

Dube, Arindrajit; Lindner, Attila

Working Paper

Minimum wages in the 21st century

RF Berlin - CReAM Discussion Paper Series, No. 25/24

Provided in Cooperation with:

Rockwool Foundation Berlin (RF Berlin)

Suggested Citation: Dube, Arindrajit; Lindner, Attila (2024) : Minimum wages in the 21st century, RF Berlin - CReAM Discussion Paper Series, No. 25/24, Rockwool Foundation Berlin (RF Berlin) and Centre for Research & Analysis of Migration (CReAM), Berlin and London

This Version is available at:

<https://hdl.handle.net/10419/307683>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Minimum Wages in the 21st Century

Arindrajit Dube, Attila Lindner

OCTOBER 2024

Minimum Wages in the 21st Century *

Arindrajit Dube

Umass Amherst, NBER

Attila Lindner

UCL, IFS, CERS

September 11, 2024

Abstract

This chapter surveys the literature on the impact of minimum wages on low-wage labor markets. We describe and critically review the empirical methods in the new minimum wage literature, particularly those leveraging quasi-experimental variation. We provide a quantitative overview of the most recent evidence on the employment and wage effects of the policy, while also exploring emerging research on its impact on other margins, including amenities, other inputs (such as capital and high-skilled workers), firm entry and exit, output prices and demand, profits, and productivity. This approach allows us to present a comprehensive picture of how minimum wage policies affect firms, workers, and labor markets. We also review the evidence on the policy's impact on wage inequality and income distribution. Finally, we discuss how these effects can vary depending on the economic context and the level of a country's development.

keywords: minimum wages, labor demand, employment and wage effects, margins of adjustment, inequality

JEL-codes:

*We thank Akash Bhatt, Annie McGrew, Jon Piqueras, and Linda Wu for excellent research assistance. We are also grateful to David Card, Christian Dustmann, Larry Katz, Pat Kline, Thomas Lemieux, Alan Manning, Michael Reich, Liyang Sun and the participants of the Conference for the Handbook of Labor Economics, Volume 5 in Berlin for the valuable comments and insights. We are grateful to the Rockwool Foundation Berlin (RFBerlin) for making this conference possible. This research has been also funded by the European Union's Horizon 2020 research and innovation program (grant agreement number 949995).

1 Introduction and overview

Minimum wage policies have evolved from their initial introduction, and have become an important tool used by many countries to address various economic and social challenges. The debate and perspectives on minimum wages have also evolved considerably, influenced by both changing economic theories and empirical research. Although economists were highly skeptical about the benefits of the policy throughout much of the previous century, that view was seriously challenged beginning in the 1990s. Today, in the 21st century, there is widespread interest in making greater use of minimum wages (Manning, 2021). Countries around the world are introducing or raising their minimum wages at an unprecedented rate. Major international organizations such as the OECD, the IMF, and the World Bank have advocated for the policy as a means to alleviate inequality and even boost employment (ILO et al., 2012).

In the research community, there has been a proliferation of empirical studies that develop new techniques, take advantage of new data sources, and consider new aspects of the policy. The literature has also expanded from primarily based on evidence from the United States to examining the impact of the policy worldwide. Meanwhile, there have been significant theoretical advances in understