign is the same. Significant coefficients at the ten-percent level are less robust: out of five, one is not significant in the replication and one is significant with the opposite sign. The latter is the effect of unionization after 1887, which is unsurprising, since Rosenbloom (1998) uses the sample sample as Friedman (1988). Weighted estimates are less precise than their unweighted counterparts: 21 coefficients are significant in the unweighted replication, against 12 weighted ones. Moreover, they are less likely to agree with original in sign.

1.6.3 Card and Olson (1995)

Card and Olson (1995) fit an attrition model for the success rate and the change in wages after a successful strike. Unlike Friedman (1988) or Rosenbloom (1998), they do not draw a random sample from the reports. They use all rows for three states in the *Third Report* instead: Illinois, Massachusetts and New York. This design helps me separate the role of greater accuracy in the new file from the role of greater coverage.

Table 1.5 shows the original coefficients and mine. I replicate their binary model of the probability of success. Although it is not their main specification, it is simpler to interpret than their attrition model. Note that their estimates are not directly comparable to Friedman's, since he treats a compromise as a success, whereas Card and Olson (1995) treat it as a failure.

The first three columns use the restricted sample. This replication is closer to the original than any other by some margin. All coefficients are similar (including the impact of union involvement), whether they are significant or not. The only noticeable discrepancy is that the effect of the number of strikers becomes insignificant. Unlike Friedman (1988) or Rosenbloom (1998), even weighted estimates resemble their unweighted counterparts in magnitude. Their

Table 1.5: Determinants of the success rate (Card and Olson, 1995)

	Original sample			Full sample	
	Original coefficients	Unweighted replication	Weighted replication	Unweighted replication	Weighted replication
Union strike	0.49 (0.11)***	0.47 (0.11)***	0.44 (0.21)**	0.19 (0.02)***	0.22 (0.06)***
Fraction of employees on strike	0.34 (0.16)**	0.42 (0.16)***	0.63 (0.40)	0.52 (0.04)***	0.46 (0.11)***
Strikers (log)	−0.06 (0.03)**	−0.05 (0.03)	−0.02 (0.05)	−0.04 (0.01)***	−0.04 (0.02)**
Fraction of female employees	−0.81 (0.33)***	−0.78 (0.33)**	−0.78 (0.68)	−0.28 (0.08)***	−0.78 (0.23)***
Generic employees (indicator)	0.11 (0.11)	0.09 (0.11)	0.43 (0.26)	−0.04 (0.03)	0.09 (0.08)
Strike in Massachusetts	−0.07 (0.19)	−0.16 (0.15)	0.69 (0.32)**	−0.23 (0.04)***	−0.38 (0.13)***
Strike in Illinois	−0.91 (0.20)***	−0.98 (0.20)***	−1.01 (0.33)***	−0.43 (0.06)***	−0.45 (0.17)***
Strike in Chicago	0.87 (0.22)***	0.93 (0.22)***	1.55 (0.35)***	0.18 (0.07)***	0.17 (0.20)
Start date: May 1–7, 1886	−0.44 (0.19)**	−0.46 (0.19)**	0.01 (0.41)	−0.77 (0.08)***	−0.31 (0.23)
Start date: after May 1, 1886	−0.21 (0.16)	−0.14 (0.17)	0.16 (0.39)	−0.30 (0.06)***	0.07 (0.20)
Observations	1026	1032	1032	15,874	15,874
Parameters	27	27	27	78	78

Notes: The table shows probit estimates and standard errors (in parentheses). The second column shows coefficients from Table 4 in Card and Olson (1995). The fourth and sixth columns weight estimates by the number of establishments (truncated at 75). All regressions include industry, state and year effects. The original sample comprises strikes for an increase in wages, from 1881 to 1886, in Illinois, Massachusetts and New York.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

standard errors are larger again though.

The full sample paints a different picture. Point estimates change considerably, perhaps because of a lack of external validity – Illinois, Massachusetts and New York were indeed atypical in their levels of industrialization and labor activism. Nonetheless, all significant coefficients in the original have the same sign in the replication. Moreover, the weighted estimates are still close to the unweighted ones.

1.6.4 Currie and Ferrie (2000)

Currie and Ferrie (2000) investigate the implications of labor law for industrial conflict. They consider five legislative interventions: the legalization of unions; limits on the number of working hours per week; bans on intimidation and boycotts; bans on blacklisting; and the use of injunctions against strikers.

The paper examines a range of outcomes. I selected four for replication: union involvement in strikes, their duration, the subsequent change in hours and the use of strikebreakers. Because I do not have legislative data from other states, I restrict the sample to the thirteen states in original analysis. Table 1.6 shows the original coefficients and mine.

The first panel displays results for union involvement. Currie and Ferrie (2000) find a marginally significant effect from the legalization of unions, which disappears in the replication. On the other hand, the highly significant impact of injunctions is robust: it is –0.073 in the original, –0.090 in the unweighted and –0.063 in the weighted replication.

The second panel considers duration. Results are mixed. The effect of hour limits survives with similar magnitude and significance level. By contrast, the coefficients of blacklist bans and injunctions are not significant in the replication. Unlike the original, I estimate a significant impact from outlawing intimidation and boycotts.

The third panel examines the change in hours after a strike. Currie and Ferrie (2000) find little effect from the five interventions under consideration. I estimate that they had a significant joint impact. Most coefficients are larger in magnitude and more precise. I compute individually significant effects from blacklist bans, injunctions (without weights) and outlawing intimidation and boycotts (with weights).

The fourth panel analyzes the incidence of strikebreaking. Three interventions exhibit significant coefficients in the original: union legalization, hour limits and blacklist bans. The first two are similar in the unweighted replication. The impact of blacklist banks becomes essentially zero. It is interesting that weighting does not yield any significant treatment effect. Given my results for Friedman (1988), Rosenbloom (1998) and Card and Olson (1995), it seems that the regressions of the use of strikebreakers is particularly sensitive to weighting.

TABLE 1.6: THE LAW AND LABOR STRIFE

	Original coefficients		Unweighted replication		Weighted replication	
Outcome: union involvement						
Unions legal	0.045	(1.78)*	−0.013	(0.387)	0.002	(0.042)
Maximum hours legislation	0.011	(0.457)	−0.082	(3.352)***	−0.046	(1.314)
Intimidation or boycotts illegal	0.019	(0.505)	0.043	(1.180)	−0.018	(0.340)
Blacklists illegal	0.008	(0.019)	−0.068	(2.876)***	−0.085	(2.916)***
Injunction used	−0.073	(2.81)***	−0.090	(3.113)***	−0.063	(1.850)*
F-*test for five laws*	2.07	[0.066]*	5.530	[0.000]***	2.935	[0.012]**
R^2	0.229		0.229		0.240	
Outcome: log duration						
Unions legal	0.076	(1.02)	−0.003	(0.037)	0.089	(0.563)
Maximum hours legislation	−0.451	(6.12)***	−0.289	(4.304)***	−0.580	(3.740)***
Intimidation or boycotts illegal	0.139	(1.23)	0.275	(2.543)**	0.461	(1.846)*
Blacklists illegal	0.138	(2.72)***	−0.030	(0.461)	0.150	(0.882)
Injunction used	0.211	(2.72)***	−0.137	(1.698)*	0.035	(0.200)
F-*test for five laws*	9.77	(0.00)***	5.521	[0.000]***	3.566	[0.003]***
R^2	0.144		0.167		0.229	
Outcome: change in hours (%)						
Unions legal	−0.146	(0.416)	1.276	(1.190)	−0.663	(0.830)
Maximum hours legislation	−0.178	(0.178)	−0.253	(0.326)	0.504	(0.668)
Intimidation or boycotts illegal	−0.345	(0.345)	−1.957	(1.092)	−2.696	(2.044)**
Blacklists illegal	−0.145	(0.145)	−1.393	(2.306)**	−2.372	(2.794)***
Injunction used	0.202	(0.202)	−2.420	(3.924)***	−0.822	(0.662)
F-*test for five laws*	0.504	[0.774]	4.425	[0.000]***	2.480	[0.030]**
R^2	0.037		0.095		0.095	
Outcome: use of strikebreakers						
Unions legal	0.056	(1.96)**	0.096	(2.743)***	0.095	(1.199)
Maximum hours legislation	−0.083	(2.91)***	−0.059	(2.228)**	−0.118	(1.928)*
Intimidation or boycotts illegal	0.006	(0.147)	0.071	(1.676)*	−0.049	(0.461)
Blacklists illegal	0.062	(2.92)***	0.009	(0.357)	0.035	(0.516)
Injunction used	0.022	(0.718)	−0.036	(1.178)	0.066	(0.785)
F-*test for five laws*	4.05	[0.001]***	3.369	[0.005]***	1.351	[0.240]
R^2	0.057		0.063		0.113	
Observations	12,695		12,532		12,532	

Notes: The table shows linear estimates, t-statistics (in parentheses in absolute value, robust to heteroscedasticity) and p-values (in brackets, robust to heteroscedasticity). The second column shows coefficients from Tables 6 and 7 in Currie and Ferrie (2000). The fourth column weights estimates by the number of establishments (truncated at 75). All regressions control for: prestrike employment; prestrike hours; the fraction of women in the workforce; the number of strikes by industry, state and year; city effects; industry effects; state effects; and time trends by state.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

2 Organizing Collective Action: Labor Strife in the U.S. in the 1880s

Look, my comrades, see the union banners waving high:
Reinforcements now appearing, victory is nigh.

— Extract of "Hold the Fort"[1]

2.1 Introduction

Industrial action underlies the bargaining power of labor unions. Without the credible threat of a work stoppage, employers have no reason to recognize unions or offer them concessions. Yet unorganized workers are also capable of collective action: for example, they undertook 38 percent of strikes in the United States in 1900, 42 percent in Austria-Hungary and 61 percent in Germany (U.S. Bureau of Labor, 1906).[2] Why do workers unionize then? How does organization improve on wildcat picketing?

This paper explores one explanation: organization helps workers win strikes. I hypothesize that unions facilitate coordination, helping members deploy a wider tactical inventory. First,

I am indebted to Lori Beaman, Nicola Bianchi, Joseph Ferrie, Carola Frydman, Robert Margo, Joel Mokyr, Matthew Notowidigdo, Nancy Qian, Daniel Rees and Max Tabord-Meehan for advice. Seminar participants at École Polytechnique, EPGE, Insper and Northwestern University and participants in the 2018 Annual Meeting of the Economic History Association offered valuable feedback as well. I am grateful to Joseph Ferrie for sharing his data. All mistakes are mine.

[1] Published in *Labor Songs Dedicated to the Knights of Labor* (Chicago, IL: J. D. Tallmadge, 1886).

[2] The strike of the freight handlers in Chicago in April 1881 is illustrative of wildcat walkouts. Under an informal leadership, workers discussed their plans at lunch and after work. They circulated a petition for a wage raise for several days, which they presented to the railroad companies. After the employers denied their request, workers struck the Illinois Central Railroad. Turnout was mixed elsewhere. Most companies promised to match concessions by the Illinois Central Railroad and other lines if employees did not quit work. Strikebreakers were hired, but they were inexperienced and suffered intimidation. Although newspapers repeatedly announced the imminent defeat of the strikers, the railroads offered an unconditional concession after five days of negotiations. (This account is based on daily reports in the *Chicago Tribune* and the *Inter Ocean*.)

they buttress the picket line by raising defense funds, fostering solidary, etc. Second, they weaken employers by calling boycotts, increasing turnout, etc. Third, they negotiate better settlements on the strength of their bargaining experience and reputation. Fourth, they facilitate the exchange of information and decision making through conventions, journals, etc.

It is difficult to test this hypothesis in modern labor markets. Figure 2.1 shows organized strikes as a percentage of strike activity in the U.S. from 1881 to 1957. As the labor movement matured, wildcat stoppages dwindled: unions participated in 92 percent of walkouts by the time of the National Labor Relations Act of 1935, accounting for 98 percent of idled man-days.[3] As a consequence, there is not enough variation in postwar data to identify the impact of unionization on industrial conflict.

Historical data overcome this deficiency. I borrow rich microdata from the earliest nationwide sample of work stoppages in the U.S., the *Third Annual Report of the Commissioner of Labor* (U.S. Bureau of Labor, 1888). My analysis encompasses 2172 unorganized and 3191 organized strikes, ranging from 1881 to 1886. Like Card and Olson (1995) and Friedman (1988), I focus on the impact of organization on the success rate – i.e. the probability that strikers extract concessions from management.

Endogeneity poses a second challenge. The probability of success of a strike comprises a baseline rate and an organization effect (if applicable). Striking is a strategic decision (Hayes, 1984): workers walk out when victory seems likely enough. For instance, they might feel that an employer is vulnerable on account of perishable inventory or outstanding orders. To the extent that the organization effect offsets worse baseline odds, unionized workers might undertake riskier strikes than the unorganized and confront stronger employers. In other words, bargaining

[3] Wildcat strikes decreased first in relative and later in absolute terms (Peterson, 1938). There were fluctuations. For example, many unions disbanded in recessions, increasing the proportion of unorganized strikes. For context, Friedman (1999) puts the unionization rate among American industrial workers at 3.75 percent in 1880, 9.68 percent in 1890, 6.35 percent in 1899 and 16 percent in 1909. Comparable series for other countries are scarce. According to second-hand data from the U.S. Bureau of Labor (1906), the percentage of organized strikes increased from 27 in 1894 to 55 in 1905 in Austria-Hungary and from 58 in 1899 to 75 in 1905 in Germany.

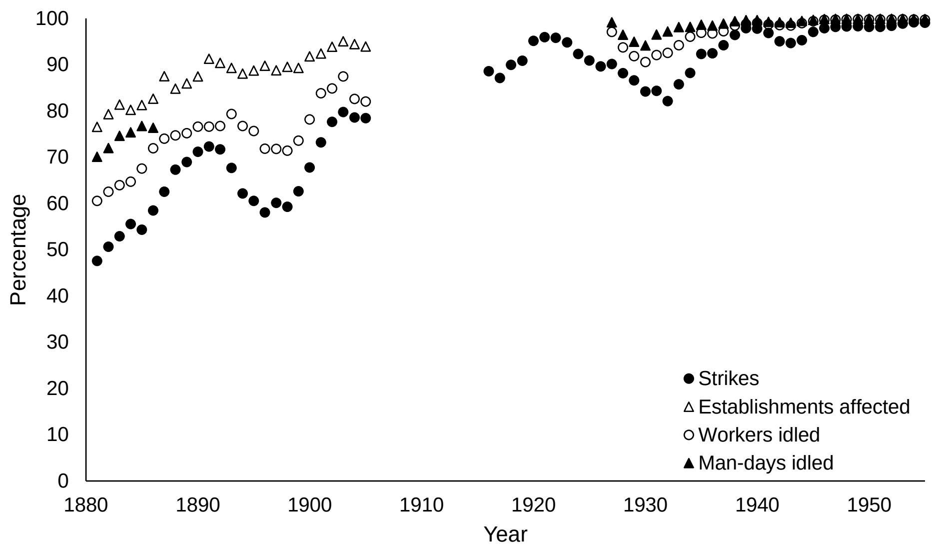

Figure 2.1: Organized strikes as a percentage of strike activity in the U.S. by measure

Notes: The figure shows three-year moving averages. For years 1881 to 1905, it shows strikes ordered by labor organizations as a percentage of all strikes. For years 1916 to 1957, it shows stoppages involving unionized workers as a percentage of strikes and lockouts. No data were collected between 1906 and 1915.
Sources: 1881 to 1905: U.S. Bureau of Labor (1888, 1906). 1916 to 1957: various editions of the *Monthly Labor Review*, by the U.S. Bureau of Labor Statistics.

strategies are endogenous. Therefore, organized strikes might exhibit a lower baseline success rate on average than the wildcat, which would downward bias estimates of the organization effect.

I exploit an instrumental variable for identification: the existence of a local assembly of the Knights of Labor (KOL) in the locality of the dispute.[4] The KOL constituted the foremost labor society of the 1880s (Friedman, 1988; Voss, 1993), peaking at a fifth of the industrial workforce in 1886. Their presence indicates that the local workforce had unionized and that union officers operated in the community. Hence, it should correlate with union involvement in strikes. I justify the exclusion restriction on four accounts: (1) assembly creation depended on recruitment opportunities and the proximity to existing offices rather than strike prospects; (2) the instrument incorporates a lag between observations and changes in the location of the

[4] See Enflo and Karlsson (2018) for a related identification strategy.

Knights, so it is orthogonal to the dynamics of conflict outcomes; (3) my results are robust to balance adjustments, sample restrictions and variations in the instrument; and (4) the instrument passes the test of the exogeneity condition of Machado, Shaikh and Vytlacil (2018).[5]

I estimate that union sponsorship increased the probability of success by 31 percentage points from a baseline rate of 40 percent, rationalizing the preponderance of organized strikes in modern industrial relations. Ordinary regression yields a lower estimate, 12 percentage points, which suggests that unionized workers responded to the higher success rate by undertaking riskier walkouts. The difference is statistically significant. Furthermore, unions reduced the probability of job loss by 22 percentage points from a baseline rate of 56 percent. On the other hand, there is no evidence that successful organized workers achieved larger wage gains or hour cuts than successful wildcat strikers. I find no change in the duration of standoffs either.

This paper pertains to three literatures. First, it furthers the scholarship on labor unions. Economists have long debated their impact on the wage structure and firm performance, from Freeman and Medoff (1984) to recent work by Card, Lemieux and Riddel (2004), Collins and Niemesh (2018), DiNardo and Lee (2004), Farber et al. (2018) and Lee and Mas (2012). I investigate an explanation for the differential bargaining power of organized labor, which underlies their estimates. Second, I add to the literature about strikes. Gramm and Schnell (1994) show that union officers affect the chances of strikebreaking, which supports my hypothesis. Card and Olson (1995) and Geraghty and Wiseman (2008) estimate attrition models from a subset of my sample. Their framework is useful in interpreting my parameters. Other empirical research has mostly focused on the interplay between work stoppages, wage outcomes and market conditions.[6] Third, this paper contributes to the historiography of the American labor movement in the 1880s. This decade saw an unprecedented experiment in radical mass unionism

[5] This test is based on the intuition that the correlation between the instrument and outcomes should be neither too small nor too large if the instrument only affects outcomes through the treatment. See Subsection 2.6.2 for further discussion.

[6] For surveys, see Card (1990), Cramton and Tracy (2003), Kennan (1986) and Kennan and Wilson (1989).

by the KOL, culminating in the strike wave of May 1886. Their subsequent collapse entrenched conservative craft unionism in the U.S. Historians have partly blamed this reversal on a string of defeats of the KOL (Friedman, 1988; Kaufman, 2001; Perlman, 1918; Voss, 1993). This paper nuances this view: I argue that the Knights were successful strike leaders, but they could not impose discipline on the rank and file. Kremer and Olken (2009) propose a similar explanation in the context of an evolutionary model of unionization. I complement recent research on environmental constraints on labor activism, such as the government (Currie and Ferrie, 2000; Friedman, 1988), employers (Schmick, 2018; Voss, 1993), market integration (Ansell and Joseph, 1998) and rival associations (Kaufman, 2001).

The paper continues as follows. Section 2.2 summarizes the historical background. Section 2.3 describes the data. Section 2.4 discusses identification. Section 2.5 explains the econometric strategy. Sections 2.6 and 2.7 present the results. Section 2.9 concludes.

2.2 Historical background

2.2.1 Origins of the American labor movement

Labor unrest was only sporadic in the United States in the early decades after independence (Saposs, 1918).[7] The first known wildcat strike implicated journeymen printers in New York in 1776. Cordwainers pioneered the organized strike in Philadelphia in the 1790s. Printers and shoemakers went on to establish associations across the northeast, but other trades did not organize until the late 1810s. Evidence of incipient working-class consciousness dates to 1827 (Saposs, 1918; Sumner, 1918), when trade societies agitated for the ten-hour day in Philadelphia. The campaign led to the creation of a citywide federation of labor unions and a labor party. Solidarity crossed occupational lines as workers learned to articulate their common grievances.

[7] There were work stoppages in colonial times. For example, bakers struck in New York in 1741. However, Saposs (1918) argues that these disputes pitted master artisans against local authorities, rather than employees against employers.

Early unions restricted membership to skilled craftsmen. As workshops gave way to factories, artisans blamed mechanization and the division of labor for a perceived loss in autonomy, competency and status (Katz and Margo, 2014; Voss, 1993). Organization was their response. They yearned for a republic of independent producers (Hallgrímsdóttir and Benoit, 2007; Voss, 1993): claiming moral superiority over the “subordinate laborer” as well as the “idle classes”, craftsmen maintained that “wage slavery” was incompatible with a free citizenry. Unions translated ideology into collective action: party politics, collective bargaining, worker cooperatives, mutual insurance, industrial education, etc.

The labor movement foundered as a recession took hold in the 1830s (Mittelman, 1918). It would slowly rebuild and evolve. The earliest national trade associations emerged in the 1850s, as rail links and the telegraph stimulated market integration (Andrews, 1918; Ansell and Joseph, 1998). They grew in importance after the Civil War, overshadowing local unions. Industrialization had so realigned interests by the late 1860s that labor leaders took the first steps to organize the “subordinate laborer” (i.e. blacks, women and the unskilled). In one such effort, Uriah S. Stephens founded the Noble and Holy Order of the Knights of Labor in Philadelphia in 1869.

2.2.2 The Knights of Labor

Stephens blamed the degradation of labor on internal divisions (Grob, 1958; Wright, 1887). Capital tended toward concentration, through which it accumulated bargaining power and political influence, whereas workers fragmented into uncoordinated trade unions. Stephens saw strength in numbers: if wage earners coalesced into a single organization, they would have the clout to defend their common interests against combined capital. He envisioned a more integrated labor movement as well as a wider constituency.

The KOL combined traits of labor unions, fraternal brotherhoods and political parties (Birdsall, 1953; Kaufman, 2001). They advocated incremental progress toward the abolition of the wage system (Grob, 1958; Wright, 1887). Their agenda included such intermediate goals as the eight-hour week, equal pay for equal work, a ban on child labor, graduated income taxation, antitrust law, public ownership of utilities and the creation of labor bureaus. As Stephens proposed, the order sought to mobilize a critical mass toward socioeconomic reform. In consequence, it adopted a distinctively inclusive admission policy. It recruited unskilled laborers as well as craftsmen, irrespective of nationality, religion or occupation. It would extend affiliation to blacks in 1878 and women in 1882.[8]

The KOL assumed a dual role (Birdsall, 1953; Wright, 1887). As educators, they commended gradualism and nurtured solidarity. As coordinators, they encouraged collective action in three forms. First, political activism would win reform at the ballot box.[9] Second, worker cooperatives would offer an alternative to wage employment. Third, organization would help workers bargain for better work conditions. On the other hand, they were ambivalent about industrial conflict. National officers warned that strikes should be a last resort, recommending arbitration instead, but local cadres had much freedom to interpret those guidelines and the ranks were keen to strike. (For further discussion, see Section 2.7.)

At their first constitutional convention in 1878, the Knights structured their government in three tiers on a territorial basis (Birdsall, 1953). The local assembly (LA) was the basic unit of organization. Its size ranged from a minimum of ten members to the thousands (Garlock, 1982). Each defined its own jurisdiction, which could span from a single establishment to a large city. Some restricted admission further: for instance, LA 5327 recruited wood workers in East Boston and LA 8072 affiliated Germans in Holyoke. The second level was the district assembly, though local assemblies did not necessarily belong to one. Supreme authority lied with the

[8] There were limits to its inclusiveness: e.g., the Knights rejected politicians, liquor distributors, lawyers and the Chinese.

[9] The Knights did not support a specific party. They occasionally endorsed candidates and members could run for office, though they did not affiliate professional politicians.

General Assembly.[10] However, the national executive board wielded little power over lower assemblies in practice. They played two main roles: settling disputes between districts and discouraging strikes from the bully pulpit. As a federation of autonomous assemblies of varying scopes, the order sacrificed a coherent national strategy for the flexibility to accommodate its diverse membership (Birdsall, 1953; Grob, 1958; Voss, 1993).

The KOL faced headwinds in their early years (Kaufman, 2001; Wright, 1887), including a downturn in the 1870s and opposition from the Catholic Church. These challenges prompted change. Stephens ceded the headship to Terence V. Powderly in 1879, clearing the way for the elimination of secret oaths in stages by 1882.[11] In an additional effort to boost recruitment, the General Assembly created the organizer in 1878. Organizers received a paid commission to found new assemblies. Because existing locals could be jealous of their constituencies and districts had a right of oversight within their jurisdiction (Birdsall, 1953; Voss, 1993), organizers had an incentive to operate in unclaimed territory, which stimulated the geographic expansion of the KOL.

The order was the third national labor federation in the United States.[12] While it had much in common with its predecessors (e.g., the emphasis on political action and independent producers), it innovated in embracing all wage earners. Europe underwent a similar transition from craft movements to radical mass unionism at the end of the 19th century (Friedman, 1988; Voss, 1993). Like the KOL, a moderate strand espoused gradualistic politics, as the British Fabians exemplify. Others preached revolution, such as the French Confédération Générale du Travail.

---

[10] This hierarchy became more complex with the advent of state assemblies in 1883 and national trade assemblies in 1884. See Birdsall (1953) and Grob (1958).

[11] The KOL kept strict secrecy until 1879. They were partly motivated by the fear of retaliation and partly influenced by a fascination with the freemasonry. However, the Catholic Church was no friend of secret societies, which proved problematic. See Wright (1887) and Kaufman (2001).

[12] Its predecessors were the National Trades' Union (1834–37) and the National Labor Union (1866–73). The International Workingmen's Association (1864–72) maintained local branches in the United States too.

2.2.3 *Labor at a crossroads*

The Panic of 1873 triggered a severe recession. Labor activism withered amid high unemployment and pay cuts, but deteriorating work conditions sowed the seed of recovery (Voss, 1993). Workers' discontent found dramatic expression in the Great Railroad Strike of 1877 (Lloyd, 2009), when wildcat protests spread nationwide after the B&O Railroad reduced wages for a third time in six months. The working class demonstrated mobilization potential and latent solidarity, which unions set out to cultivate.

Friedman (1999) estimates that total union membership rose from 168,000 in 1880 to 500,000 in 1885. The unionization rate reached 4.6 percent.[13] The Knights grew from nine thousand in 1878 to a hundred thousand in 1885 (Perlman, 1918). This expansion owed much to the subsumption of independent unions and the reorganization of extinct ones. It involved considerable turnover: for instance, 18,104 workers joined the KOL in 1880, but 10,056 quit. Labor strife affected an yearly average of 2630 establishments and 176,513 workers between 1881 and 1885 (U.S. Bureau of Labor, 1888).

In October 1884, the Federation of Organized Trades and Labor Unions (FOTLU) announced that the eight-hour day should become standard by May 1, 1886. The plan energized the labor movement (Kemmerer and Wickersham, 1950; Perlman, 1918). Further momentum built after impressionistic press reports about the KOL and successful stoppages of the railroads. As a strike wave loomed, union membership attained 1.2 million in the spring of 1886 (Friedman, 1999), including one in five industrial workers. The Knights increased sevenfold between new recruits and readmissions (Perlman, 1918), surpassing 700,000 members and becoming the largest labor organization in the United States. More than 300,000 protesters marched on May Day. Pickets continued into the following weeks, but the campaign lost impetus after the Haymarket Affair (the fatal bombing of Chicago police on May 4). Notwithstanding concessions from individual

[13] Friedman (1999) speculates that these figures understate union membership and overstate growth.

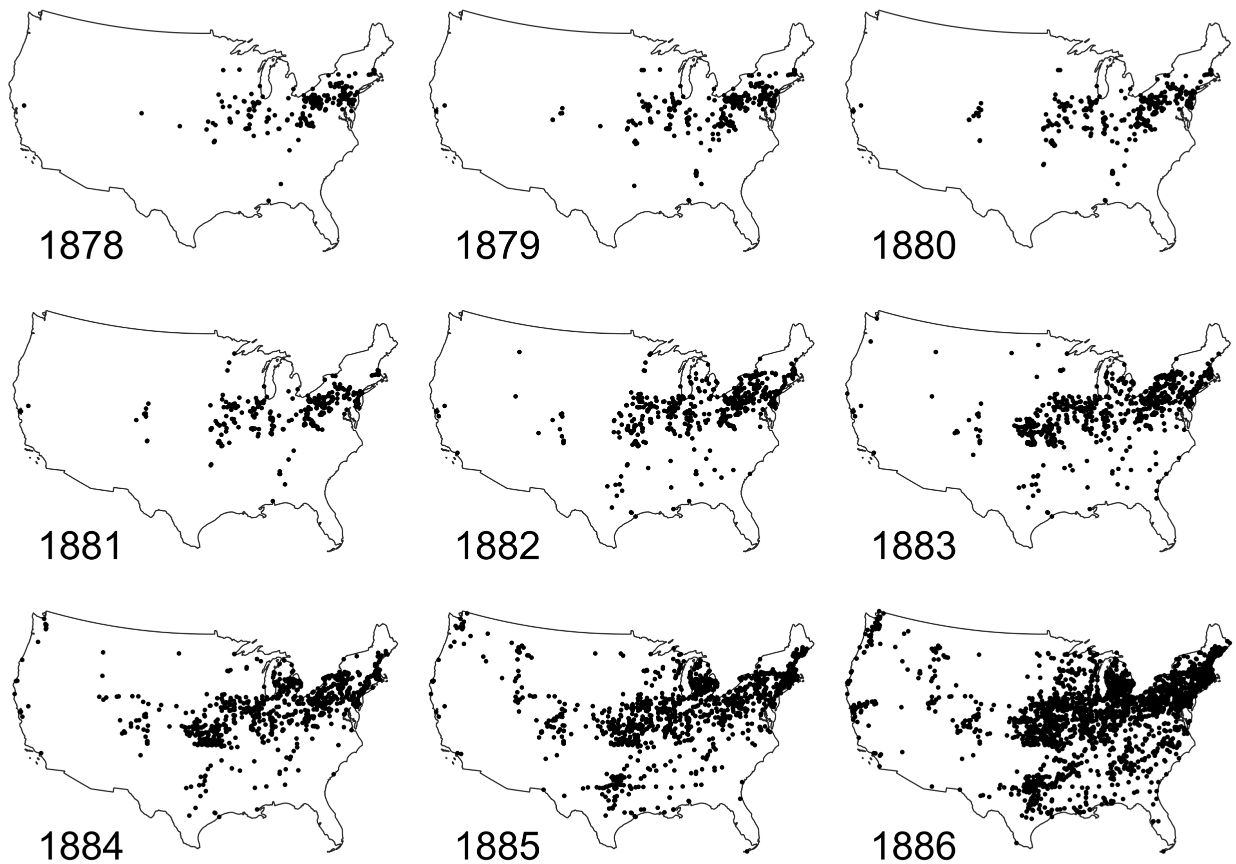

Figure 2.2: Local assemblies of the Knights of Labor by year (select years)

Source: Garlock (1982, 2009).

employers, activists yielded without achieving the statutory eight-hour day as the violence caused a backlash in public opinion against the strikers.[14]

May 1886 transformed organized labor. The Knights entered rapid decline (Oestreicher, 1984), dwindling to twenty thousand by 1900. In December 1886, Samuel Gompers forged the American Federation of Labor (AFL) from the FOTLU. The AFL was a league of conservative craft unions (Friedman, 1988; Grob, 1960), which eschewed social reform to focus on workplace issues. It would soon dominate the American labor movement, whereas the unskilled remained mostly unorganized until the 1930s. By contrast, radical industrial unions recovered from early setbacks in Europe. This divergence is a topic of ongoing debate. Recent research suggests that environmental constrains limited the effectiveness of mass unionism in the U.S. (Ansell and Joseph, 1998; Friedman, 1988; Kaufman, 2001; Voss, 1993). This paper nuances this view: the Knights and other unions helped workers win strikes in fact, yet they failed to impose discipline on the rank and file. As Kremer and Olken (2009) argue, weak leadership may have imperiled their survival on the long run. For further discussion, see Sections 2.7 and 2.9.

2.3 Data

This section describes the data sources. For further detail about the construction of the sample, see Section 2.A. Table 2.1 shows summary statistics.

2.3.1 Strikes and lockouts

I use strike data from the *Third Annual Report of the Commissioner of Labor* (U.S. Bureau of Labor, 1888).[15] The report was the second nationwide survey of work stoppages in the United

[14] See Biggs (2002) for a study of multistage bargaining in the context of May 1886.

[15] This project exploits the full sample. The data incorporate the subset of Card and Olson (1995), Currie and Ferrie (2000) and Geraghty and Wiseman (2008). Friedman (1988) and Rosenbloom (1998) used a different subsample.

States,[16] following the postal inquiry for the *Tenth Census* by Weeks (1886), and covers the period from 1881 to 1886.[17] Agents collected data in two stages. They first compiled a list of strikes and lockouts from newspapers, trade journals and other sources. They then visited each locality on the list to interview managers, workers and union officers. These inquiries were not limited to the episodes in the initial list. The only exclusion criterion was a minimum duration of one day. This investigation yielded detailed information about each dispute, including: localities, industry, dates, causes and outcome; involvement of unions or employers' associations; affected workers, their occupation and average wages; affected establishments, their size, average wages and weekly hours; and establishment closures, idled hands and financial losses.[18] The report gives employment, wages and hours before and after the conflict as well as by gender.

The final sample consists of 5363 observations. I exclude lockouts (358 rows), incomplete strikes (4 rows), general strikes (48 rows) and strikes in imprecise localities (36 rows). To account for price differences across time and regions, I adjust wages by the monthly price index of Warren and Pearson (1932) and the state price index of Haines (1989). I impute hours for 23 observations with irregular workweeks. I chose 60 hours, the mode in both the *Third Report* and the Census of Manufactures (Atack and Bateman, 1992).

My primary outcome is an indicator of success. In accordance with contemporaneous practice (Card and Olson, 1995), the report classifies walkouts as successes (43 percent), compromises (9 percent) or failures (49 percent). Since workers' initial demand is a strategic variable, I combine compromises and successes, following Friedman (1988). Hence, success consists in extracting concessions from employers for my purposes. For future reference, successful strikes

[16] Statewide surveys took place in Ohio (1878), Massachusetts (1879) and Pennsylvania (1882). Massachusetts' commissioner was Carroll D. Wright at the time, who became the federal commissioner in 1885.

[17] The *Tenth Report* (U.S. Bureau of Labor, 1896) covers the period between 1886 and 1894. I do not use this sample because the unit of observation changes and no wage information is available. Later reports contain only aggregate data, as does Weeks (1886). Their original schedules could not be located.

[18] The report gives the number of new employees at affected establishments after the conflict. This series does not include temporary replacement workers. Currie and Ferrie (2000) and Rosenbloom (1998) use it as a proxy for the number of strikebreakers, though it is a lower bound.

constitute 51.5 percent of the sample. My treatment is union sponsorship. In the words of the *Third Report*, treated units were "ordered by a labor organization". It is unclear whether the Bureau abided by this restrictive label. It may also have flagged unauthorized stoppages by unionized workers, especially when the union aided its members in some form.[19] For future reference, organized strikes constitute 59.5 percent of the sample.

The *Third Report* has two shortcomings. First, the unit of observation is inconsistent. Because general strikes can often be difficult to delimit (Peterson, 1938), the Bureau of Labor planned separate entries for each affected establishment. This design proved overly ambitious. Therefore, the Bureau aggregated related stoppages instead, provided that they coincided in industry, dates and resolution. A row may thus represent either a strike or part of one. Following the literature,[20] I treat each line as an observation, which has the advantage of underweighting outliers, but my results are robust to the choice of weighting scheme.[21] I account for the correlation between rows in inference (see Subsection 2.5.2). Second, Bailey (1991) shows that the Bureau undercounted stoppages: local newspapers record twice as many disputes in Terre Haute. He could not identify systematic differences between included and excluded stoppages.[22] As Card and Olson (1995) argue, missing data do not jeopardize statistical analysis so long as they are random.

2.3.2 Assemblies of the Knights of Labor

Garlock (1982, 2009) gathered information on the local assemblies of the Knights of Labor from primary sources. Two official publications cover my sample period: the *Journal of United Labor*

[19] The line between authorized and unauthorized strikes had become so unclear by the 1920s that the Bureau preferred to record the mere involvement of unionized workers instead (Peterson, 1938). This distinction might have been ambiguous in the 1880s as well.

[20] Contemporaneous authors weighted rows by establishments or employees (U.S. Bureau of Labor, 1888).

[21] Table 2.11 in Section 2.C explores the sensitivity of my main estimate to weighting by establishments.

[22] I attempted to replicate the exercise for three cities: Chicago, Decatur and Milwaukee. I focused on the first six months of 1881. I believe that omissions are most likely in this early period, since the data collection took place between 1886 and 1887. I found six additional strikes in Chicago (against 46 in the report), none in Decatur (vs. 3) and three in Milwaukee (vs. 3). Hence, it appears that 50 percent is an upper bound on the omission rate.

(from 1880 to 1885) and the proceedings of the General Assembly (from 1879 to 1885). Therefore, his list should be exhaustive or nearly so. The data include location, years of operation and fragmentary membership statistics.[23]

My instrument is the existence of an assembly in the locality of the strike in the year before the strike. I use a lag for two reasons. First, it ensures that locals were not chartered post factum, given that I do not observe exact organization dates. Second, it allays concerns about endogenous entry decisions in anticipation of a standoff. Section 2.4 discusses identification in greater detail. To test the exogeneity condition, Section 2.8 considers two alternative instruments: the existence of an assembly in the locality of the strike in 1880 and the log distance to the nearest assembly outside the locality.

It is sometimes unclear whether all sources distinguished between two localities. For example, the *Third Report* contains both Knoxville and Knoxville Junction (IA), whereas Garlock (1982, 2009) could only find mention of Knoxville in the files of the KOL. I treat each pair in question as one place. When an assembly or a strike spanned multiple localities (23 and 103 observations, respectively), I base the instrument on the first entry. If the locality of the strike is a county (20 observations), I consider whether a local existed anywhere within the county. My estimates are robust to these choices.

2.3.3 Market conditions

To account for heterogeneity across local labor markets, I draw aggregate statistics from the *Tenth Decennial Census* of 1880 and other sources. These variables are all measured at the county level.

I construct demographic statistics from the full-count microdata of the population census

[23] Locals ought to submit quarterly membership figures, which appear in abridged form in the annual proceedings of the General Assembly, but this duty was often neglected. According to Kaufman (2001), water leaks destroyed the original forms.

(Ruggles et al., 2018). I restrict the sample to industrial workers. This subset is more representative of potential strikers, since labor activism was marginal in agriculture, trade and services.[24] My control set includes: industrial workers as a percentage of the total labor force; urban workers as a percentage of the industrial workforce; and gender, ethnic and trade fragmentation indexes. I construct average firm sizes and average daily wages from the tables of the census of manufactures (Haines and ICPRS, 2010).[25] I adjust wages by the state price index of Haines (1989). These covariates help me address the correlation between unionization, industrialization and urbanization. Note that they do not vary over time. Moreover, I borrow data from Atack (2016) to compute the ratio of railways to land area. Transportation links facilitated access to replacement workers, which could influence labor strife. Finally, I construct an indicator of past labor strife from the *Third Report*, viz. the occurrence of a successful stoppage in the previous year in the pertinent sector and county. This variable is not available for 1881, as the report does not include data for 1880. My results are robust to alternative specifications, e.g., considering all strikes rather than successful strikes or using data for 1880 at the state level from Weeks (1886). This indicator proxies for unobserved determinants of strike incidence and conflict outcomes.

2.4 Identification

To estimate the causal impact of union intervention on conflict outcomes, one must account for endogeneity in bargaining strategies. This section presents my identification strategy. To simplify the exposition, I focus on the success rate and ignore the nature of the payoff in dispute.

Striking is a strategic decision (Hayes, 1984; Kennan, 1986): workers walk out if victory seems

[24] For my purposes, industry comprises mining, construction, manufacturing, transportation and utilities. Agriculture, trade and services represent 46 observations in my sample but over two thirds of the labor force. I use the industrial classification of the 1950 Census, imputed by IPUMS (Ruggles et al., 2018).

[25] I do not adjust estimates for differences in capital stock, child labor or unemployment incidence. Although my results are robust to such adjustments, these variables are poorly measured, so I omit them. Additional estimates are available upon request.

likely enough. The probability of success is the sum of a baseline rate and an organization effect. The baseline rate reflects the circumstances of each confrontation. For example, Newark leather workers faced worse odds after their employers committed in 1886 to pay a fine to the industry association if they should offer concessions in any future standoff (Voss, 1993). The organization effect applies if a union defends the strikers. Because it raises the probability of victory (hence, expected payoffs), it might influence the likelihood of a breakdown in negotiations: union members might strike when the unorganized would rather accommodate. In other words, bargaining strategies are endogenous. Therefore, unions have a direct effect on the probability of success (given a baseline rate, they increase the total rate) and an indirect effect (they enable strikes with lower baseline rates). This indirect impact lowers the average success rate of organized walkouts, which would downward bias estimates of the organization effect.[26] It is important to distinguish these two channels: if I estimate the organization effect to be zero, is it because unions did not help workers win or because their indirect effect offset the direct one?

I use an instrument to establish causality: the existence of an assembly of the KOL in the locality of the dispute in the preceding year. It should correlate with union intervention for two related reasons: it shows that the local workforce had unionized to some extent and it indicates that workers had easy access to union officers. As Imbens and Angrist (1994) note, it identifies effects on compliers, for which union support depended on the presence of the KOL in the community. The identification is mostly due to geographic variation: I compare the success rate in localities with an assembly to the rate in localities without them. A longitudinal analysis is inviable because of the incidence of labor strife is so low in most towns that I do not observe stoppages before the entry of the Knights and afterwards.

A valid instrument must satisfy the exclusion restriction: it must only correlate with outcomes

[26] Card and Olson (1995) raise a different concern. Unions were not keen on industrial action, since a defeat could trigger defections and imperil the organization. (See also Friedman (1988) and Kaufman (2001).) If unions avoided endorsing riskier strikes, naive estimates would be upward biased. There is plausible anecdotal evidence in support of this hypothesis, but my results suggest otherwise.

through the treatment. My instrument incorporates a lag between observations and the creation of new assemblies. This construction ensures that it is orthogonal to the baseline success rate so far as its determinants change over time. For instance, the Knights could not anticipate in 1885 that Newark leather manufacturers would later associate. My results are robust to the choice of lag (see Section 2.8). Two threats to the exogeneity condition remain. First, the presence of a local assembly might correlate with static covariates. The KOL had an incentive to win strikes, which could help them attract and retain members. Therefore, they might have targeted communities where strikers were most likely to succeed, causing upward bias. There is no anecdotal evidence though that they took strike outcomes into consideration as they expanded. Recall from Subsection 2.2.2 that the Knights pursued membership growth in the hope of advancing labor legislation via the ballot box. They did not direct organizers' efforts, which were mainly constrained by recruitment opportunities and the location of existing assemblies. Still, the presence of an assembly might unintentionally correlate with determinants of conflict outcomes. For instance, locals were more common in urban areas (due to the abundance of manufacturing workers). I allay this concern by reweighting the sample for imbalances in strike characteristics and market conditions (see Section 2.5). These corrections do not affect my results. Second, union members might not be representative of strikers at large. Selection should work against me though in that a weak bargaining position gives workers an incentive to unionize before striking. The Knights were especially susceptible to negative selection, as they did not restrict admissions by skill.

Machado, Shaikh and Vytlacil (2018) propose a statistical test of the exclusion restriction. As Subsection 2.6.2 shows, I reject the null hypothesis of an invalid instrument at any conventional significance level.

In addition to the exclusion restriction, Abadie (2003) establishes three identification conditions: the existence of compliers, monotonicity and common support. The first condition means that the instrument should correlate with the treatment, which I can easily assess through the first-stage F-statistic. The monotonicity assumption states that the existence of an assembly must not reduce the probability of union intervention, which is plausible. The support condition requires that we observe all covariate values for both values of the instrument. We cannot otherwise separate the sources of variation in the treatment.[27]

For further insight, I examine covariates. Table 2.1 displays covariate averages by KOL presence and organization status. It presents coefficients from logistic regressions as well. I find similar patterns for the instrument and the treatment. As the logistic analysis shows, the instrument improves balance in such idiosyncratic characteristics as industry and firm size. Discrepancies remain in the proportion of strikes against multiple establishments and in wages at affected establishments. On the other hand, it aggravates imbalances in market conditions. It is particularly associated with industrialized urban communities with a history of successful walkouts. This correlation arises for two reasons: first, KOL assemblies and labor strife were both concentrated in manufacturing centers; second, there is little variation in the instrument within localities (unlike the treatment). As Subsections 2.6.2 and 2.7.1 show, individual circumstances influence conflict outcomes more than market conditions. Therefore, the instrument seems helpful.

In addition, Table 2.1 characterizes compliers.[28] Compliers may differ from the population even if the instrument is valid, which could distort my estimates. I find that these strikes are more likely to be offensive, to involve demands for fewer hours and to occur in urban areas. The Mideast is overrepresented, at the expense of New England and the Midwest, as is the construction industry, at the expense of mining and food, drink and tobacco. Nonetheless, compliers are broadly similar to other observations. Hence, my estimates might plausibly

[27] Common support is equivalent to the assumption of full rank or no collinearity in linear models.

[28] I characterize compliers via Abadie's (2003) weighting method. I reweight the sample for one covariate at a time. The weights are based on a local polynomial estimate of the conditional probability of KOL presence.

TABLE 2.1: COVARIATE MEANS AND LOGISTIC ANALYSIS OF KOL PRESENCE AND ORGANIZATION

	All	Compliers	KOL present		Organized		Coefficients from logit model of	
			No	Yes	No	Yes	KOL presence	Organization
Strike characteristics								
Strike of generic employees (%)[a]	26.981	32.194	22.945	28.181	26.565	27.264	−0.038 (0.142)	0.266 (0.144)
Women as a share of the workforce (%)[a]	8.822	8.295	10.126	8.434	11.049	7.306	−0.290 (0.324)	−0.774 (0.290)***
Strike against multiple establishments (%)[a]	20.380	23.182	14.972	21.988	12.569	25.697	0.492 (0.145)***	0.896 (0.130)***
Average size of affected establishments (log)	4.226	3.839	4.550	4.130	4.695	3.907	−0.015 (0.044)	−0.145 (0.048)***
Average wage at affected establishments (log)	0.706	0.816	0.567	0.747	0.588	0.786	0.773 (0.225)***	2.017 (0.261)***
Weekly hours at affected establishments – 60	0.152	0.085	0.457	0.062	0.430	−0.036	−0.012 (0.008)	0.001 (0.008)
Defensive strike (%)[a]	25.340	18.627	27.746	24.625	26.703	24.412	0.152 (0.128)	0.101 (0.117)
Cause: compensation (%)[a]	68.525	62.119	76.810	66.062	75.184	63.992	−0.230 (0.123)	−0.039 (0.157)
Cause: hours (%)[a]	15.961	27.032	5.370	19.110	7.689	21.592	0.202 (0.267)	0.738 (0.250)***
Cause: union rights (%)[a]	7.533	2.253	5.370	8.176	2.302	11.094	−0.146 (0.218)	1.600 (0.238)***
County characteristics								
Industrial workers (percentage, all workers)	34.756	35.233	31.186	35.818	33.986	35.281	0.027 (0.017)	0.009 (0.013)
Urban industrial workers (percentage, all workers)	25.821	30.018	15.171	28.987	22.344	28.187	0.053 (0.014)***	−0.016 (0.007)***
Gender fragmentation (index, industrial workers)	26.708	30.762	23.823	27.565	23.859	28.646	−0.007 (0.006)	0.021 (0.005)***
Ethnic fragmentation (index, industrial workers)	67.480	73.034	58.181	70.244	63.783	69.996	0.001 (0.006)	0.001 (0.005)
Trade fragmentation (index, industrial workers)	92.700	95.355	89.178	93.747	90.920	93.911	0.022 (0.013)	0.021 (0.012)
Average establishment size (log, manufacturing)	2.723	2.783	2.446	2.805	2.621	2.792	−0.340 (0.215)	−0.407 (0.172)***
Average daily wage (log, manufacturing)	0.142	0.221	−0.039	0.196	0.069	0.191	0.963 (0.432)***	1.575 (0.377)***
Railroad tracks (km / km^2)	0.203	0.274	0.108	0.231	0.163	0.231	4.135 (1.710)***	1.416 (0.672)***
Past labor conflict								
Successful strike in previous year (%)[a]	53.387	59.414	31.300	59.294	43.810	59.651	0.342 (0.148)***	0.241 (0.128)

Continues...

Covariate means and logistic analysis of KOL presence and organization (continued)

	All	Com- pliers	KOL present No	KOL present Yes	Organized No	Organized Yes	Coefficients from logit model of KOL presence	Coefficients from logit model of Organization
Period								
1881 (%)[a]	11.635	8.827	18.633	9.555	13.720	10.216		
1882 (%)[a]	10.517	9.037	17.738	8.370	12.155	9.401		
1883 (%)[a]	11.281	7.549	12.693	10.861	11.188	11.344	0.730 (0.287)***	–0.090 (0.201)
1884 (%)[a]	10.554	14.470	8.788	11.079	10.727	10.436	1.202 (0.307)***	–0.205 (0.206)
1885 (%)[a]	14.954	14.932	16.762	14.417	15.976	14.259	0.863 (0.285)***	–0.209 (0.195)
Before May, 1886 (%)[a]	12.381	14.047	9.032	13.377	13.812	11.407	1.690 (0.327)***	–0.307 (0.226)
After May, 1886 (%)[a]	28.678	30.331	16.355	32.342	22.422	32.936	1.195 (0.316)***	–0.263 (0.178)
Region								
New England (%)[a]	11.710	6.321	23.190	8.297	16.114	8.712		
Mideast (%)[a]	42.812	52.473	30.513	46.468	36.648	47.007	0.839 (0.366)***	0.527 (0.281)
Great Lakes (%)[a]	29.778	19.623	23.515	31.640	28.039	30.962	1.220 (0.440)***	0.653 (0.331)***
Plains (%)[a]	8.913	9.356	13.181	7.644	12.155	6.706	0.852 (0.510)	–0.418 (0.368)
South (%)[a]	4.643	3.400	8.055	3.628	5.157	4.293	1.481 (0.517)***	1.313 (0.435)***
West (%)[a]	2.144	2.680	1.546	2.322	1.888	2.319	2.188 (0.546)***	1.121 (0.454)***
Sector								
Mining and quarrying (%)[a]	15.626	7.455	29.699	11.442	22.744	10.780		
Construction (%)[a]	9.621	23.136	5.858	10.740	6.860	11.501	0.089 (0.258)	–0.408 (0.239)
Food, drink and tobacco (%)[a]	12.922	5.482	10.334	13.691	3.269	19.492	0.304 (0.225)	2.317 (0.307)***
Light manufacturing (%)[a]	34.253	37.384	30.757	35.293	32.505	35.443	0.015 (0.237)	0.266 (0.213)
Heavy manufacturing (%)[a]	19.746	16.581	17.575	20.392	21.501	18.552	–0.140 (0.253)	–0.246 (0.228)
Services (%)[a]	7.831	4.857	5.777	8.442	13.122	4.231	–0.034 (0.309)	–1.610 (0.332)***
McFadden's pseudo R^2							0.311	0.274

[a] Regression coefficients were divided by a hundred.
Notes: The means of covariates among compliers are estimated via reweighting (Abadie, 2003). Data about past strikes are not available for 1881. Both regressions include an intercept and exclude observations from 1881 (for a total of 37 parameters and 4739 observations). Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

generalize to the entire sample.

For comparison, I estimate the organization effect under selection on observables too (Rosenbaum and Rubin, 1983). The identification framework is similar to Abadie's (2003): the treatment should satisfy conditional independence and common support. These estimates are only consistent if unionization does not affect the probability of a stoppage.

2.5 Empirical strategy

2.5.1 Estimation

I am interested in the impact of union involvement on strike outcomes – in particular, the success rate. Because my primary outcome is binary, linear regression is inconsistent. Hence, I adopt a two-stage weighting approach instead. I draw on Abadie (2003), who shows that valid instruments identify the entire marginal distributions of compliers' potential outcomes, and Frölich and Melly (2013), who apply this insight to the estimation of unconditional treatment effects.[29]

For each observation i, write y_i for the outcome of interest, d_i for union sponsorship (the treatment), z_i for the presence of an assembly of the KOL (the instrument) and $\boldsymbol{x}_i$ for the control vector. For concreteness, suppose that y_i is an indicator of success in the following.

In the first stage, I construct estimation weights w_i. The weighting scheme depends on the identification assumptions. Under selection on observables, I use inverse probability weighting (Hirano, Imbens and Ridder, 2003):

$$w_i = 1/\,\mathrm{P}(d_i \mid \boldsymbol{x}_i).$$

[29] Abadie (2003) develops a similar estimator of conditional effects. I follow Frölich and Melly (2013) for three reasons: first, unconditional effects are easier to interpret; second, his estimator requires that I model expected outcome values; and, third, his estimator seems imprecise, as his empirical illustration attests. Clarke and Windmeijer (2012) and Lewbel, Dong and Yang (2012) discuss the relative merits of alternative estimators of the effect of endogenous treatments on binary outcomes.

In the case of endogenous selection, the weighting scheme is due to Abadie (2003) and Frölich and Melly (2013):

$$w_i = (2d_i - 1)(2z_i - 1) / \mathrm{P}(z_i \mid \boldsymbol{x}_i).$$

These weights have two key properties. First, they overweight underrepresented observations in each instrument group, which improves balance in covariates. Second, they are negative when d_i differs from z_i, which helps us recover the treatment effect on compliers by cancelling the contribution of noncompliers.[30] This first step requires an estimate of the conditional probabilities $\mathrm{P}(d_i \mid \boldsymbol{x}_i)$ and $\mathrm{P}(z_i \mid \boldsymbol{x}_i)$. I use a logistic specification. Alternative parametric estimators yield similar results. My sample is too small for more flexible models.

In the second stage, I regress y_i on d_i and an intercept:

$$(\hat{\alpha}, \hat{\beta}) = \arg\min_{\alpha,\beta} \left\{ \hat{\mathrm{E}}\left[w_i (y_i - \alpha - \beta d_i)^2 \right] \right\}.$$

The constant $\hat{\alpha}$ estimates the average baseline success rate – i.e. the success rate of wildcat walkouts. The coefficient $\hat{\beta}$ gives the average treatment effect. It is equal to the difference in weighted mean outcomes between unorganized and organized strikes. Under selection on observables, $\hat{\alpha}$ and $\hat{\beta}$ pertain to the entire sample; under endogenous selection, to compliers.[31]

How does weighting improve on linear regression? Weighting does not require parametric assumptions in principle; hence, it readily accommodates binary responses. Fully nonparametric estimation is difficult in practice though because conditional probabilities are subject to the curse of dimensionality. The choice of a model for $\mathrm{P}(d_i \mid \boldsymbol{x}_i)$ and $\mathrm{P}(z_i \mid \boldsymbol{x}_i)$ is important, as one

[30] Hirano, Imbens and Ridder (2003) develop an early application of inverse probability weighting to treatment evaluation. See also Firpo and Pinto (2016). Note that the two weighting schemes coincide when the instrument is the treatment itself: Frölich and Melly (2013) thus generalize inverse probability weighting in the same sense as two-stage least squares generalizes ordinary least squares.

[31] In terms of averages, $\hat{\alpha} = \hat{\mathrm{E}}[(1 - d_i) w_i y_i] / \hat{\mathrm{E}}[(1 - d_i) w_i]$ and $\hat{\beta} = \hat{\mathrm{E}}\left[d_i w_i (y_i - \hat{\alpha})\right] / \hat{\mathrm{E}}(d_i w_i)$.

must forecast individual probabilities in constructing the weights w_i. On the other hand, linear regression imposes linearity on the conditional expectation of outcomes. While ordinary least squares give the best linear approximation to the average effect of conditionally exogenous treatments, this property does not extend to instrumental regression under endogenous selection (Abadie, 2003; Lewbel, Dong and Yang, 2012). Note that these methods yield the same treatment effect without covariates (i.e. $\boldsymbol{x}_i = 1$).[32]

2.5.2 *Inference*

Correct inference must account for correlated errors. My sample features dependence by design: as Subsection 2.3.1 explains, the *Third Report* could present a single strike against multiple establishments as several rows. Local shocks and dynamic bargaining may also induce intrinsic correlation over time and space. For example, Biggs (2002) analyzes the marches of May 1886 in Chicago as a sequence of interactions, spanning preemptive concessions, violent pickets and uneasy truces.

My approach is based on Conley (1999). I assume that the maximum possible residual correlation between observations decreases with distance in time and space. In other words, I allow for arbitrary correlation between two stoppages if they begin in the same place on the same day, but distant episodes must be effectively independent. It seems plausible that labor strife in New York would be more likely to spill over to Brooklyn than San Francisco.

By Theorem 6.1 of Newey and McFadden (1994), my estimators are asymptotically normal. Their limit variances take the form $\mathrm{E}(v_{ij}\boldsymbol{h}_i\boldsymbol{h}_j^\top)$ for some weights v_{ij} and some vector function $\boldsymbol{h}_i$. The weights v_{ij} capture the residual correlation between observations. (See Section 2.B for the formula for $\boldsymbol{h}_i$.) Write $r_t(i, j)$ for the difference in start dates and $r_s(i, j)$ for the spatial distance

[32] Inverse probability weighting and ordinary least squares yield the same intercept as well (equal to the raw average success rate of unorganized strikes). On the other hand, Frölich and Melly (2013) do not compute the same intercept as two-stage least squares: weighting gives compliers' average baseline outcome, whereas regression estimates a mixture of compliers' and never-takers' (Abadie, 2003).

between observations i and j.[33] Let k be a kernel function. Let b_t and b_s be bandwidths. My variance estimator is:

$$\hat{\mathrm{E}}\left[k\left(\sqrt{[r_t(i,j)/b_t]^2+[r_s(i,j)/b_s]^2}\right)\hat{\boldsymbol{h}}_i\hat{\boldsymbol{h}}_j^\top\right].$$

The kernel term $k(\ldots)$ bounds v_{ij} in absolute value. Other than regularity conditions, consistency requires that b_t and b_s increase with the sample size at an appropriate rate, relaxing the bound on v_{ij}.

I set b_t to one year, so the bound on the correlation between two observations in the same locality is no lower than 0.75 if they start within three months of each other. I set b_s to 380 km, so the bound on the correlation between two observations in the same county is no lower than 0.75 if they start on the same date. I use the Parzen kernel for k. My findings are robust to these choices.[34]

2.6 Impact of the KOL on labor strife

2.6.1 Effect on strike incidence

This subsection investigates the impact of unionization on the incidence of labor strife. Recall from Section 2.4 that organized workers might undertake riskier strikes than the unorganized if union support increases their probability of success. This possibility suggests that unionization could increase the frequency of conflict. I cannot properly test this hypothesis because I do not observe unionization rates in the 1880s. I can however address a related question: whether

[33] The start date is incomplete for 13 observations, in which case I impute the first day of the month. If a strike affected multiple localities (110 observations), I base distances on the first entry. Since the *Third Report* does not seem to list localities in a logical pattern, I assume that the first entry was the main theater of events. If the locality is a county, I base distances on the coordinates of its centroid.

[34] The standard error on my benchmark estimate of the average effect of union sponsorship on the success rate is 0.095. If I set b_s to 570 km and b_t to 547 days (an increase of fifty percent), the standard error becomes 0.094. Additional results are available upon request.

Table 2.2: Impact of KOL assemblies on strike incidence

	(1)	(2)	(3)	(4)	(5)
KOL effect					
Two-stage least squares	0.451***	0.217***	0.223***	0.156***	0.147***
	(0.025)	(0.026)	(0.034)	(0.029)	(0.049)
Ordinary least squares	0.328***	0.160***	0.173***	0.105***	0.060***
	(0.016)	(0.011)	(0.014)	(0.010)	(0.013)
Fit					
First-stage F-statistic	441.194	186.050	166.670	149.084	121.862
Adjusted R^2 (OLS)	0.198	0.378	0.349	0.422	0.536
Controls					
Lagged outcome		×		×	×
County characteristics			×	×	
County effects					×
Region effects		×	×	×	
Year effects		×	×	×	×
Sample					
Counties	2519	2519	2519	2519	2519
Years	5	5	5	5	5
Parameters	2	12	19	20	2525

Notes: See Subsection 2.6.1 for details. Standard errors, in parentheses, are clustered at the county level. The first-stage F-statistic tests the effect of the instrument on the treatment in a linear specification.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

strike incidence correlates with the presence of the KOL.

To that end, I construct a panel of counties from 1882 to 1886. The outcome is an indicator of the occurrence of a walkout. The treatment is an indicator of the existence of a local assembly of the KOL in the previous year. To address attenuation bias from measurement error, I construct two instruments: an indicator of the existence of assemblies in neighboring counties in the previous year and the fraction of the labor force in neighboring counties with assemblies in the previous year. Table 2.2 shows my estimates. Since I combine two instruments and one is continuous, I use linear regression in lieu of reweighting. Standard errors are clustered at the county level.

The first specification does not include covariates. KOL presence is associated with an increase in strike incidence from 1 to 46 percent per year. This effect decreases considerably once I correct it for differences in economic development between counties, which correlates with both unionization and striking. The fourth specification adds controls for past stoppages, market characteristics, region and year. KOL presence is now associated with a shift in strike incidence from 6 to 22 percent. The last specification includes county effects. The coefficient drops from 17 to 15 percentage points.

These estimates are not causal, so one should not take them at face value. Nonetheless, they provide circumstantial evidence in favor of my hypothesis. Moreover, it seems that my covariates are a good proxy for market conditions so far as they do not change over time.

2.6.2 Effect on union sponsorship and the success rate

This subsection explores the impact of the KOL on work stoppages. First, I estimate their effect on the probability of union intervention. This exercise is analogous to the first stage of linear regression. Second, I estimate their effect on the success rate, which is the reduced form of my main specification. Unlike the previous subsection, the unit of observation is the strike. Table 2.3 presents my results. Table 2.10 in Section 2.C shows their linear counterparts.

The existence of an assembly in the locality of the dispute increased the probability of union involvement by 26.1 percentage points from a baseline of 39.4 percent to 65.5 percent. This estimate is robust to balance adjustments. The third specification is the exception. It takes idiosyncratic strike characteristics into account, which absorb much of the variation in union support, so the coefficient falls from 26.1 to 16.8 percentage points. These results are precise, so the instrument should satisfy the correlation condition (see Section 2.4).

The Knights had a more modest impact on the success rate. Their presence raised it by 8.6 percentage points from a baseline of 44.9 percent to 53.5 percent. This estimate is robust to balance corrections as well, though I lose precision if I adjust it for differences in market

TABLE 2.3: IMPACT OF KOL ASSEMBLIES ON ORGANIZATION AND SUCCESS RATES

	(1)	(2)	(3)	(4)	(5)	(6)
Outcome: organization						
KOL effect	0.262***	0.223***	0.168***	0.234***	0.269***	0.223***
	(0.032)	(0.029)	(0.029)	(0.052)	(0.034)	(0.028)
Baseline probability	0.393***	0.409***	0.448***	0.376***	0.392***	0.413***
	(0.021)	(0.023)	(0.024)	(0.044)	(0.023)	(0.025)
Adjusted R^2 (OLS)	0.050	0.159	0.282	0.202	0.050	0.166
Outcome: success						
KOL effect	0.086***	0.089***	0.065***	0.093***	0.112***	0.110***
	(0.024)	(0.022)	(0.023)	(0.040)	(0.025)	(0.024)
Baseline probability	0.449***	0.439***	0.461***	0.428***	0.418***	0.409***
	(0.017)	(0.020)	(0.023)	(0.043)	(0.018)	(0.023)
Adjusted R^2 (OLS)	0.005	0.045	0.074	0.049	0.008	0.049
Controls						
Period, region & sector		×	×	×	×	×
Strike characteristics			×			
County characteristics				×		
Past labor conflict					×	×
Sample						
Sample size	5363	5363	5363	5363	4739	4739
Year 1881	×	×	×	×		
Parameters	2	18	28	26	2	18

Notes: See Section 2.3 for information about the data. Column (6) excludes observations from 1881 for lack of strike microdata for 1880. Unless noted, the table shows reweighted estimates (see Subsection 2.5.1). Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

characteristics. As the coefficient of determination demonstrates, covariates exert little influence on the success rate.[35]

Machado, Shaikh and Vytlacil (2018) develop tests of the exclusion restriction when the instrument, the outcome and the treatment are all binary. It exploits the fact that exogeneity bounds the coefficient from the reduced form: if the instrument only affects outcomes through the treatment, the correlation between the instrument and the outcome should be neither too small nor too large. I implement their test of the null hypothesis of an invalid instrument under the monotonicity assumption (see Section 2.4). Because the procedure uses bootstrapped critical values, I cannot follow my preferred inference strategy (see Subsection 2.5.2). I cluster critical values at the county level instead. I reject the null hypothesis at any conventional significance level: the test statistic is 3.58, which is comfortably larger than the critical value of 2.36 at the one-percent level. This result provides additional evidence in favor of my identification strategy.

2.7 Impact of union sponsorship on labor strife

2.7.1 Effect on the success rate

Table 2.4 presents my main results: the average effect of union sponsorship on the probability of success of a strike. The first two rows contain estimates by the weighting method for endogenous treatments of Frölich and Melly (2013). The first row shows the average treatment effect and the second row shows the baseline success rate. The following two rows contain analogous estimates by inverse probability weighting for conditionally exogenous treatments (Hirano, Imbens and Ridder, 2003). The fifth row displays the m-statistic of Hausman (1978), which tests the estimates of the union effect for equality. The sixth row gives the coefficient of determination. The seventh

[35] This finding might seem surprising, as covariates explain much of the variation in conflict incidence and union support. To understand it, consider the following schematic model. Workers draw a probability of success, which is either 0.4 or 0.6. Covariates influence their probability of drawing 0.4 or 0.6. Workers strike if it is 0.6. Then, there is no variation in the success rate for covariates to explain, but they influence the likelihood of a strike.

row shows the first-stage partial F-statistic.[36] The control set changes across columns. The last specification excludes observations from 1881 for lack of stoppage data for 1880. For comparison, the fifth column shows estimates under the benchmark specification without observations from 1881 as well. Table 2.10 in Section 2.C shows their linear counterparts.

Column (1) is my benchmark. It ignores imbalances in covariates. Wildcat strikers' mean success rate was 37.8 percent across compliers (second row) and 44.2 percent across the entire sample (fourth row). These estimates are similar, which suggests that compliers form a representative subsample of the population. They confirm the intuition that workers walked out when they stood a reasonable chance of winning (Biggs, 2002). The average causal effect of organization was 32.7 percentage points (first row), which implies that union intervention increased the probability of victory from 38 to 71 percent. In line with Card and Olson (1995) and Friedman (1988), the naive estimate is 12 percentage points (third row). These two coefficients are statistically different at the five-percent level: the m-statistic is 2.436 (fifth row).

Columns (2) through (6) take covariates into account. The second specification balances the instrument across periods, regions and sectors. The causal organization effect rises from 32.7 to 39.7 percentage points. I obtain similar numbers after including controls for idiosyncratic characteristics and market conditions, but I lose precision. The adjusted coefficients are not statistically different from the benchmark.[37] The last specification includes an indicator of successful past walkouts in the same sector and county. The average treatment effect becomes 43.8 percentage points. Note that this specification excludes observations from 1881, since I

[36] The m-statistic is asymptotically normally distributed. Given a linear regression of the treatment on the instrument and covariates, the first-stage partial F-statistic is the squared t-statistic for the zero null hypothesis. It follows a χ^2_1 distribution asymptotically. It is a measure of instrument strength and relates to the share of compliers in the sample (cf. Section 2.4).

[37] As Column (4) shows, my estimate is exceptionally imprecise if I reweight the sample for differences in county characteristics. The additional noise is due to the overlap between assemblies' location and urbanization (cf. Table 2.1), which threatens the assumption of common support. Unlike other specifications, this estimate is moreover sensitive to the choice of urbanization measure and of the model of $\mathrm{P}(d_i \mid \boldsymbol{x}_i)$ (see Subsection 2.5.1). The table shows my lowest and noisiest estimate.

TABLE 2.4: EFFECT OF UNION SPONSORSHIP ON THE STRIKE SUCCESS RATE

	(1)	(2)	(3)	(4)	(5)	(6)
With instrument						
Union effect	0.327***	0.397***	0.364***	0.286	0.414***	0.446***
	(0.095)	(0.106)	(0.155)	(0.820)	(0.102)	(0.127)
Baseline rate	0.378***	0.365***	0.431***	0.255***	0.340***	0.347***
	(0.053)	(0.078)	(0.108)	(0.061)	(0.056)	(0.108)
No instrument						
Union effect	0.122***	0.108***	0.109***	0.099***	0.123***	0.096***
	(0.026)	(0.022)	(0.026)	(0.024)	(0.027)	(0.024)
Baseline rate	0.442***	0.447***	0.441***	0.450***	0.432***	0.444***
	(0.014)	(0.017)	(0.022)	(0.019)	(0.014)	(0.017)
Fit						
Hausman m-statistic	2.436	2.844	1.642	0.227	3.163	2.819
First-stage F-statistic	68.462	57.804	27.225	16.362	63.185	46.679
Adjusted R^2 (OLS)	0.014	0.049	0.077	0.053	0.014	0.051
Controls						
Period, region & sector		×	×	×		×
Strike characteristics			×			
County characteristics				×		
Past labor conflict						×
Sample						
Sample size	5363	5363	5363	5363	4739	4739
Year 1881	×	×	×	×		
Parameters	2	18	28	26	2	18

Notes: See Section 2.3 for information about the data. Column (6) excludes observations from 1881 for lack of strike microdata for 1880. Unless noted, the table shows reweighted estimates (see Subsection 2.5.1). Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2). The m-statistic tests the difference between estimates of the union effect (Hausman, 1978). The first-stage F-statistic tests the effect of the instrument on the treatment.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

lack strike data for 1880. This restriction explains most of the change in the causal estimate: if I compute the organization effect without controls (like Column (1)) and without observations from 1881 (like Column (6)), I obtain 41 percentage points (Column (5)).

These results support the hypothesis that unions help workers win strikes. This advantage rationalizes unions' growing role in industrial conflict (cf. Figure 2.1). Moreover, I find evidence of downward bias in the naive estimate of the organization effect, which suggests that workers adjust their bargaining strategies in response to the availability of union support and confront stronger employers on average than the unorganized. In other words, unionization expands their tactical inventory, enabling them to strike in less favorable circumstances. (See Subsection 2.7.3 for a discussion of possible mechanisms.)

Unions were actually ambivalent about industrial action in the 1880s. Besides the financial toll, stoppages were fraught with danger: job loss, blacklisting, violence against picket lines, jail terms and more (Currie and Ferrie, 2000; Rosenbloom, 1998). Officers worried that a defeat might threaten the survival of the association (Kremer and Olken, 2009), as it could depress morale and wreck leaders' prestige. Competition for members among associations was fierce (Kaufman, 2001). Moreover, industrial disruption antagonized sympathetic employers and public authorities (Friedman, 1988; Voss, 1993). Therefore, union executives had reason to avoid conflict (especially if victory was uncertain).[38] However, these fears clashed with the interests of the rank and file. Workers unionized to maximize their own welfare by increasing wages, decreasing hours and improving work conditions (Eichengreen, 1987; Kremer and Olken, 2009). Unions could hardly disavow strikers, lest it weaken their appeal to existing members and potential recruits (Perlman, 1918). Local officers were particularly willing to endorse unauthorized picketing (Card and Olson, 1995; Kremer and Olken, 2009). My results

[38] Postwar commentators made the opposite argument: for personal and ideological reasons, union leaders were more belligerent than the rank and file. This view motivated legislation to condition industrial action on secret ballots. For example, see Moore (2013, 2016) or Olofsgård (2012). This difference may be due to the institutionalization of union rights in the 20th century.

imply that unions were not an effective moderating force in this period: workers were able to impose facts on the ground and extract support for difficult confrontations.

There is evidence of such tension within the KOL. For example, the leading article of the *Journal of United Labor* of June 1882 bemoaned that the mechanic hung on "to the *old* barbarous, clumsy, unyielding, and treacherous system, known as strike, for his own personal benefit [emphasis in the original]".[39] It went on to berate the cost and uncertain benefits of work stoppages. The *Journal* later quoted the *Chicago Express*: "The striking mania among the workers has partly yielded to judicious counsel. Organization is regulating it, and will presently control it fully [...]. Strikes are voted down as disorderly and leading to bloodshed." Yet this very edition contains an appeal for aid from embattled miners in Maryland. It begins: "Whilst I have condemed [sic] without stint the strike system, it is not without purpose, or to no good in all cases, when I witness the action of capital in demanding of their employees that they work twelve hours for a day's work." It is clear that local assemblies paid lip service to official guidelines against walkouts. At the General Assembly of 1882, Grand Master Powderly declared: "One cause for the tidal wave of strikes that has swept over my Order comes from the exaggerated reports of the strength of the Order, numerically and financially, given by many of my organizers. Such a course may lead men into the Order, but by a path that leads them out again [...]" (Wright, 1887). Nonetheless, assembly representatives seized the occasion to legalize strike relief. They reversed this position in 1884.[40] Nor were these disagreements exclusive to the KOL. Chicago's *Inter Ocean* reported the following resolution in May 1881: "The Tinners and Sheet-iron Workers' Union, No. 1, of Chicago, are not organized in the spirit of a strike; [...]

[39] The *Journal of United Labor* was the official journal of the KOL. It circulated between 1880 and 1889. This edition is the second number of the third volume, published by Robert D. Layton in Pittsburgh (PA).

[40] On this occasion, Powderly observed: "[...] many new Assemblies are deceived on being organized; they are told by the Organizer that the assistance fund is laying idle [...]. These members, thinking that they are entitled to this fund, become obnoxious and troublesome to their employers [...];the result is a lock-out and trouble" (Powderly, 1884). Although he refers to lockouts, this quote shows that workers could become bellicose if they believed that union support was forthcoming.

There is a spirit of discontent prevalent among the different branches of our trade; therefore, [...] the Union will not hold itself responsible for the acts of individual members."

Note that my findings do not support theories of asymmetric information between officers and the membership. Ashenfelter and Johnson (1969) and Olofsgård (2012) argue that union leaders have access to private information (e.g., the company's books), which could help them forecast conflict outcomes. Organized workers should then learn that certain disputes are hopeless, so their baseline success rate would be higher than wildcat strikers' and naive estimates, upward biased. Information asymmetries must thus have been relatively unimportant in the 1880s, though they may have grown with the institutionalization of collective bargaining in the 20th century.

2.7.2 Effect on payoffs

The previous subsections analyzed the probability of success. This subsection focuses on payoffs. Table 2.5 presents my findings.

Workers did not always return to their jobs at the end of hostilities (Currie and Ferrie, 2000; Rosenbloom, 1998). Some found alternative employment during the standoff. Others were permanently replaced by strikebreakers. Some employers refused to reinstate strike leaders in particular, though they might offer concessions to other workers. Discharged employees were often blacklisted as well. Column (1) investigates unions' influence over dismissals. The *Third Report* does not specify job losses, but it gives the change in firm size and the number of new employees by gender, which allows me to approximate the incidence of layoffs. As the baseline rate shows, there were layoffs in more than half of all disputes. Organization offered workers protection: the causal estimate implies that unions decreased the incidence of job loss by 22.6 percentage points or nearly half. This coefficient is four times greater than the naive estimate, which conforms with the hypothesis that organized labor took more risk than the

TABLE 2.5: EFFECT OF UNION SPONSORSHIP ON STRIKE PAYOFFS

	Job loss (1)	Weekly hours (2)	Daily wage (3)	Daily wage (4)
With instrument				
Union effect	–0.226***	–0.984	0.018	–0.025
	(0.070)	(3.070)	(0.024)	(0.054)
Baseline outcome	0.561***	–8.344***	0.129***	–0.101***
	(0.051)	(2.595)	(0.013)	(0.041)
No instrument				
Union effect	–0.057**	–2.558***	0.018***	–0.003
	(0.023)	(1.163)	(0.006)	(0.010)
Baseline outcome	0.502***	–5.295***	0.122***	–0.128***
	(0.016)	(0.899)	(0.003)	(0.006)
Fit				
Hausman m-statistic	–2.486	0.578	0.019	–0.443
First-stage F-statistic	68.462	15.115	37.126	6.803
Sample				
Cause	Any	Hours	Wage raise	Wage cut
Result	Any	Success	Success	Defeat
Sample size	5363	381	1498	367

Notes: See Section 2.3 for information about the data. Wage and hours regressions exclude strikes after which all strikers lost their jobs. Wage regressions use log wages. Unless noted, the table shows reweighted estimates (see Subsection 2.5.1). Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2). The m-statistic tests the difference between estimates of the union effect (Hausman, 1978).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

wildcat.

Column (2) examines the change in weekly hours. Following Card and Olson (1995), I restrict the sample to successful strikes for a shorter workweek. Unorganized compliers achieved an average reduction of 8.3 hours. Unionization had no significant impact. Column (3) reports similar findings for the change in daily pay after a successful stoppage for a wage raise: the mean baseline increase is 12.9 percent and the organization effect is insignificant. On the other hand, the naive estimates are significant (–2.6 hours and 1.8 percentage points, respectively), which suggests that officers may have been sensitive to pressure from the ranks over the terms of settlement as well as the decision to strike.[41]

I find no significant organization effect on the payoff of successful stoppages. Note however that there was a significant effect on the expected outcome of a walkout, since it depends on the probability of success in addition to the realized payoff. Moreover, a rough estimate indicates that the benefits outweighed the costs of organized strikes on average. The KOL charged $15 to charter a new assembly. The minimum membership was ten workers, so suppose that my hypothetical worker contributed $1.5. There were also an induction fee ($1) and a quarterly membership fee ($0.25). Suppose that they struck after a year. Unionization cost $3.5. Its benefit is a higher success rate by 0.33, times a wage raise of 15 percent, times a mean initial daily wage of $2 for male strikers – i.e. 9.6 cents per day or $7.2 per quarter.[42]

The last column considers the decrease in daily wages after an unsuccessful stoppage against a wage cut. Daily pay fell by 10 percent on average among compliers. Neither estimator yields a significant organization effect. This result is unsurprising: employers announced wage cuts before workers struck, so union intervention should only affect payoffs through the probability of victory.

[41] Note that layoffs may induce spurious changes in wages and hours, which I cannot account for.

[42] This calculation ignores the impact of organization on strike duration. As Subsection 2.7.3 shows, I find no significant effect on duration.

TABLE 2.6: EFFECT OF UNION SPONSORSHIP ON STRIKE DEVELOPMENT

	With instrument		No instrument		Hausman
	Baseline	Union effect	Baseline	Union effect	*m*–stat.
Entire workforce on strike (indicator)	0.179***	0.506***	0.298***	0.111***	3.333
	(0.053)	(0.135)	(0.029)	(0.033)	
New workers from other places (indicator)	0.204***	–0.202***	0.121***	0.010	–3.543
	(0.030)	(0.065)	(0.009)	(0.017)	
New workers after strike (w.r.t. initial workforce)	0.093***	0.096***	0.115***	0.028***	1.864
	(0.023)	(0.039)	(0.008)	(0.010)	
Shutdown of affected firms (indicator)	0.701***	–0.280***	0.591***	–0.013	–2.398
	(0.064)	(0.120)	(0.031)	(0.027)	
Financial assistance (indicator)	0.098***	0.398***	0.084***	0.331***	0.816
	(0.032)	(0.101)	(0.010)	(0.028)	
Duration (log days)	2.329***	–0.186	1.965***	0.542***	–2.943
	(0.158)	(0.255)	(0.048)	(0.070)	

Notes: See Section 2.3 for information about the data. Each row presents presents estimates for a different outcome. The table shows reweighted estimates (see Subsection 2.5.1). Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2). The *m*-statistic tests the difference between estimates of the union effect (Hausman, 1978).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

2.7.3 Mechanisms

Subsection 2.7.1 argued that unions raised the probability of success of a strike. This subsection explores the mechanisms behind it. Table 2.6 reports my results.

I would ideally quantify the contribution of different channels to the organization effect. This exercise is infeasible though because strikers' tactics are endogenous. Consider for example financial assistance. A naive regression would implausibly have us believe that it lowered the success rate. Selection bias is the likely culprit: for instance, unions may have prioritized the most difficult confrontations in allocating funds. An accurate decomposition would thus require a separate instrument for each mechanism of interest. Given the limitations of my data, I adopt a simpler approach and compute the impact of unionization on the course of each dispute.

Association was partly an answer to such challenges to collective action as coordination failures and free riding. First, unions fostered solidarity through lectures, meetings, parades, songs, etc., which helped workers internalize their contribution to others' welfare and increased the social fallout of crossing the picket line. Second, officials could leverage their experience and the threat of expulsion to impose discipline, improve coordination and overcome mistrust. They could also accumulate bargaining expertise, which helped them negotiate better settlements. Thirdly, there were logistical advantages: for example, labor journals expanded the reach of boycotts.

As the first row of Table 2.6 shows, organization succeeded in boosting turnout: nearly seventy percent of union strikes involved the entire workforce of affected establishments, against 18 percent of the wildcat. Unions had a more complex impact on strikebreaking. As Subsection 2.7.2 noted, layoffs were rarer in organized stoppages. Unions were especially effective against outside replacement workers (second row). However, firms hired more permanent replacements if they hired them at all (third row): new employees represented a fifth of the initial workforce on average, against a tenth for unorganized stoppages. This difference may reflect increased participation in union strikes. Furthermore, affected establishments were less likely to close by 28 percentage points despite higher turnout (fourth row), which suggests that employers procured either temporary strikebreakers or help from other firms.

Few strikers had enough savings for a prolonged standoff. Some found alternative employment during stoppages, but many relied on outsiders for financial relief. Labor societies were the main providers, building resistance funds in peacetime and pooling risks across branches. When unions' resources proved insufficient, they coordinated donations (e.g., KOL assemblies pleaded for aid on the *Journal of United Labor*). As the fifth row of Table 2.6 shows, half of authorized walkouts received financial assistance, against a tenth of the unorganized. The difference between estimators is not significant, which is interesting in that strikers exerted little influence over aid.

The last row examines duration. Most disputes were short in this period: a quarter ended within three days and half ended within ten days. I find no causal effect on duration. This result is surprising to some extent: one would think that organization helped workers endure longer stoppages (through financial relief, for instance). However, employers might concede defeat earlier if they expect greater resistance from unionized workers. Therefore, the organization effect is ambiguous a priori.

2.7.4 Subgroup effects

Friedman (1988) studies the impact of union intervention on industrial conflict across two phases of the American labor movement: the radical experiment (1881–86), under the aegis of the KOL, and the return of craft unionism (1887–94), led by the AFL. He finds an increase in the organization effect after 1886, which he attributes to a change in strategy. The Knights sought the strength in numbers to bully employers into submission, he argues, but they could not provide strikers with adequate assistance or win the sympathy of hostile public authorities. By contrast, craft unions restricted membership to skilled workers and increased fees. Their walkouts were fewer, smaller, better planned and better funded – hence, more successful.

Table 2.7 presents estimates by strike size and skill level. The first column divides the sample according to the average workforce of affected establishments. As Friedman (1988) argued, organized workers had a lower baseline success rate against large employers in the early 1880s. However, union intervention was significantly more effective, compensating the lower intercept. Large strikes were probably susceptible to coordination problems, which unions could mitigate. I obtain the same pattern if I classify observations according to the number of strikers instead of firm size (second column). The last two columns are based on proxies for strikers' skill level: whether the *Third Report* specified their trade and whether their mean wage was higher than $2 before the conflict.[43] Differences between estimates are neither significant nor consistent. We find little evidence overall in support of Friedman (1988).

[43] In microdata from the *Census of Manufactures* (Atack and Bateman, 2016), $2 is the median average wage of skilled workers and the 99th percentile of the average wage of unskilled labor.

Table 2.7: Effect of union sponsorship on the strike success rate by subgroup

	Estab. size above median	Strike size above median	Strike of generic workers	Strikers' wage above $2
With instrument				
Union effect if outside subgroup	0.174	0.253***	0.349***	0.366***
	(0.103)	(0.099)	(0.098)	(0.114)
Baseline rate if outside subgroup	0.470***	0.417***	0.412***	0.309***
	(0.058)	(0.068)	(0.062)	(0.073)
Union effect if in subgroup	0.544***	0.418***	0.308***	0.311***
	(0.163)	(0.148)	(0.146)	(0.152)
Baseline rate if in subgroup	0.243***	0.332***	0.277***	0.497***
	(0.099)	(0.073)	(0.084)	(0.088)
Equality test (union effect)	1.918	0.926	−0.233	−0.287
Equality test (total probability)	1.008	0.564	−1.309	0.871
No instrument				
Union effect if outside subgroup	0.134***	0.150***	0.121***	0.127***
	(0.035)	(0.030)	(0.027)	(0.030)
Baseline rate if outside subgroup	0.458***	0.419***	0.457***	0.437***
	(0.019)	(0.016)	(0.017)	(0.015)
Union effect if in subgroup	0.100***	0.090***	0.128***	0.108***
	(0.027)	(0.032)	(0.048)	(0.038)
Baseline rate if in subgroup	0.430***	0.470***	0.402***	0.458***
	(0.016)	(0.020)	(0.022)	(0.028)
Equality test (union effect)	−0.754	−1.396	0.129	−0.395
Equality test (total probability)	−1.603	−0.241	−1.089	0.037
Share of sample in subgroup	0.492	0.491	0.270	0.395

Notes: The table shows reweighted estimates for each subgroup (see Subsection 2.5.1). For example, the first two estimates in the first column concern establishments whose size is below the median. Medians are taken within regions and sectors. The equality test is the t-statistic for the difference between subgroups. Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

2.8 Robustness tests

Subsection 2.7.1 examined the robustness of the estimate of unions' effect on the success rate to different control sets. This section explores additional robustness tests.

Table 2.8 investigates sample restrictions. The first row reduces the sample to strikes whose locality had not suffered stoppages in the previous year. It addresses the concern that the Knights may have targeted areas with a higher success rate. If a locality experienced few strikes in the past, it would have been harder for them to predict conflict outcomes. The organization effect is here equal to 45.5 percentage points. As this sample excludes observations from 1881, the relevant benchmark is Column (5) in Table 2.4, 41 percentage points.[44] The second row excludes localities which the KOL had not organized by 1886. The estimate becomes 30 percentage points, against 32.7 for my benchmark. The last two rows split the sample into a subsample without assemblies (third row) and a subsample with assemblies (fourth row). They indicate that unorganized strikes had similar success rates across instrument values. The naive organization effect is lower in locations without assemblies by four percentage points, but this estimate is imprecise, so the difference is not statistically significant.

Table 2.9 considers two alternative instruments. Because one is continuous, this table uses linear regression in lieu of reweighting. The first two columns use the existence of an assembly in the locality of the dispute in 1880 (instead of the preceding year). This specification should decrease the correlation between the instrument and the baseline success rate so far as its determinants change over time. The first column does not take covariates into account. The resulting estimate is 31 percentage points, against 32.7 for my benchmark. The second column uses the main instrument as a control, which should further diminish any residual correlation with determinants of the probability of success. The coefficient remains similar

[44] Note that this test is a more stringent version of Column (6) in Table 2.4, in which we reweight observations for the incidence of successful past strikes in the same sector and locality. Note too that this subset excludes most large cities (cf. Subsection 2.7.1).

TABLE 2.8: ROBUSTNESS OF THE UNION EFFECT TO SAMPLE RESTRICTIONS

	Sample size	Instrument	Baseline rate	Union effect
Localities without strikes in the previous year	1459	Yes	0.436***	0.455***
			(0.139)	(0.222)
Localities with a KOL assembly by 1886	5035	Yes	0.364***	0.300***
			(0.061)	(0.104)
Localities without KOL assemblies	720	No	0.436***	0.083
			(0.027)	(0.053)
Localities with KOL assemblies	4643	No	0.444***	0.125***
			(0.015)	(0.027)

Notes: See Section 2.3 for information about the data. The table shows reweighted estimates (see Subsection 2.5.1). The first row excludes observations from 1881 for lack of strike microdata for 1880. Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

at 28.1 percentage points. The last two columns replace the instrument with the log distance between the locality and the nearest assembly in 1880.[45] This formulation should lower the correlation between the instrument and the success rate to the extent that its determinants are specific to each locality. The estimates are somewhat larger than the benchmark: 40.9 and 44.1 percentage points. Because these instruments have less power, these coefficients are less precise and less robust than the benchmark. Nonetheless, they provide additional evidence in support of the exclusion restriction.

2.9 Conclusion

This paper explored the effect of unionization on strike outcomes in the United States in the early 1880s. To identify causal effects, I constructed an instrument from the location of the assemblies of the Knights of Labor. Organized strikers were significantly more successful than wildcat strikers: union sponsorship increased the probability of success of a strike by 32 percentage points from a baseline rate of 38 percent. This result rationalizes unions' leading

[45] Assemblies in the locality in question are ignored. Assemblies are also ignored if they were founded after an assembly was established in the locality.

TABLE 2.9: ROBUSTNESS OF THE UNION EFFECT TO ALTERNATIVE INSTRUMENTS

	(1)	(2)	(3)	(4)
Union effect				
Two-stage least squares	0.310***	0.281	0.409	0.441
	(0.115)	(0.203)	(0.295)	(0.409)
Ordinary least squares	0.122***	0.111***	0.122***	0.111***
	(0.026)	(0.025)	(0.026)	(0.025)
Fit				
Hausman m-statistic	1.813	0.868	0.989	0.816
First-stage F-statistic	39.469	9.967	9.823	6.781
Adjusted R^2 (OLS)	0.014	0.016	0.014	0.016
Controls				
KOL presence in year of strike		×		×
Instrument				
KOL presence in 1880	×	×		
Distance to nearest assembly in 1880 (log)			×	×
Sample				
Sample size	5363	5363	5363	5363
Parameters	2	3	2	3

Notes: See Section 2.3 for information about the data. Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2). The m-statistic tests the difference between estimates of the union effect (Hausman, 1978). The first-stage F-statistic tests the effect of the instrument on the treatment in a linear specification.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

role in collective bargaining in the postwar period. Organization reduced the probability of job loss as well. On the other hand, I found no effect on the payoff of successful walkouts.

Because wildcat stoppages are so few today, strike theory has not paid much attention to the interaction between unions and workers.[46] Most models assume pairwise bargaining between a firm and a union. Empiricists evaluate their predictions about duration or the impact of aggregate shocks (Card, 1990). This paper provides theorists with additional empirical evidence. Organization is a twofold shock: it lowers the cost of a standoff to workers (through financial assistance, etc.) and increases its cost to firms (by reducing strikebreaking, etc.). My results are consistent with an attrition model (Card and Olson, 1995; Geraghty and Wiseman, 2008; Kennan and Wilson, 1989). In this framework, firms and workers dispute a known indivisible surplus.[47] They pay a fixed delay cost per period of stoppage. They know their own costs, but not each other's. This model captures the effect of organization on the success rate as well as the lack of an effect on the payoff of successful strikes or duration.

This paper sheds new light on the American labor movement in the 1880s. This decade saw an unprecedented experiment in radical mass unionism under the aegis of the KOL. The KOL entered rapid decline in 1886, which entrenched conservative craft unionism and the American Federation of Labor, whereas radical inclusive unions rebounded from similar setbacks in Europe at the end of the century. This divergence is a topic of ongoing debate. Recent research has emphasized environmental constraints in the U.S. (Ansell and Joseph, 1998; Friedman, 1988; Kaufman, 2001; Voss, 1993). I find evidence that organized workers undertook riskier confrontations than the unorganized, which indicates a discipline problem within unions. Kremer and Olken (2009) make a similar point in the context of an evolutionary model of unionization. They argue that democratic unions are evolutionarily disadvantaged because they focus on maximizing members' welfare instead of their own survival. The Knights of Labor

[46] Exceptions include Ashenfelter and Johnson (1969) and Olofsgård (2012).
[47] The surplus need not be indivisible: see the behavioral model of Abreu and Gul (2000), for example.

were a loose federation of nearly autonomous assemblies, whereas the American Federation of Labor centralized power. Greater discipline may therefore elucidate the triumph of craft unionism in the U.S.

2.A Sample construction

The *Third Report* has 5809 rows. I exclude lockouts (358 rows), unfinished strikes (4 rows), general strikes (48 rows) and strikes in imprecise localities (36 rows). The general strikes are: the strike of the Amalgamated Association of Iron and Steel Workers of 1882, the nationwide strike of the Brotherhood of Telegraphers of 1883 and the Great Southwest Railroad Strike of 1886. The telegraphers' strike appears as one full row (in New York) and 44 empty rows (in other states). The imprecise localities are: Jersey Meadows (1 row), Hocking Valley (3 rows) and Western Pennsylvania (28 rows).

I follow these definitions in constructing covariates from the *Third Report* and the *Tenth Census*:

- *Strikes over pay:* mostly for a wage raise (70 percent) or against a wage cut (22 percent), but also over effective compensation (change of screen, payment in script, etc.).
- *Generic employees:* "employés", laborers or helping hands (labels from the *Third Report*).
- *Fragmentation index:* $100 \times (1 - \sum_i s_i^2)/(1 - 1/N)$, where s_i is the share of group i in the industrial workforce and N is the number of groups. Following census reports, I distinguish eight ethnic groups: Black American, White American, British, Canadian, Irish, German, Nordic and other. (German includes Austrians and the Swiss.) I use the trade classification of the *Tenth Census*.
- *Incidence of unemployment:* the percentage of workers that experienced at least one month of unemployment in the twelve months before the census.

- *Successful strike in previous year:* the occurrence of a successful strike in the same sector and county in the preceding year.

I construct two variables from the tables of the census of manufactures: the average establishment size and the average daily wage.[48] Because the census reports more establishments than employees for a few counties, I add one employee to each establishment before taking averages. If there were fewer than five establishments in a county, I substitute state figures. To compute railways per square kilometer, I calculate the land area of each county from boundary files (Manson et al., 2018).

I divide the sample into seven periods. The first five are yearly (1881–85). Following Card and Olson (1995), I divide the eight-hour campaign of 1886 into two stages: January to April (buildup) and May to December (fallout). I aggregate states into six regions. Four correspond to the definitions of the Bureau of Economic Analysis (BEA): the Great Lakes, the Mideast, New England and the Plains. The South includes the BEA region of the same name, Oklahoma and Texas. The West combines two BEA regions, Far West and Rocky Mountain, Arizona and New Mexico. I aggregate industries into six sectors: mining and quarrying (coal, ice, metal and stone), construction (building trades, public ways and public works), food, drink and tobacco (agriculture, food, drink and tobacco), light manufacturing (ceramics, clothing, leather, paper, printing, rubber, textiles, wood and other manufacturing), heavy manufacturing (chemicals, coke, gas, machinery, metals and transportation equipment) and services (communications, government, services, trade and transportation).

[48] To compute average daily wage, I divided total yearly wages across manufacturing firms in each county by the number of employees times 300.

2.B Variance of weighting estimators

My parametric implementation of the weighting estimators of Hirano, Imbens and Ridder (2003) and Frölich and Melly (2013) are asymptotically normally distributed by Theorem 6.1 of Newey and McFadden (1994). Their limit variances take the form $\mathrm{E}(v_{ij}\boldsymbol{h}_i\boldsymbol{h}_j^\top)$ for some weights v_{ij} and some vector function $\boldsymbol{h}_i$. The weights v_{ij} capture the residual correlation between observations.

I use the notation of Section 2.5. Since the weights of Frölich and Melly (2013) simplify to inverse probability weighting when z_i is d_i, I focus on the more general case. Let $\Lambda(\cdot) \equiv \exp(\cdot)/[1+\exp(\cdot)]$ be the logistic function. Recall that I set $\mathrm{P}(z_i \mid \boldsymbol{x}_i) = \Lambda(\boldsymbol{x}_i^\top \boldsymbol{\gamma})$ for some vector $\boldsymbol{\gamma}$. Define $\boldsymbol{d}_i \equiv (1, d_i)$ and $\boldsymbol{\beta} \equiv (\alpha, \beta)$.

Newey and McFadden (1994) give the formula for $\boldsymbol{h}_i$:

$$\mathrm{E}(w_j \boldsymbol{d}_j \boldsymbol{d}_j^\top)^{-1} \left(w_i \boldsymbol{d}_i (y_i - \boldsymbol{d}_i^\top \boldsymbol{\beta}) - \mathrm{E}\left[\boldsymbol{d}_j (y_j - \boldsymbol{d}_j^\top \boldsymbol{\beta}) \, \mathrm{D}_{\boldsymbol{\gamma}^\top} w_j \right] \mathrm{E}(\boldsymbol{x}_j \boldsymbol{x}_j^\top)^{-1} \boldsymbol{x}_i [d_i - \Lambda(\boldsymbol{x}_i^\top \boldsymbol{\gamma})] \right),$$

where $\mathrm{D}_{\boldsymbol{\gamma}^\top} w_i = -\boldsymbol{x}_i (2d_i - 1) \Lambda(\boldsymbol{x}_i^\top \boldsymbol{\gamma})[1 - \Lambda(\boldsymbol{x}_i^\top \boldsymbol{\gamma})] \left(z_i / \Lambda(\boldsymbol{x}_i^\top \boldsymbol{\gamma})^2 + (1 - z_i)/[1 - \Lambda(\boldsymbol{x}_i^\top \boldsymbol{\gamma})]^2 \right)$.

2.C Additional results

Table 2.10 shows linear estimates of the treatment effects in Tables 2.3 and 2.4. These estimates are given for completeness, though they are inconsistent. Linear regression yields larger coefficients than the weighting method of Frölich and Melly (2013), partly because it does not restrict outcomes to the unit interval and partly because it is more sensitive to limited overlap (Imbens, 2015). For example, the third specification yields 205 fitted values outside the unit interval (3.82 percent of observations); the fourth, 479 (8.93 percent); the fifth, 625 (13.19 percent). Ordinary least squares and inverse probability weighting give similar estimates of the union effect on success rates. Weighted estimates of KOL effects are more stable across specifications than their linear counterparts.

TABLE 2.10: LINEAR ESTIMATES OF THE EFFECTS OF KOL PRESENCE AND UNION SPONSORSHIP

	(1)	(2)	(3)	(4)	(5)
KOL effect on organization (OLS)	0.262***	0.184***	0.111***	0.103***	0.173***
	(0.032)	(0.024)	(0.021)	(0.025)	(0.025)
KOL effect on success rate (OLS)	0.086***	0.078***	0.065***	0.063***	0.087***
	(0.024)	(0.019)	(0.023)	(0.020)	(0.021)
Union effect on success rate (TSLS)	0.318***	0.410***	0.576***	0.704**	0.483***
	(0.096)	(0.106)	(0.183)	(0.303)	(0.128)
Union effect on success rate (OLS)	0.122***	0.100***	0.094***	0.084***	0.087***
	(0.026)	(0.021)	(0.021)	(0.021)	(0.022)
Hausman *m*-statistic	2.307	2.998	2.634	2.054	3.147
Controls					
Period, region & sector		×	×	×	×
Strike characteristics			×		
County characteristics				×	
Past labor conflict					×
Sample					
Sample size	5363	5363	5363	5363	4739
Parameters	2	18	28	26	18

Notes: See Section 2.3 for information about the data. Column (5) excludes observations from 1881 for lack of strike microdata for 1880. Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2). The *m*-statistic tests the difference between estimates of the union effect (Hausman, 1978).
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

As Subsection 2.3.1 notes, the unit of observation is ambiguous in the *Third Report*: each row may represent an entire strike or a subset of the affected establishments. Following the literature, I treat each line as an observation. This approach has the advantage of underweighting outliers. Moreover, unions may strategically strike additional establishments to put pressure on recalcitrant employers; therefore, weighting estimates by establishments could introduce endogeneity bias. Nonetheless, Table 2.11 investigates the sensitivity of my benchmark specification for completeness. For reference, a row may represent up to 1500 establishments, 80 percent represent a single establishment and 99 percent represent 50 or fewer. My findings are qualitatively robust: the union effect is large and downward biased. The coefficients are sensitive to the bound on the weights. Unreported results show that they are more stable if I account for imbalances across periods, regions and sectors. They are then close to the unweighted estimate in Column (2) of Table 2.4, 0.368.

TABLE 2.11: SENSITIVITY OF MAIN ESTIMATES TO WEIGHTING BY ESTABLISHMENTS

	(1)	(2)	(3)	(4)	(5)
KOL effect on organization (IPW)	0.262***	0.324***	0.329***	0.333***	0.364***
	(0.032)	(0.035)	(0.047)	(0.050)	(0.053)
KOL effect on success rate (IPW)	0.086***	0.088***	0.138***	0.145***	0.098
	(0.024)	(0.034)	(0.050)	(0.056)	(0.077)
Union effect on success rate (FM)	0.327***	0.273***	0.418***	0.435***	0.269
	(0.095)	(0.104)	(0.149)	(0.165)	(0.214)
Union effect on success rate (IPW)	0.122***	0.162***	0.219***	0.235***	0.183***
	(0.026)	(0.031)	(0.037)	(0.040)	(0.066)
Fit					
Hausman m-statistic	2.436	1.166	1.389	1.267	0.492
First-stage F-statistic	68.462	84.457	49.439	45.108	46.672
Weighting scheme					
Bound on weights	1	10	50	100	
Sum of weights	5,363	10,550	15,612	17,514	21,593

Notes: See Section 2.3 for information about the data. Standard errors, in parentheses, are robust to correlation across time and space (see Subsection 2.5.2). The m-statistic tests the difference between estimates of the union effect (Hausman, 1978). The first-stage F-statistic tests the effect of the instrument on the treatment in a linear specification.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent. Row labels distinguish estimators: "FM" refers to Frölich and Melly (2013); "IPW", to inverse probability weighting.

3 The Task Content of Occupations

With Francis Kramarz* and Alexis Maitre†

3.1 Introduction

The task approach has attracted considerable attention in labor economics since the seminal work of Autor, Levy and Murnane (2003). Tasks are the building blocks of production. Firms discharge some through machines and contractors. They combine the remainder into jobs, whose content depends on employees' abilities and market conditions. A study of task assignment can thus provide valuable insight into the evolution of labor markets.

Task data are seldom available at the individual level. Therefore, economists have typically examined jobs after they have been grouped into occupational classifications of mostly administrative origin. This approach treats each occupation as a bundle of tasks, the demand for which shifts with such shocks as automation and offshoring. Yet occupations evolve (Autor, Levy and Murnane, 2003; Levy and Murnane, 1996; Spitz-Oener, 2006). Autor (2015) writes: "As the routine cash-handling tasks of bank tellers receded [...], banks recognized the value of tellers [...] as salespersons, forging relationships with customers and introducing them to additional bank services like credit cards, loans, and investment products." This observation suggests that job content is flexible: firms adapt assignments to changes in the relative costs of production

We are indebted to David Autor, Pierre Cahuc, Élise Coudin, David Dorn, Joseph Ferrie, Michael J. Handel, Joel Mokyr, Matthew Notowidigdo and Corinne Prost for advice. All mistakes are ours. Luca Bittarello acknowledges support from the Balzan Foundation and the Center for Economic History at Northwestern University. Francis Kramarz acknowledges support from the ERC Advanced Grant FIRMNET.

* CREST-INSEE: francis.kramarz@ensae.fr. Francis Kramarz acknowledges support from the ERC Advanced Grant FIRMNET.

† Sciences Po: alexis.maitre@sciencespo.fr.

factors, as bank clerks exemplified by assuming more cognitive tasks.[1] In consequence, there is no exact mapping from a job title to a set of tasks (Arntz, Gregory and Zierahn, 2017; Autor and Handel, 2013): today's tellers share few duties with their counterparts from the 1970s, just as they cater to different clients at multinational banks and regional institutions, yet their jobs receive the same occupational code. Unlike tasks, occupations are not a precise economic concept: they are statistical tools, the result of complex algorithms and specific classifications.

This paper explores the relation between job content and market conditions. In particular, we assess the impact of changes in the supply of skilled labor in France from 1991 to 2013. Thanks to public investment in higher education, university graduates increased their share of the workforce from 18 to 36 percent over this period. We exploit individual data from five surveys of work conditions, which allows us to compare jobs within occupations. Table 3.1 in Section 3.2 presents our task measures. Following the literature, we group them into three indexes for analysis: routine, cognitive and social.

Our argument is threefold. First, job content is heterogeneous within occupations (Arntz, Gregory and Zierahn, 2017; Autor and Handel, 2013). For example, consider again bank clerks. As Figure 3.1 shows, there is significant variation in their tasks. If we divide the subsample by the number of tasks in each group, no cell contains more than 13 percent of observations and 90 percent report tasks in multiple categories. Second, university graduates hold a comparative advantage in cognitive work (Acemoglu and Autor, 2011; Spitz-Oener, 2006). Third, higher average educational attainment increased the supply of cognitive tasks and reduced their relative price, so workers spent more time on routine tasks instead. Figure 3.1 illustrates this shift: bank clerks were given more routine and fewer cognitive tasks as the share of university graduates rose from 14 percent of tellers in 1991 to 58 in 2013.

Section 3.4 formalizes these ideas into a model. Workers supply one unit of labor, which they

[1] We use "job content" as a synonym for workers' tasks throughout the paper.

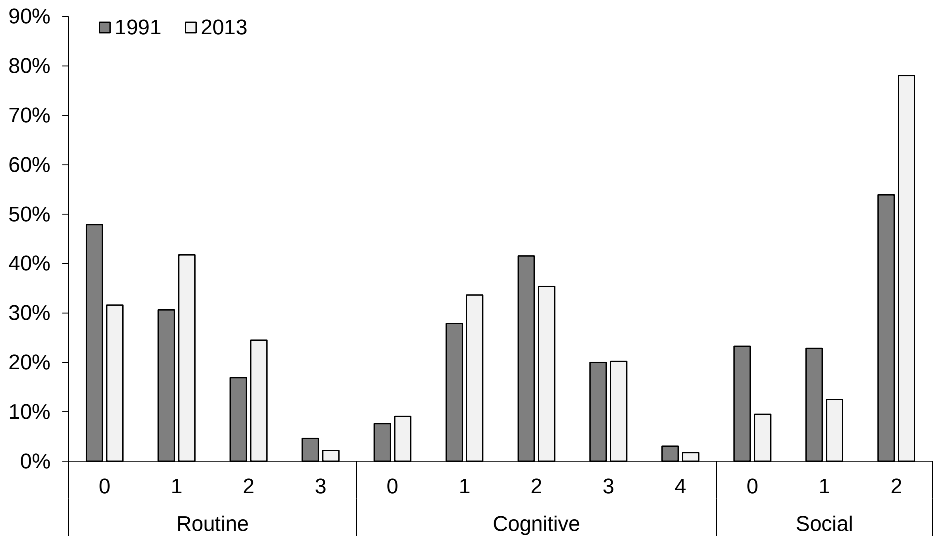

Figure 3.1: Distribution of the sum of task indicators for bank clerks by year

Notes: Section 3.2 discusses task categories. *Sources:* Authors' calculations, based on the Work Conditions Survey by INSEE and DARES.

share between a routine and a cognitive task. Skilled workers hold a comparative advantage in the cognitive task. Unlike Autor, Levy and Murnane (2003) or Acemoglu and Autor (2011), we assume that workers do not specialize. Firms combine tasks into output. The model predicts two effects from an increase in the skill supply. Because skilled workers perform more cognitive work than the unskilled, a composition effect raises the cognitive content of aggregate output. On the other hand, a substitution effect obtains at the individual level: as the relative price of cognitive tasks decreases, each worker supplies more routine and fewer cognitive tasks.

To test these predictions, we regress our task indexes on the share of university graduates within each labor market by year. We assume that a separate labor market exists for each occupation within each region of France. Since schooling, migration and labor supply are endogenous, we forecast the graduate share on the basis of previous surveys to construct instruments. The first is the graduate share among workers who will still be under the minimum retirement age by the next survey. The second supposes that the contingents of each skill group in each labor market will evolve at the national rate between surveys (Bartik, 1991). We define both instruments in terms of birth regions rather than region of residence on account of migration. The exclusion restriction assumes that temporary local shocks are orthogonal to the initial distribution of graduate shares across labor markets (Goldsmith-Pinkham, Sorkin and Swift, 2018).

Our first specification includes fixed effects for occupation, region and year. Therefore, we obtain identification from variation in task assignment across workers within each labor market and the coefficients inform us about the presence of a substitution effect at the individual level. The second does not include occupation effects. The resulting coefficients combine variation in job content at the individual level with variation between occupations (hence, labor markets). As a consequence, they inform us about the existence of a composition effect.

Our results are twofold. For a given occupation, a higher graduate share is associated with more routine, fewer cognitive and fewer social tasks. The opposite pattern holds across occupations: the average job involves fewer routine, more cognitive and more social tasks. Hence, we find evidence for both theoretical predictions of an individual substitution effect and an aggregate composition effect. The estimates are significant but modest: the task indexes shift by 2 to 16 percent of a standard deviation for a rise in the graduate share of 10 percentage points around the nationwide share in 1990.

We examine task compensation as well. We show that an increase in the routine index lowers hourly wages by 0.6 to 1 percent, an increase in the cognitive index of one standard deviation raises them by 0.9 to 2.1 percent and an increase in the social index lowers them by 0.4 to 1.33 percent. The wage effects of routine and cognitive tasks decrease in magnitude between surveys, which is also consistent with the model. Our estimates are similar to Autor and Handel's (2013), though theirs are based on different measures of job content and American data.

The task literature has greatly improved our understanding of labor markets. For example, Autor, Levy and Murnane (2003) argue that computers replaced labor in routine activities, raising the cognitive content of occupations and reshaping the occupational structure within industries. Similar analyses have shed light on automation (Agrawal, Gans and Goldfarb, 2019; Atack, Margo and Rhode, 2019; Gregory, Salomons and Zierahn, 2019; Spitz-Oener, 2006), employment polarization (Acemoglu and Autor, 2011; Autor and Dorn, 2013; Firpo, Fortin and Lemieux, 2011), gender gaps (Black and Spitz-Oener, 2010), immigration (Peri and Sparber, 2009), mobility (Gathmann and Schönberg, 2010), offshoring (Blinder, 2009; Jensen and Kletzer, 2010), part-time work (Elsayed, de Grip and Fourge, 2017), social skills (Deming, 2017), trade (Autor, Dorn and Hanson, 2015) and more. This paper shows that market conditions influence task assignment within occupations. This finding highlights the need for nuance in discussing the future of work (Acemoglu and Restrepo, 2018, 2019; Arntz, Gregory and Zierahn, 2017). It does not suffice to examine the typical tasks in an occupation at present to forecast its susceptibility to automation or outsourcing. As we noted earlier, occupations evolve: workers may perform unautomated tasks more intensively, firms may develop new tasks for idle employees, etc. Rising educational attainment may facilitate this adjustment by preparing workers for lifelong learning and flexible roles.

The paper continues as follows. Section 3.2 presents the data. Section 3.3 discusses stylized facts. Section 3.4 introduces the model. Section 3.5 describes our empirical approach. Section 3.6 contains the results. Section 3.7 concludes.

3.2 Data

This section describes our data. Our sources are the French Labor Force Survey (*Enquête Emploi*, LFS), the Work Conditions Survey (*Enquête Conditions de Travail*, WCS) and the Work Organization Survey (*Enquête Techniques et Organisation du Travail*, WOS).

3.2.1 The Labor Force Survey

The National Institute of Statistics and Economic Studies (*Institut national de la statistique et des études économiques*, INSEE) developed the Labor Force Survey in 1950 in an effort to measure employment between census years (Goux, 2003). It was mostly yearly until 2002. It averaged 146 000 respondents per year between 1990 and 2002, renewed by thirds. Data collection became continuous in 2003. Results are quarterly. The sample averaged 71 500 respondents per quarter between 2002 and 2008, renewed by sixths. It increased to an average of 104 000 respondents per quarter between 2010 and 2012.

The LFS collects information about workers' characteristics, their jobs and their households. We construct the following covariates for the empirical analysis: female; married; foreign born; age and age squared; tenure and tenure squared; multiple jobs; part-time job; fixed-term contract; and civil servant. Except for age and tenure, all covariates are binary indicators. Furthermore, we include fixed effects for education level,[2] occupation and region of residence.[3] We use two-digit occupations, since INSEE changed the four-digit classification in 2003. We do not include industry effects because the classification changed in 1993 and 2008. Other than covariates, the LFS gives us the share of university graduates by year, region and occupation. Because certain cells are small, we pool observations across three years at a time for additional precision. For example, we estimate the graduate share in 1991 with data from 1990–92.

The LFS gathers data about monthly wages after social charges. A third of respondents provide intervals instead of precise numbers. A small percentage refuses to answer at all (less than three percent of wage workers). INSEE imputes wages for these observations. The resulting distribution is similar to the distribution across the *Déclarations annuelles de données sociales* (the reference for French wage data). Because of the reduction of the workweek from 39 to 35

[2] We use five education levels: less than middle school, middle school, high school, college and postgraduate.

[3] Because the sample contains few observations from Corsica, we merge it into Provence. In constructing the instrument, we use regions of birth. We create a synthetic region for the foreign born.

hours between 1999 and 2002, monthly wages are not directly comparable across years. For this reason, we use weekly hours to construct hourly wages.[4] We truncate hours at the legal limit (60 hours per week). We also adjust them if the employer extended holidays in lieu of shortening the workweek. If the respondent reported an interval, we use its half point. If they did not answer at all, we use median hours by occupation and part-time status.

We restrict the sample to wage workers by excluding interns, apprentices, artisans, agricultural workers, the self-employed, business owners and the clergy. Wage regressions exclude workers whose hourly wages are smaller than four fifths of the minimum wage and outliers.[5] We use sampling weights throughout the paper. We normalize the sum of weights across the final sample of each year to unity.

3.2.2 The Work Conditions Survey and the Work Organization Survey

INSEE conducted its first WCS in 1984. A supplementary survey of the outgoing group of the LFS, it enquired into sources of stress at work, whether physical (e.g., loud noises) or psychological (e.g., interacting with the public). INSEE repeated the exercise in 1991, 1998 and 2005. The WOS was a similar supplement to the LFS, focused on job content and the organization of work. It was undertaken in 1987 and 1993. The WCS and WOS averaged 20 000 respondents per wave. The Directorate for Research, Studies and Statistics at the Labor Ministry (*Direction de l'animation de la recherche, des études et des statistiques*, DARES) took responsibility over the WCS in 2013. It became an independent survey and involved 33 673 respondents in its first wave.

Researchers have often drawn task data from two sources from the U.S.: the Dictionary of

---

[4] The yearly survey collected information about regular weekly hours. The quarterly survey has distinguished between contractual hours and regular hours. We use contractual hours whenever they are available.

[5] Following Crépon and Gianella (1999), outliers are observations for which $|\hat{u}_i| > 5 \times (q_{75} - q_{25})$, where $\hat{u}_i$ is the residual from a linear regression of log hourly wages and q_x is x-th centile of residuals. The regression uses data from the LFS between 1990 and 2012. Years are given equal weight.

Occupational Titles and the O*NET. Both files provide scores for a large number of occupations in terms of activities, aptitudes and requirements (Autor, Levy and Murnane, 2003; Jensen and Kletzer, 2010). The WCS and WOS offer a significant advantage over these data: access to individual responses. As a consequence, we can explore heterogeneity within occupations and the joint task distribution across workers.[6]

We measure job content along three dimensions: routine, cognitive and social. We borrow this approach from the extensive literature on automation and offshoring (Autor, Levy and Murnane, 2003; Handel, 2012; Jensen and Kletzer, 2010). Following Spitz-Oener (2006), we construct task indexes by selecting relevant variables from the WCS and the WOS, transforming them into indicators and averaging the indicators. We selected variables on three criteria: they unambiguously pertain to one of our three categories, they are available across all years and the underlying questions are identically phrased across surveys. Note that they surveys report respondents' original answers and interviewers did not help them interpret the questions.

Table 3.1 shows the means of each indicator by category and year. Routine tasks denote a lack of autonomy or initiative. Cognitive tasks involve decision making.[7] Social tasks require interaction with clients or the public.[8] Note that we limit the sample to the period from 1991 to 2013. We discard the 1984 WCS and the 1987 WOS because the LFS did not contain all of the variables of interest at the time. See Section 3.3 for further discussion.

We also construct an indicator of computer usage from the WCS and the WOS. We use it as a control variable to account for the influence of technological shocks on the skill supply and the

[6] Similar samples are available for Germany: see Spitz-Oener (2006). For cross sections, see Arntz, Gregory and Zierahn (2017) and Autor and Handel (2013).

[7] Note that our cognitive tasks capture autonomy (e.g., handling incidents) rather than intellectual difficulty (e.g., complex calculations). If companies rarely update their hierarchies, these measures may exhibit inertia, which could partly explain the lack of an increase in the cognitive score over the sample period (cf. Table 3.1). Note furthermore that the growth in the skill supply is due to young workers, whose typical job involves less decision making than seniors'.

[8] Our measures of social tasks include the fact that external demands determine one's work rhythm, since it indicates that workers and clients interacted.

TABLE 3.1: TASK MEASURES BY CATEGORY

	Occupation with highest incidence	Percentage of positive responses					
		1991	1993	1998	2005	2013	All
Routine tasks							
Production norms to be fulfilled within the day	Drivers	38.0	42.6	43.0	42.1	45.9	42.3
Repetitive movements	Unskilled manufacturing workers	29.6	24.5	28.7	27.9	41.2	30.4
Work rhythm determined by machinery	Unskilled manufacturing workers	12.8	11.5	13.6	13.9	18.0	14.0
Cognitive tasks							
Choosing strategy to achieve goals	Senior technicians in the private sector	83.5	83.9	86.9	81.4	80.2	83.2
Departing from deadlines	Senior technicians in the private sector	35.7	37.1	36.2	36.6	34.3	36.0
Departing from instructions	Professionals in arts and culture	24.6	22.4	28.0	30.3	28.2	26.7
Handling incidents	Senior managers in the private sector	50.2	53.4	56.6	52.1	50.7	52.6
Social tasks							
Contact with the public	Sales workers	60.7	61.3	62.4	68.5	70.9	64.8
Work rhythm determined by external demands	Sales workers	45.9	45.0	54.3	53.5	58.0	51.4

demand for tasks.

3.3 Stylized facts

This section presents stylized facts about the French labor market.

Figure 3.2 illustrates the increase in educational attainment in France since 1990. University graduates constituted 36 percent of the employed workforce in 2012, up from 18 percent. The proportion of workers with secondary degrees rose from 12 to 19 percent in this period. By contrast, the fraction of workers without degrees fell from 41 to 20 percent. This upskilling process largely due to sustained public investment in higher education. As education minister under President François Mitterrand in the mid 1980s, Jean-Pierre Chevènement initiated an effort to raise the high-school graduation rate to 80 percent. Modernization plans for tertiary education followed in 1990 (*Université 2000*) and 1999 (*Université du troisième millénaire*), which included the creation of eight universities and dozens of technical colleges.

Table 3.2 summarizes the evolution of job content in this period. Workers performed 1.05 out of three possible routine tasks on average in 2013, up from 0.8 in 1991; in consequence, our index increased by thirty percent between 1991 and 2013. We observe a smaller rise in the incidence of social tasks, 21 percent. Cognitive tasks remained stable. These trends may seem surprising, but they are broadly consistent with cross-country evidence from the European Working Conditions Survey in Handel (2012). Moreover, there is little evidence of polarization in the French labor market (Van Reenen, 2011; Verdugo, 2014; Verdugo, Fraisse and Horny, 2012)[9] and the impact of computerization was limited (Card, Kramarz and Lemieux, 1999; Goux and Maurin, 2000). Hence, it is plausible to find an increase in routine tasks and no change in cognitive tasks.

Table 3.2 also reveals that covariates do not provide much insight into the distribution of job content. It shows the coefficient of determination from yearly linear regressions of task indexes

[9] See Bozio, Breda and Guillot (2016) for evidence of polarization in terms of labor costs.

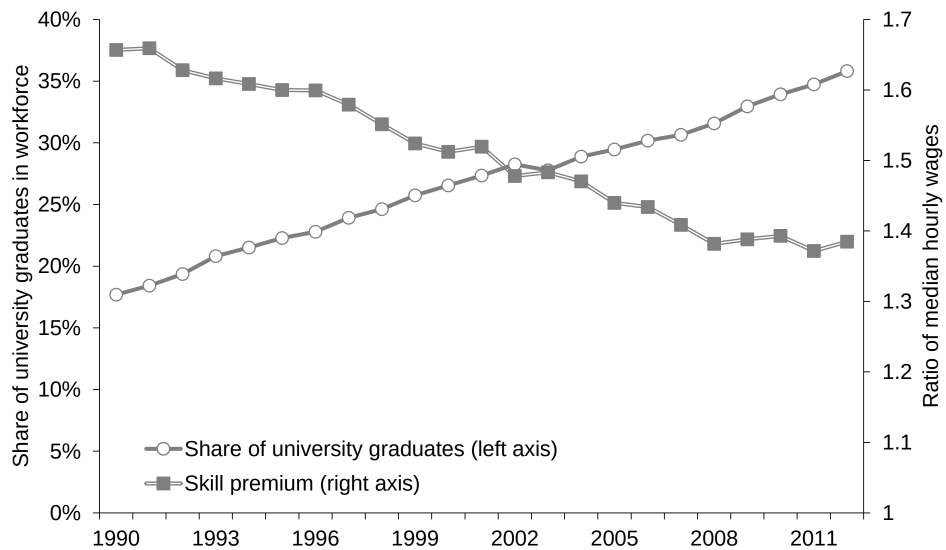

Figure 3.2: Share of university graduates and skill premium by year

Notes: The sample consist of employed wage workers. The skill premium is the ratio of median hourly wages of university graduates and less educated workers. *Sources:* Authors' calculations, based on the Labor Force Survey by INSEE.

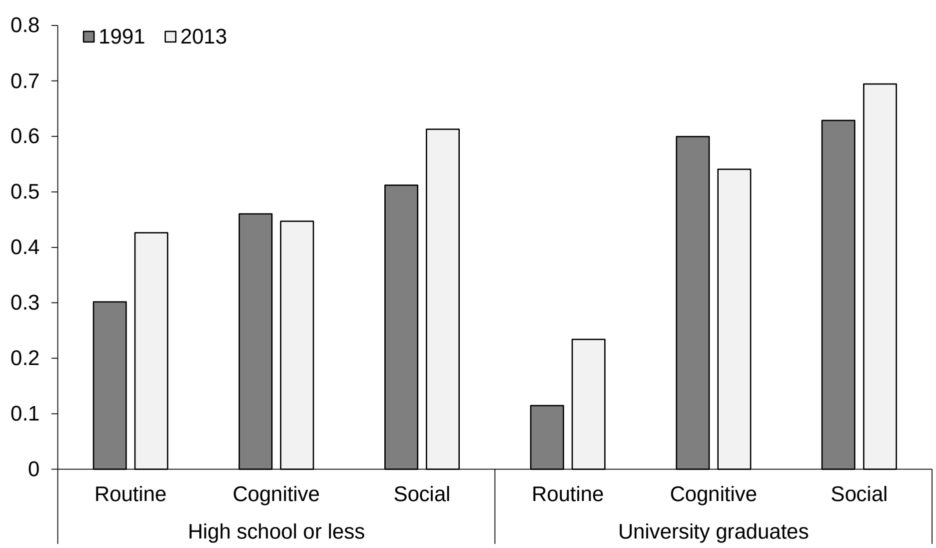

Figure 3.3: Change in task assignment by education level

Notes: The figure shows the average of task indicators in each group by education level and year (q.v. Section 3.2). *Sources:* Authors' calculations, based on the Work Conditions Survey by INSEE and DARES.

TABLE 3.2: VARIATION IN TASK ASSIGNMENT

	Routine tasks			Cognitive tasks			Social tasks		
	Mean	St. dev.	R^2	Mean	St. dev.	R^2	Mean	St. dev.	R^2
1991	0.268	0.305	0.228	0.485	0.272	0.200	0.533	0.415	0.246
1993	0.262	0.300	0.231	0.492	0.257	0.178	0.532	0.405	0.180
1998	0.284	0.309	0.243	0.519	0.259	0.164	0.584	0.405	0.196
2005	0.280	0.309	0.203	0.501	0.267	0.119	0.610	0.393	0.154
2013	0.350	0.327	0.217	0.484	0.265	0.151	0.645	0.377	0.175
All	0.289	0.312	0.223	0.496	0.264	0.156	0.581	0.402	0.193

Notes: The table shows summary statistics for the average of task indicators in each group (q.v. Section 3.2). The R^2 refers to a linear regression on individual characteristics and fixed effects (education, occupation, region and year).

on individual characteristics and fixed effects for education level, occupation and region. This model explains a quarter of the variation in task assignment at most – further evidence that job content is heterogeneous within occupations (Arntz, Gregory and Zierahn, 2017; Autor and Handel, 2013).

As Figure 3.2 shows, the skill premium shrank throughout the period: university graduates' median hourly wage was 38 percent larger than other workers' in 2012, down from 66 percent in 1990. Wage inequality fell as a result (Charnoz, Coudin and Gaini, 2011; Verdugo, 2014; Verdugo, Fraisse and Horny, 2012). Table 3.3 displays the growth in each decile of hourly wages between 1991 and 2005.[10] The ratio of the ninth to the first decile decreased by 8.3 percent. As the table shows, differences in job content between wage deciles diminished too. Routine tasks are more frequent in the bottom of the wage distribution; however, the ratio of the average routine score in the ninth wage decile to the average in the first rose by 18 percent between 1991 and 2005. Cognitive and social tasks are more common in the top of the distribution, but the ratio of cognitive scores decreased by 15 percent and that of social scores, by 10 percent.

Figure 3.3 presents average task indexes by education level in 1991 and 2013. As Spitz-Oener (2006) notes, university graduates perform fewer routine, more cognitive and more social

[10] The table excludes the 2013 WCS because it did not collect comparable wage data.

TABLE 3.3: TASK ASSIGNMENT AND HOURLY WAGES BY DECILE OF HOURLY WAGES

	Hourly wages		Routine tasks		Cognitive tasks		Social tasks	
	Level, 1991	Change, 91–05	Level, 1991	Change, 91–05	Level, 1991	Change, 91–05	Level, 1991	Change, 91–05
1	6.625	0.215	0.336	–0.011	0.396	0.123	0.482	0.220
2	7.494	0.193	0.327	–0.002	0.409	0.109	0.495	0.202
3	8.351	0.181	0.317	0.008	0.424	0.092	0.507	0.187
4	9.207	0.168	0.302	0.028	0.445	0.065	0.520	0.169
5	10.192	0.151	0.287	0.042	0.464	0.049	0.531	0.156
6	11.246	0.154	0.267	0.060	0.489	0.033	0.544	0.142
7	12.712	0.154	0.240	0.084	0.521	0.016	0.559	0.125
8	14.988	0.142	0.202	0.117	0.562	–0.008	0.573	0.109
9	19.485	0.114	0.151	0.171	0.627	–0.047	0.583	0.094

Notes: Wages are shown in constant euros (base 2015). The table shows average task indexes within each wage decile and the proportional change in averages from 1991 to 2005.
Sources: Authors' calculations, based on the Labor Force Survey and the Work Conditions Survey by INSEE.

tasks than other workers. However, these patterns weakened over time. Routine and social activities increased in incidence in both groups, but the change was larger among graduates in proportional terms. On the other hand, graduates discharged fewer cognitive tasks in 2013 than 1991.[11] As we discuss in the following sections, upskilling may partly explain these changes: higher educational attainment may have reduced the relative price of cognitive tasks, leading workers to spend more time on routine tasks instead.

We conclude this section with an overview of economic conditions in the 1990s and 2000s. Growth was slow and unsteady. Real GDP per capita expanded at an average yearly rate of 1.8 percent between 1990 and 2013. There were four recessions in this period (in 1992, 2001, 2008 and 2012). Unemployment was persistently high, averaging 8.8 percent, and the shares of both fixed-term contracts and part-time jobs increased (from 6 to 11 percent and from 11 to 17 percent, respectively). These decades are also noteworthy for the reduction of the legal

[11] Spitz-Oener (2006) finds different patterns in Germany. Nonroutine tasks became more common at all education levels, but the proportional change was larger for the uneducated. Routine manual tasks exhibit a larger decrease for the uneducated as well. On the other hand, she observes a larger cut in routine cognitive tasks for the educated. It is unclear whether these discrepancies are due to fundamentals or differences in task indexes.

workweek from 39 to 35 hours between 2000 and 2002, which inflated hourly wages and compressed their distribution (Aeberhardt, Givord and Marbot, 2016).

3.4 Theoretical framework

This section develops a simple theoretical framework for our empirical analysis of the interaction between the supply of skilled workers and task assignment. We adapt the model by Peri and Sparber (2009).

3.4.1 Task demand

Consider an economy in autarchy. A representative firm combines tasks into a consumption good (y). Tasks may be routine (r) or cognitive (c). For simplicity, we assume that production does not require capital. The production technology is:

$$y = \left[r^{(\sigma-1)/\sigma} + c^{(\sigma-1)/\sigma}\right]^{\sigma/(\sigma-1)},$$

where σ controls the elasticity of substitution between inputs (n.b. $\sigma > 0$).[12] Production does not require both tasks (unless $\sigma \to 1$), but this functional form implies that the firm will always mix them in equilibrium, as we observe in the data (cf. Section 3.2).

The firm purchases task services on frictionless labor markets. The consumption good is the numeraire. Therefore, profits are: $y - w_r r - w_c c$, where w_r is the price of a unit of routine tasks (analogously for w_c). By combining the necessary conditions for profit maximization, we find the relative demand for tasks:

$$\frac{r^*(\omega)}{c^*(\omega)} = \frac{1}{\omega^\sigma}, \tag{3.1}$$

[12] For a discussion of this production function, see Acemoglu and Autor (2011).

where ω is the price ratio: $\omega \equiv w_r/w_c$. The firm decreases the routine content of production in response to an increase in the cost of routine tasks. This inverse relationship is important for our results, though its precise functional form is not.

3.4.2 Task supply

The economy comprises a measure p of skilled workers ($s = 1$) and a measure $1 - p$ of unskilled workers ($s = 0$). Each worker is endowed with one unit of labor. They do not derive utility from leisure. Hence, they apportion x_s of their time to the supply of r_s in routine tasks and $1 - x_s$ to the supply of c_s in cognitive tasks. The resulting task supply is:

$$r_s = \alpha_{rs} x_s^{\beta} \quad \text{and} \quad c_s = \alpha_{cs}(1 - x_s)^{\beta}, \tag{3.2}$$

where $\beta \in (0, 1)$, $\alpha_{rs} > 0$ and $\alpha_{cs} > 0$. The curvature parameter β implies that workers become less productive as they repeat tasks, which may reflect technical limitations (e.g., fatigue) or a preference for variety at work. The scale parameters α_{rs} and α_{cs} determine total productivity. We assume that skilled workers enjoy a relative advantage at cognitive tasks: $\alpha_{c1}/\alpha_{r1} > \alpha_{c0}/\alpha_{r0}$.

Because savings bear no interest and the model is static, workers do not save. Therefore, they maximize utility by maximizing their income,

$$w_r \alpha_{rs} x_s^{\beta} + w_c \alpha_{cs}(1 - x_s)^{\beta},$$

through the choice of x_s. The optimal allocation ($x_s^*(\omega)$) satisfies:

$$\frac{x_s^*(\omega)}{1 - x_s^*(\omega)} = \left(\frac{\alpha_{rs}\omega}{\alpha_{cs}}\right)^{\frac{1}{1-\beta}}, \tag{3.3}$$

where $x_s^*(\cdot)$ is an increasing function.

Equation (3.3) has three implications. First, the supply of routine tasks increases with their relative price. Second, unskilled workers perform more routine tasks than the skilled: $r_0 > r_1$; conversely, $c_1 > c_0$. This property is a consequence of their relative advantages and finds support in the data (cf. Figure 3.3). Third, workers do not specialize: $r_s > 0$ and $c_s > 0$ for all s. This feature is due to the nonlinearity in the task supply (q.v. equation (3.2)). Empirical evidence support it too: fewer than 15 percent of workers in our sample report tasks in a single category. Our model thus differs from Acemoglu and Autor's (2011) or Autor, Levy and Murnane's (2003), where the task supply is linear and workers perform one task each.[13]

3.4.3 Equilibrium and comparative statics

Equilibrium obtains when prices, w_r and w_c, ensure that each task market and the goods market clear:

$$
\begin{aligned}
r &= (1-p)r_0 + pr_1, \\
c &= (1-p)c_0 + pc_1, \\
y &= w_r[(1-p)r_0 + pr_1] + w_c[(1-p)c_0 + pc_1].
\end{aligned}
$$

We can find the equilibrium in two steps. Equations (3.1) and (3.3) fix relative prices. We can then determine absolute prices by clearing the goods markets.

This paper investigates the impact of changes in the supply of skilled labor on task assignment. The model has implications for our empirical analysis. To see this, first combine equations (3.1)

[13] Deming (2017) proposes a model of partial specialization, in which each worker performs a subset of all tasks and outsources the remainder. The mechanism is different from ours: workers do not specialize because of trade costs in his model, whereas we assume that the task supply exhibits decreasing returns.

and (3.3):

$$\frac{(1-p)\alpha_{r0}x_0^*(\omega)^\beta + p\alpha_{r1}x_1^*(\omega)^\beta}{(1-p)\alpha_{c0}[1-x_0^*(\omega)]^\beta + p\alpha_{c1}[1-x_1^*(\omega)]^\beta} = \frac{1}{\omega^\sigma}.$$

Implicit differentiation then reveals that the price ratio, ω, is an increasing function of the share of skilled workers, p.[14] For a fixed ω, a higher p induces an expansion in the aggregate supply of cognitive tasks, since skilled workers spend more time on cognitive tasks than the unskilled. Equilibrium requires that cognitive tasks become relatively cheaper – i.e. ω must go up.[15] As the price ratio rises, each worker supplies more routine and fewer cognitive tasks, while the firm demands more cognitive and fewer routine tasks. Therefore, the new equilibrium differs in two aspects. There is a substitution effect toward the routine at the worker level by equation (3.3). Nonetheless, the cognitive content of aggregate output increases through a composition effect. These effects are generally nonlinear in both the initial share of skilled workers and the magnitude of the shift in p between equilibria.

3.4.4 Discussion

To develop intuition, it is useful to adopt the firm's perspective. Because skilled workers perform more cognitive tasks than the unskilled, upskilling implies an oversupply of cognitive tasks at constant prices, so their relative price comes down. Hence, the firm has an incentive to use more cognitive tasks. However, its employees want fewer cognitive tasks, since they now pay less. Therefore, the firm replaces unskilled workers with skilled ones: although each skilled worker discharges fewer cognitive tasks than before, they still perform more cognitive tasks than unskilled workers in the former equilibrium, so the cognitive content of production increases.

We have followed the literature in assuming that there are distinct markets for each task

[14] The proof is available from the authors upon request.
[15] Note that an increase in ω entails a falling skill premium (in accordance with Figure 3.2).

and that workers control task supply. We could rewrite the model so that markets separate by skill instead and firms assign tasks to their employees. Although the exposition is more cumbersome, our results go through: an increase in the skill supply would again shrink the skill premium, generating a substitution effect at the worker level and a countervailing composition effect at the aggregate level.

The model assumes that skill is binary and exogenous. A more complex setup could instead treat education as an endogenous function of one's aptitude for cognitive tasks. As the graduate share increased, the marginal graduate would have an ever smaller comparative advantage in cognitive tasks in this framework, reinforcing the substitution and composition effects. We ignore this mechanism because we can not measure innate ability in the data. In interpreting our results, one should nevertheless keep in mind that the quantity of graduates might affect their quality.

The model is also silent on the role of capital and technology. A growing literature analyzes automation (Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2019; Atack, Margo and Rhode, 2019), computerization (Autor and Dorn, 2013; Autor, Levy and Murnane, 2003; Spitz-Oener, 2006), artificial intelligence (Agrawal, Gans and Goldfarb, 2019) and like shocks. Innovation could both increase the skill supply through endogenous schooling and distort the demand for tasks. Other than parsimony, we leave capital out for two reasons.[16] First, previous studies have only found a limited impact of computerization on the French labor market (Card, Kramarz and Lemieux, 1999; Goux and Maurin, 2000). Second, technological shocks are likely to work against us. As Spitz-Oener (2006) argues, computerization decreases the demand for routine tasks and increases the demand for skilled labor. If schooling is endogenous, we should therefore observe a negative correlation between routine tasks and the skill supply, whereas the model predicts the opposite relationship within occupations.

[16] Recall that we include computer usage as a control in the empirical analysis.

3.5 Empirical approach

The model implies that an increase in the skill supply should correlate with more routine and fewer cognitive tasks at the individual level (the substitution effect) but with fewer routine and more cognitive tasks at the aggregate level (because of a composition effect). To test the first prediction, we use variation in job content across individual workers in each labor market. To test the second, we use variation between labor markets. Although social tasks are not part of the model, we include them in the empirical analysis in light of the rising importance of social skills in modern labor markets (Deming, 2017).

We assume that workers segregate into distinct labor markets by administrative region, two-digit occupation and year. Our analysis encompasses 2415 markets by this definition ($21 \times 23 \times 5$).[17] We use the share of university graduates as a proxy for the skill supply.

Consider observation i in occupation o_i, region r_i and year t_i. Write y_i for i's individual task score, $p_{o_i r_i t_i}$ for the graduate share in their labor market, $\boldsymbol{x}_i$ for a vector of individual characteristics (q.v. Section 3.2) and u_i for the residual. We consider two specifications. The first includes fixed effects for occupation, region and year:

$$y_i = \gamma_1 p_{o_i r_i t_i} + \gamma_2 p^2_{o_i r_i t_i} + \boldsymbol{\beta} \cdot \boldsymbol{x}_i + \alpha_{o_i} + \alpha_{r_i} + \alpha_{t_i} + \alpha + u_i. \tag{3.4}$$

We are interested in γ_1 and γ_2. The quadratic term helps us take nonlinearities into account (cf. Subsection 3.4.3). Because it includes a full set of fixed effects, this regression only exploits variation in task assignment within labor markets for identification; hence, it assesses the extent of substitution between tasks at the individual level. The second specification leaves occupation

[17] In French terminology, we define markets in terms of *régions* and *catégories socioprofessionnelles*.

effects out:

$$y_i = \gamma_1' p_{o_i r_i t_i} + \gamma_2' p^2_{o_i r_i t_i} + \boldsymbol{\beta}' \cdot \boldsymbol{x}_i + \alpha'_{r_i} + \alpha'_{t_i} + \alpha' + u'_i. \tag{3.5}$$

This regression mixes variation at the individual level and variation between occupations (hence, between labor markets), so it provides insight into the task content of aggregate output. So far as γ_1' and γ_2' differ from γ_1 and γ_2, it informs us about the existence of a composition effect at an aggregate level.

Ordinary least squares need not yield consistent estimates of γ_1, γ_2, γ_1' and γ_2'. One concern is measurement error, given that we estimate the graduate share from the LFS. Although we pool observations across three years for additional precision (see Section 3.2), we may still lack power for some occupations in less populated regions. Another worry is endogeneity in schooling decisions, migration and workforce participation. For example, a local technological shock could affect both the relative demand for cognitive tasks (because routine tasks are automated, say) and the skill supply (because skilled workers immigrate from other regions, say).

Therefore, we use instrumental variables for identification. We construct two instruments by projecting the graduate share on the basis of earlier waves of the LFS.[18] For concreteness, consider observation i from 1991. We use three surveys to compute the instruments for i: 1983, 1984 and 1985.[19] The first instrument exploits retirements. It is the graduate share among such workers as were born in the same region as i, have the same occupation as i and will not reach the minimum retirement age by 1991 (viz. 60 years). This definition helps us cancel the effect of endogenous education (by dropping incoming cohorts) and migration (by using birth regions). It does not involve actual retirements, lest we introduce bias from participation decisions. This instrument is relevant because retirements boosted the share of graduates in

[18] Because the regression equations are quadratic, we use each instrument as described and its square.

[19] For 1993, we use data from 1983–85 as well. For 1998, we use data from 1990–92. For 2005, we use data from 1997–99. For 2013, we use data from 2004–06. We update the base year for symmetry.

this period by removing less educated cohorts from the labor force. The second is a Bartik instrument (Bartik, 1991). Consider such workers as were born in the same region and have the same occupation as i. To construct the instrument, we multiply the contingent of each skill group within this population by the corresponding growth rate across the entire workforce and compute the implied graduate share. In computing the national growth rates, we exclude i's birth region and occupation. The exclusion restriction requires the same assumption for both instruments: the initial distribution of graduate shares across labor markets must be orthogonal to local shocks (Goldsmith-Pinkham, Sorkin and Swift, 2018).

Table 3.4 shows coefficients from linear regressions of the graduate share on the instruments. Both are highly correlated with the graduate share. The Bartik instrument is slightly stronger, perhaps because it takes the average education of incoming cohorts into account. Covariates reduce the coefficients, but they remain significant at any conventional level.[20]

3.6 Results

3.6.1 The impact of the skill supply on task assignment

Table 3.5 displays our main results: estimates of the impact of changes in the graduate share on job content. Each column presents one combination of task index and covariates. The first two columns show coefficients from regressions of the routine score; columns (3) and (4), of the cognitive score; columns (5) and (6), of the social score; the last two columns, of the ratio between the routine score and the sum of the three scores.[21] Odd columns display coefficients from regressions with fixed effects for occupation; even columns, without them. Each column contains estimates by both two-stage and ordinary least squares. The table shows separate

[20] Even columns use fewer observations than odd columns because we do not observe tenure for all observations. Our estimates are robust to dropping incomplete observations altogether or imputing tenure.

[21] Columns (7) and (8) exclude 661 observations for which the sum of task indexes is zero.

TABLE 3.4: LINEAR REGRESSION OF THE GRADUATE SHARE ON INSTRUMENTS

	(1)	(2)	(3)	(4)
Retirement instrument	1.000***	0.418***		
	(0.006)	(0.022)		
Bartik instrument			0.960***	0.476***
			(0.005)	(0.022)
Partial F-statistic	3.2×10^4	3.5×10^2	3.5×10^4	4.9×10^2
Controls				
Individual characteristics		16		16
Occupation fixed effects		22		22
Region fixed effects		20		20
Year fixed effects		4		4
Fit				
Observations	94,990	94,253	94,990	94,253
Adjusted R^2	0.963	0.980	0.972	0.982

Notes: The table shows coefficients from ordinary linear regressions. Standard errors (in parentheses) are clustered by occupation, region and year. The outcome is the share of university graduates by occupation, region and year. Both instruments are a projection of the share of graduates by birth region and occupation. The retirement instrument supposes that the graduate share will only evolve between surveys because of retirements. The Bartik instrument supposes that the contingents of graduates and nongraduates will evolve at the national rate in each local market between surveys. See Section 3.5 for further detail. The partial F-statistic tests the joint significance of the instruments.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

TABLE 3.5: IMPACT OF CHANGES IN THE GRADUATE SHARE ON TASK ASSIGNMENT

	Routine tasks		Cognitive tasks		Social tasks		Routine tasks (share)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Retirement instrument								
Share of university graduates	0.207***	–0.676***	–0.043	0.416***	0.040	0.742***	0.052	–0.647***
by occupation, region and year	(0.075)	(0.032)	(0.086)	(0.028)	(0.122)	(0.054)	(0.048)	(0.029)
Squared share of university graduates	–0.433***	0.448***	–0.439***	–0.254***	–0.404***	–0.720***	0.028	0.487***
by occupation, region and year	(0.090)	(0.036)	(0.123)	(0.032)	(0.161)	(0.062)	(0.056)	(0.032)
Bartik instrument								
Share of university graduates	0.136*	–0.735***	–0.100	0.447***	0.018	0.783***	0.043	–0.693***
by occupation, region and year	(0.079)	(0.033)	(0.078)	(0.027)	(0.119)	(0.054)	(0.051)	(0.031)
Squared share of university graduates	–0.334***	0.521***	–0.338***	–0.292***	–0.459***	–0.768***	–0.044	0.543***
by occupation, region and year	(0.084)	(0.037)	(0.104)	(0.032)	(0.151)	(0.062)	(0.054)	(0.034)
No instrument								
Share of university graduates	0.026	–0.729***	–0.061	0.404***	–0.004	0.810***	0.004	–0.686***
by occupation, region and year	(0.054)	(0.031)	(0.045)	(0.025)	(0.082)	(0.049)	(0.033)	(0.028)
Squared share of university graduates	–0.132***	0.512***	–0.046	–0.239***	–0.292***	–0.797***	–0.009	0.533***
by occupation, region and year	(0.054)	(0.034)	(0.044)	(0.029)	(0.084)	(0.057)	(0.031)	(0.031)
Controls								
Individual characteristics	16	16	16	16	16	16	16	16
Occupation fixed effects	22		22		22		22	
Region fixed effects	20	20	20	20	20	20	20	20
Year fixed effects	4	4	4	4	4	4	4	4
Fit								
Observations	94,253	94,253	94,253	94,253	94,253	94,253	93,592	93,592
Adjusted R^2 (OLS)	0.219	0.144	0.156	0.114	0.193	0.085	0.301	0.172

Notes: Standard errors (in parentheses) are clustered by occupation, region and year. The share of routine tasks is the ratio of routine tasks to the sum of task indexes. See Section 3.2 for a description of the task indexes. See Section 3.5 for a description of the instruments.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

results for each instrument.[22]

Our estimates may be difficult to interpret because of the quadratic term in equations (3.4) and (3.5). For convenience, we define the standardized effect as the change in a given task index for an increase of ten percentage points in the graduate share around the nationwide share in 1990 (i.e. from 13 to 23 percent). The discussion focuses on estimates by two-stage least squares with the retirement instrument for parsimony's sake. Our conclusions are robust to the choice of estimator.

Consider routine tasks first. As column (1) shows, the effect of an increase in the graduate share on routine tasks is concave within occupations. According to the causal estimates, it peaks when the graduate share nears 24 percent and turns negative when it reaches 48 percent. Since the nationwide graduate share was 18 percent in 1991, we conclude that rising educational attainment caused an increase in the routine job content in France, corroborating the prediction of a substitution effect at the worker level from our model. As a reminder, workers perform more routine tasks in the model because their relative price goes up as the skill supply expands. The magnitude is modest: the standardized effect is 0.005 or 2 percent of a standard deviation.[23] The causal coefficients are larger than the estimates by linear regression, but they agree in direction. Column (2) repeats this exercise without occupation indicators. We find the opposite pattern: the impact of an expansion in the graduate share is convex and uniformly negative, bottoming out when the graduate share is just past 75 percent. Unlike the previous regression, this specification uses variation in job content between occupations for identification. Therefore, it captures a mixture of the change in the incidence of routine tasks within occupations and growing employment in cognitive occupations (in which skilled workers specialize). It constitutes

[22] We do not use both instruments together for two reasons: first, they are so correlated that we gain little power (the correlation is 0.99); second, the comparison of the results for each instrument helps us assess the robustness of our estimates.

[23] For comparison, the Bartik instrument implies that the effect of an increase in the graduate share on routine tasks peaks when the graduate share nears 20 percent and turns negative when it reaches 41 percent. The implied standardized effect is 0.002 (0.6 percent of a standard deviation).

evidence of the composition effect in the model. The standardized effect is –0.051 or 16 percent of a standard deviation.

As theory suggests, cognitive tasks mirror the routine. Column (3) implies that the impact of an increase in the graduate share on the cognitive score is negative and concave. The standardized effect is –0.02 or 8 percent of a standard deviation. By contrast, we find a positive and concave relationship upon dropping the occupation indicators, as column (4) shows. It peaks as the graduate share nears 82 percent. The standardized effect is 0.032 or 12 percent of a standard deviation. These estimates are again consistent with the two main predictions of our model: a substitution effect away from cognitive tasks at the worker level and an aggregate composition effect toward cognitive tasks.

Columns (5) and (6) examine social tasks. Unlike the routine or the cognitive, social tasks are not part of our model. We analyze them nonetheless for completeness. Perhaps because of complementarities between cognitive and social skills (Deming, 2017), our estimates are broadly similar to the regressions of the cognitive score. If we include occupation indicators, we find a negative impact of an increase in the graduate share on social tasks. The quadratic term is especially salient. The standardized effect is –0.011 or 3 percent of a standard deviation. Should we exclude occupation indicators, the response function becomes positive and concave. The maximum occurs when graduates represent 52 percent of employed workers. The standardized effect is 0.048 or 12 percent of a standard deviation.

The last two columns show regressions for the ratio of the routine score to the sum of task scores. Hence, the resulting coefficients combine the individual effects in columns (1) through (6). They are consistent with the regressions of the routine score in the first two columns, but the coefficients are not significant when occupation indicators are included.

TABLE 3.6: IMPACT OF JOB CONTENT ON WAGES

	Hourly wages (log)			Monthly wages (log)		
	1991	1998	2005	1991	1998	2005
Routine tasks	−0.033***	−0.024***	−0.022***	−0.029***	−0.030***	−0.031***
	(0.007)	(0.008)	(0.010)	(0.008)	(0.008)	(0.009)
Cognitive tasks	0.078***	0.050***	0.035***	0.099***	0.081***	0.054***
	(0.010)	(0.010)	(0.011)	(0.010)	(0.010)	(0.012)
Social tasks	−0.010*	−0.033***	−0.017**	0.015***	−0.018***	0.014*
	(0.006)	(0.006)	(0.007)	(0.006)	(0.007)	(0.008)
Controls						
Individual characteristics	16	16	16	15	15	15
Occupation fixed effects	22	22	22	22	22	22
Region fixed effects	20	20	20	20	20	20
Fit						
Sample	All	All	All	Full time	Full time	Full time
Observations	16,819	17,522	15,514	14,936	14,694	12,790
Mean outcome	2.396	2.420	2.555	7.448	7.483	7.528
Adjusted R^2	0.606	0.610	0.548	0.634	0.634	0.599

Notes: The table shows coefficients from ordinary linear regressions. Standard errors (in parentheses) are clustered by occupation and region. See Section 3.2 for a description of the task indexes.
Legend: Stars denote significance: *, at the 10 percent level; **, 5 percent; ***, 1 percent.

3.6.2 *The impact of task assignment on wages*

Our model predicts that an increase in the skill supply should raise the price of routine tasks and reduce that of cognitive tasks. As a rudimentary test of this prediction, we undertake yearly regressions of wages on task assignment. Because educational attainment increases throughout the period, there should be changes in task prices if the theory is correct. Our coefficients are not causal (Autor and Handel, 2013; DiNardo and Pischke, 1997), since task assignment is a function of workers' comparative advantages and we do not have instruments for tasks. We present them nonetheless as preliminary empirical evidence and for comparison with Autor and Handel (2013). Table 3.6 shows our results. Note that we do not use the 2013 WCS because it did not collect comparable wage data.[24]

The first three columns consider hourly wages. Routine and cognitive tasks have opposite

[24] The 1993 WOS and the 1991 WCS yield similar estimates. We only report results for 1991 for parsimony. Results for 1993 are available upon request.

effects on pay: routine tasks lower wages, whereas cognitive tasks raise them. The penalty for routine tasks falls from 3.3 percent of hourly wages in 1991 to 2.4 in 1998 and 2.2 in 2005. These estimates represent the wage loss for performing all three routine tasks in the survey as opposed to none. Conversely, the premium for cognitive tasks declines from 7.8 percent of hourly wages in 1991 to 5 in 1998 and 3.5 in 2005. The pattern of changes between surveys match our theoretical predictions. The picture is less clear for social tasks: they reduce pay as well, but the coefficient increases in magnitude from 1991 to 1998 and decreases from 1998 to 2005. Deming (2017) also finds a negative correlation between social tasks and compensation.

The last three columns examine monthly wages. We restrict the sample to full-time workers. Routine tasks have similar effects on monthly and hourly wages, but the monthly penalty is stable across years. The impact of cognitive tasks on monthly wages is slightly larger than its hourly counterpart. It too shrinks between surveys, going from 10 percent of monthly wages in 1991 to 5.4 percent in 2005. By contrast, the role of social tasks is equivocal. The coefficient is positive in 1991 and 2005 but negative in 1998. Although we find a similar dip in 1998 for hourly wages, all coefficients are negative in that case.

Our wage regressions are consistent with Autor and Handel's (2013). The authors classify tasks into three groups: abstract, routine and manual. (Their abstract tasks correspond to our cognitive. They do not discuss social tasks.) They study the impact of job content on wages within occupations with survey data from the United States. Although they examine a different labor market and measure job content through different variables, their coefficients are surprisingly similar to ours in direction and magnitude. Albeit exploratory, our analyses should provide useful guidance for future research on task assignment.

3.7 Conclusion

This paper contributes to a growing literature about task assignment. Since the seminal work of Autor, Levy and Murnane (2003), this research has provided insight into employment polarization, wage inequality and much else. It has mostly studied the influence of job content on an outcome of interest. We take the opposite approach and investigate the determination of job content in equilibrium. In particular, we show that the skill supply affects job content by analyzing the impact of an increase in the share of university graduates in the French workforce from 18 percent in 1991 to 36 percent in 2013. We find that higher average educational attainment is associated with more routine, fewer cognitive and fewer social tasks within occupations and with fewer routine, more cognitive and more social tasks across occupations.

Our results have three methodological implications for future research. First, researchers should explore variation in job content within occupations in greater depth. Second, occupations evolve, so care is needed in analyzing long-term trends in labor markets on the basis of rigid occupational classifications. Third, identification deserves attention as task assignment is endogenous to both worker aptitudes and aggregate conditions.

Our approach has two significant limitations. Although we can measure changes in task assignment at the individual level, we cannot distinguish the intensive margin (i.e. changes in tasks for a given worker in a given job) and the extensive margin (i.e. changes through job creation, job destruction and employee turnover). Moreover, we cannot observe innovation in tasks in our data. Task creation has historically offset the pressure of automation on wages and employment (Acemoglu and Restrepo, 2018, 2019), such as we experience today. Panel data would help us address these shortcomings. It would also allow us to study the influence of job content on workers' careers, wage inequality and more – a promising avenue for future research.

References

ABADIE, A. (2003): “Semiparametric instrumental variable estimation of treatment response models.” *Journal of Econometrics* 113(2), 231–263.
ABREU, D., AND F. GUL (2000): “Bargaining and reputation.” *Econometrica* 68(1), 85–117.
ACEMOGLU, D., AND D. H. AUTOR (2011): “Skills, tasks and technologies: Implications for employment and earnings.” In *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, vol. 4B, 1043–1171. Amsterdam, NL: North-Holland.
ACEMOGLU, D., AND P. RESTREPO (2018): *Artificial intelligence, automation and work*. Working Paper 24196, NBER.
——— (2019): “Automation and new tasks: How technology displaces and reinstates labor.” *Journal of Economic Perspectives* 33(2), 3–30.
AEBERHARDT, R., P. GIVORD AND C. MARBOT (2016): *Spillover effect of the minimum wage in France: An unconditional quantile regression approach*. Working Paper 2016-05, CREST.
AGRAWAL, A., J. S. GANS AND A. GOLDFARB (2019): “Artificial intelligence: The ambiguous labor market impact of automating prediction.” *Journal of Economic Perspectives* 33(2), 31–50.
ANDREWS, J. B. (1918): “Nationalisation (1860–1877).” In *History of Labour in the United States*, ed. by J. R. Commons, vol. II, part V, 1–191. New York, NY: Macmillan.
ANSELL, C. K., AND A. JOSEPH (1998): “The mass production of craft unionism: Exploring workers’ solidarity in late nineteenth-century France and America.” *Politics & Society* 26(4), 575–602.
ARNTZ, M., T. GREGORY AND U. ZIERAHN (2017): “Revisiting the risk of automation.” *Economic Letters* 159(C), 157–160.
ASHENFELTER, O., AND G. E. JOHNSON (1969): “Bargaining theory, trade unions and industrial strike activity.” *American Economic Review* 59(1), 35–49.
ATACK, J. (2016): *Historical Geographic Information Systems (GIS) database of U.S. Railroads*, URL: `my.vanderbilt.edu/jeremyatack/data-downloads/`.
ATACK, J., AND F. BATEMAN (1992): “How long was the workday in 1880?” *Journal of Economic History* 52(1), 129–160.
——— (2016): *State samples from the 1880 Census of Manufacturing*. Ann Arbor, MI. DOI: `10.3886/ICPSR09384.v2`.
ATACK, J., R. A. MARGO AND P. W. RHODE (2019): “‘Automation’ of manufacturing in the late nineteenth century: The Hand and Machine Labor study.” *Journal of Economic Perspectives* 33(2), 51–70.
AUTOR, D. H. (2015): “Why are there still so many jobs? The history and future of workplace automation.” *Journal of Economic Perspectives* 29(3), 3–30.
AUTOR, D. H., AND D. DORN (2013): “The growth of low-skill service jobs and the polarization of the US labor market.” *American Economic Review* 103(5), 1553–1597.
AUTOR, D. H., D. DORN AND G. H. HANSON (2015): “Untangling trade and technology: Evidence from local labour markets.” *Economic Journal* 125(584), 621–646.
AUTOR, D. H., AND M. J. HANDEL (2013): “Putting tasks to the test: Human capital, job tasks, and wages.” *Journal of Labor Economics* 31(S1), S59–S96.
AUTOR, D. H., F. LEVY AND R. J. MURNANE (2003): “The skill content of recent technological change: An empirical exploration.” *Quarterly Journal of Economics* 118(4), 1279–1333.
BAILEY, G. L. (1991): “The Commissioner of Labor’s *Strikes & Lockouts*: A cautionary note.” *Labor History* 32(3), 432–440.

BARTIK, T. J. (1991): *Who Benefits from State and Local Economic Development Policies?* Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.

BEVAN, G. P. (1880): "The strikes of the past ten years." *Journal of the Statistical Society of London* 43(1), 35–64.

BIGGS, M. (2002): "Strikes as sequences of interaction." *Social Science History* 26(3), 583–617.

BIRDSALL, W. C. (1953): "The problem of structure in the Knights of Labor." *Industrial and Labor Relations Review* 6(4), 532–546.

BLACK, S. E., AND A. SPITZ-OENER (2010): "Explaining women's success: Technological change and the skill content of women's work." *Review of Economics and Statistics* 92(1), 187–194.

BLINDER, A. S. (2009): "How many U.S. jobs might be offshorable?" *World Economics* 10(2), 41–78.

BOZIO, A., T. BREDA AND M. GUILLOT (2016): *Taxes and technological determinants of wage inequalities: France 1976–2010*. Working Paper 2016-05, PSE.

BUREAU OF INDUSTRIAL STATISTICS OF PENNSYLVANIA (1882): "Labor troubles in Pennsylvania." In *Annual Report of the Secretary of Internal Affairs of the Commonwealth of Pennsylvania*, vol. IX, part III, 262–391. Harrisburg, PA: Lane S. Hart.

BUREAU OF LABOR STATISTICS OF OHIO (1878): *First Annual Report of the Bureau of Labor Statistics*. Columbus, OH: Nevins & Myers.

BUREAU OF STATISTICS OF LABOR OF MASSACHUSETTS (1880): "Strikes in Massachusetts." In *Eleventh Annual Report of the Bureau of Statistics of Labor*, part I, 1–71. Boston, MA: Rand, Avery, & Co.

CARD, D. (1990): "Strikes and bargaining: A survey of the recent empirical literature." *American Economic Review: Papers and Proceedings* 80(2), 410–415.

CARD, D., F. KRAMARZ AND T. LEMIEUX (1999): "Changes in the relative structure of wages and employment: A comparison of the United States, Canada, and France." *Canadian Journal of Economics* 32(4), 843–877.

CARD, D., T. LEMIEUX AND W. C. RIDDEL (2004): "Unions and wage inequality." *Journal of Labor Research* 25(4), 519–562.

CARD, D., AND C. A. OLSON (1995): "Bargaining power, strike durations, and wage outcomes: An analysis of strikes in the 1880s." *Journal of Labor Economics* 13(1), 32–61.

CHARNOZ, P., É. COUDIN AND M. GAINI (2011): *Changes in the French wage distribution 1976–2004: Inequalities within and between education and experience groups*. Working Paper 2011-23, CREST.

CLARKE, P. S., AND F. WINDMEIJER (2012): "Instrumental variable estimators for binary outcomes." *Journal of the American Statistical Association* 107(500), 1638–1652.

COLLINS, W. J., AND G. T. NIEMESH (2018): "Unions and the Great Compression of wage inequality in the US at mid-century: Evidence from local labour markets." *Economic History Review* 72(2), 691–715.

CONLEY, T. G. (1999): "GMM estimation with cross sectional dependence." *Journal of Econometrics* 92(1), 1–45.

CRAMTON, P., AND J. TRACY (2003): "Unions, bargaining and strikes." In *International Handbook of Trade Unions*, ed. by J. T. Addison and C. Schnabel, 86–117. Cheltenham, UK: Edward Elgar.

CRÉPON, B., AND C. GIANELLA (1999): *Wage Inequality in France 1969-1992*. Working Paper G9905, DESE (INSEE).

CURRIE, J., AND J. FERRIE (2000): "The law and labor strife in the United States, 1881–1894." *Journal of Economic History* 60(1), 42–66.

DEMING, D. J. (2017): "The growing importance of social skills in the labor market." *Quarterly Journal of Economics* 132(4), 1593–1640.

DINARDO, J. E., AND J.-S. PISCHKE (1997): "The returns to computer use revisited: Have pencils changed the wage structure too?" *Quarterly Journal of Economics* 112(1), 291–303.

DINARDO, J., AND D. S. LEE (2004): "Economic impacts of new unionization on private sector employers: 1984–2001." *Quarterly Journal of Economics* 119(4), 1383–1441.

EDWARDS, A. (1933): "A social-economic grouping of the gainful workers in the United States." *Journal of the American Statiscal Association* 28(4), 529–542.

EICHENGREEN, B. (1987): "The impact of late nineteenth-century unions on labor earnings and hours: Iowa in 1894." *Industrial and Labor Relations Review* 40(1), 501–515.

ELSAYED, A., A. DE GRIP AND D. FOURGE (2017): "Computer use, job tasks and the part-time pay penalty." *British Journal of Industrial Relations* 55(1), 58–82.

ENFLO, K., AND T. KARLSSON (2018): "The importance of mediation in Swedish work stoppages 1907–1927." *European Review of Economic History*. DOI: 10.1093/ereh/hey023.

FARBER, H. S., D. HERBST, I. KUZIEMKO AND S. NAIDU (2018): *Unions and inequality over the twentieth century: New evidence from survey data*. Working Paper 24587, NBER.

FIRPO, S., N. M. FORTIN AND T. LEMIEUX (2011): *Occupational tasks and changes in the wage structure*. Discussion Paper 5542, IZA.

FIRPO, S., AND C. PINTO (2016): "Identification and estimation of distributional impacts of interventions using changes in inequality measures." *Journal of Applied Econometrics* 31(3), 457–486.

FREEMAN, R. B., AND J. L. MEDOFF (1984): *What Do Unions Do?* New York, NY: Basic Books.

FRIEDMAN, G. (1988): "Strike success and union ideology: The United States and France, 1880–1914." *Journal of Economic History* 48(1), 1–25.

——— (1999): "New estimates of union membership: The United States, 1880–1914." *Historical Methods* 32(2), 75–86.

FRÖLICH, M., AND B. MELLY (2013): "Unconditional quantile treatment effects under endogeneity." *Journal of Business & Economic Statistics* 31(3), 346–357.

GARLOCK, J. (1982): *Guide to the Local Assemblies of the Knights of Labor*. Westport, CT: Greenwood.

——— (2009): *Knights of Labor Assemblies, 1879–1889*. Ann Arbor, MI. DOI: 10.3886/ICPSR00029.v1.

GATHMANN, C., AND U. SCHÖNBERG (2010): "How general is human capital? A task-based approach." *Journal of Labor Economics* 28(1), 1–49.

GERAGHTY, T. M., AND T. WISEMAN (2008): "Wage strikes in 1880s America: A test of the war of attrition model." *Explorations in Economic History* 45(4), 303–326.

GOLDSMITH-PINKHAM, P., I. SORKIN AND H. SWIFT (2018): *Bartik instruments: What, when, why, and how*. Working Paper 24408, NBER.

GOUX, D. (2003): "Une histoire de l'Enquête Emploi." *Économie et Statistique* 362, 41–57.

GOUX, D., AND E. MAURIN (2000): "The decline in demand for unskilled labor: An empirical analysis method and its application to France." *Review of Economics and Statistics* 82(4), 596–607.

GRAMM, C. L., AND J. F. SCHNELL (1994): "Difficult choices: Crossing the picket line during the 1987 national football strike." *Journal of Labor Economics* 12(1), 41–73.

GREGORY, T., A. SALOMONS AND U. ZIERAHN (2019): *Racing with or against the machine? Evidence from Europe*. Discussion Paper 12063, IZA.

GROB, G. N. (1958): "The Knights of Labor and the trade unions, 1878–1958." *Journal of Economic History* 18(2), 176–192.

——— (1960): "Origins of the political philosophy of the A. F. of L., 1886–1896." *Review of Politics* 22(4), 496–518.

HAINES, M. R. (1989): "A state and local consumer price index for the United States in 1890." *Historical Methods* 22(3), 97–105.

HAINES, M. R., AND ICPRS (2010): *Historical, demographic, economic, and social data: The United States, 1790–2002*. Ann Arbor, MI. DOI: 10.3886/ICPSR02896.v3.

HALLGRÍMSDÓTTIR, H. K., AND C. BENOIT (2007): "From wage slaves to wage workers: Cultural opportunity structures and the evolution of the wage demands of the Knights of Labor and the American Federation of Labor, 1880–1900." *Social Forces* 85(3), 1393–1411.

HANDEL, M. J. (2012): *Trends in job skill demands in OECD countries*. Social, Employment and Migration Working Paper 143, OECD.

HAUSMAN, J. A. (1978): "Specification tests in econometrics." *Econometrica* 46(6), 1251–1271.

HAYES, B. (1984): "Unions and strikes with asymmetric information." *Journal of Labor Economics* 2(1), 57–83.

HIRANO, K., G. W. IMBENS AND G. RIDDER (2003): "Efficient estimation of average treatment effects using the estimated propensity score." *Econometrica* 71(4), 1161–1189.

IMBENS, G. W. (2015): "Matching methods in practice: Three examples." *Journal of Human Resources* 50(2), 373–419.

IMBENS, G. W., AND J. D. ANGRIST (1994): "Identification and estimation of local average treatment effects." *Econometrica* 62(2), 467–475.

JENSEN, B. J., AND L. G. KLETZER (2010): "Measuring tradable services and the task content of offshorable services jobs." In *Labor in the New Economy*, ed. by K. Abraham, M. Harper and J. Spletzer, 309–35. Chicago, IL: University of Chicago Press for the NBER.

KARLSSON, T. (2019): *Strikes and lockouts in Sweden: Reconsidering Raphael's list of work stoppages 1859–1902*. Lund Papers in Economic History 192.

KATZ, L. F., AND R. A. MARGO (2014): "Technical change and the relative demand for skilled labor: The United States in historical perspective." In *Human Capital in History: The American Record*, ed. by L. P. Boutsna, C. Frydman and R. A. Margo, 15–57, NBER. Chicago, IL: University of Chicago Press.

KAUFMAN, J. (2001): "Rise and fall of a nation of joiners: The Knights of Labor revisited." *Journal of Interdisciplinary History* 31(4), 553–579.

KEMMERER, D. L., AND E. D. WICKERSHAM (1950): "Reasons for the growth of the Knights of Labor in 1885–1886." *ILR Review* 3(2), 213–220.

KENNAN, J. (1986): "The economics of strikes." In *Handbook of Labor Economics*, ed. by O. Ashenfelter and R. Layard, vol. II, 1092–1137. Amsterdam: Elsevier.

KENNAN, J., AND R. WILSON (1989): "Strategic bargaining models and interpretation of strike data." *Journal of Applied Econometrics* 4(S1), S87–S130.

KREMER, M., AND B. A. OLKEN (2009): "A biological model of unions." *American Economic Journal: Applied Economics* 1(2), 150–175.

LEE, D. S., AND A. MAS (2012): "Long-run impacts of unions on firms: New evidence from financial markets, 1961–1999." *Quarterly Journal of Economics* 127(1), 333–378.

LEVY, F., AND R. J. MURNANE (1996): "With what skills are computers a complement?" *American Economic Review: Papers and Proceedings* 86(2), 258–262.

LEWBEL, A., Y. DONG AND T. T. YANG (2012): "Comparing features of convenient estimators for binary choice models with endogenous regressors." *Canadian Journal of Economics* 45(3), 809–829.

LLOYD, J. P. (2009): "The strike wave of 1877." In *The Encyclopedia of Strikes in American History*, ed. by A. Brenner, B. Day and I. Nes, 177–190. London, UK: Routledge.

MACHADO, C., A. M. SHAIKH AND E. J. VYTLACIL (2018): "Instrumental variables and the sign of the average treatment effect," Unpublished manuscript.

MANSON, S., J. SCHROEDER, D. VAN RIPER AND S. RUGGLES (2018): *IPUMS National Historical Geographic Information System: Version 13.0*. Minneapolis, MN. DOI: 10.18128/D050.V13.0.

MCCONNELL, S. (1990): "Cyclical fluctuations in strike activity." *Industrial and Labor Relations Review* 44(1), 130–143.

MITTELMAN, E. B. (1918): "Trade unionism (1833–1839)." In *History of Labour in the United States*, ed. by J. R. Commons, vol. I, part III, 333–484. New York, NY: Macmillan.

MOORE, C. (2013): *Margaret Thatcher: The Authorized Biography*, vol. 1: *From Grantham to the Falklands*. New York, NY: Alfred A. Knopf.

——— (2016): *Margaret Thatcher: The Authorized Biography*, vol. 2: *At Her Zenith: In London, Washington and Moscow*. New York, NY: Alfred A. Knopf.

NAIDU, S., AND N. YUCHTMAN (2018): "Labor market institutions in the gilded age of American economic history." In *The Oxford Handbook of American Economic History*, ed. by L. P. Cain, P. V. Fishback and P. W. Rhode, vol. 1, 329–354. Oxford, UK: Oxford University Press.

NEWEY, W. K., AND D. L. MCFADDEN (1994): "Large sample estimation and hypothesis testing." In *Handbook of Econometrics*, ed. by R. F. Engle and D. L. McFadden, vol. IV, 2111–2245. Amsterdam: Elsevier.

OESTREICHER, R. (1984): "A note on Knights of Labor membership statistics." *Labor History* 25(1), 102–108.

OLOFSGÅRD, A. (2012): "Union leaders as experts: Wage bargaining and strikes with union-wide ballot requirements." *Scandinavian Journal of Economics* 114(1), 200–227.

PERI, G., AND C. SPARBER (2009): "Task specialization, immigration, and wages." *American Economic Journal: Applied Economics* 1(3), 135–169.

PERLMAN, S. (1918): "Upheaval and reorganisation (Since 1876)." In *History of Labour in the United States*, ed. by J. R. Commons, vol. II, part VI, 193–537. New York, NY: Macmillan.

PETERSON, F. (1938): “Strikes in the United States, 1880–1936.” Bulletin 651, BLS. Washington, DC: Government Printing Office.

POWDERLY, T. V. (1884): “Address of the General Master Workman.” In *Proceedings of the General Assembly of the Knights of Labor of America*, 569–579. Minneapolis, MN: General Assembly.

ROSENBAUM, P. R., AND D. B. RUBIN (1983): “The central role of the propensity score in observational studies for causal effects.” *Biometrika* 70(1), 41–55.

ROSENBLOOM, J. L. (1998): “Strikebreaking and the labor market in the United States, 1881–1894.” *Journal of Economic History* 58(1), 183–205.

RUGGLES, S., S. FLOOD, R. GOEKEN, J. GROVER, J. PACAS AND M. SOBEK (2018): *IPUMS USA: Version 8.0*. Minneapolis, MN. DOI: 10.18128/D010.V8.0.

SAPOSS, D. J. (1918): “Colonial and federal beginnings (to 1827).” In *History of Labour in the United States*, ed. by J. R. Commons, vol. I, part I, 23–165. New York, NY: Macmillan.

SCHMICK, E. (2018): “Collective action and the origins of the American labor movement.” *Journal of Economic History* 78(3), 744–784.

SMITH, R. M. (1888): “Third Annual Report of the Commissioner of Labor” [review]. *Political Science Quarterly* 3(4), 709–710.

SPITZ-OENER, A. (2006): “Technical change, job tasks, and rising educational demands: Looking outside the wage structure.” *Journal of Labor Economics* 24(2), 235–270.

SUMNER, H. L. (1918): “Citizenship (1827–1833).” In *History of Labour in the United States*, ed. by J. R. Commons, vol. I, part II, 167–332. New York, NY: Macmillan.

TILLY, C., AND D. K. JORDAN (2012): *Strikes and labor activity in France, 1830–1960*. Ann Arbor, MI. DOI: 10.3886/ICPSR08421.v2.

U.S. BUREAU OF LABOR (1888): *Third Annual Report of the Commissioner of Labor: Strikes and Lockouts*. Washington, DC: Government Printing Office.

——— (1896): *Tenth Annual Report of the Commissioner of Labor: Strikes and Lockouts*, vol. I. Washington, DC: Government Printing Office.

——— (1906): *Twenty-First Annual Report of the Commissioner of Labor: Strikes and Lockouts*. Washington, DC: Government Printing Office.

VAN REENEN, J. (2011): “Wage inequality, technology and trade: 21st century evidence.” *Labour Economics* 18(6), 730–741.

VERDUGO, G. (2014): “The great compression of the French wage structure, 1969–2008.” *Labour Economics* 28(C), 131–144.

VERDUGO, G., H. FRAISSE AND G. HORNY (2012): “Évolution des inégalités salariales en France: Le rôle des effets de composition.” *Revue économique* 63(6), 1081–1112.

VOSS, K. (1993): *The Making of American Exceptionalism: The Knights of Labor and Class Formation in the Nineteenth Century*. Ithaca, NY: Cornell University Press.

WARREN, G. F., AND F. A. PEARSON (1932): “Wholesale prices in the United States for 135 years, 1797 to 1932.” In *Wholesale Prices for 213 Years, 1720 to 1932*, part I, 5–200. Memoir 142, Cornell University Agricultural Experiment Station. Ithaca, NY: Cornell University.

WEEKS, J. D. (1886): “Report on strikes and lockouts occurring within the United States during the calendar year 1880.” In *Report on the Statistics of Wages in Manufacturing Industries; with Supplementary Reports on the Average Retail Prices of Necessaries of Life, and on Trades Societies, and Strikes and Lockouts*. Washington, DC: U.S. Government Printing Office.

WRIGHT, C. D. (1887): “An historical sketch of the Knights of Labor.” *Quarterly Journal of Economics* 1(2), 137–168.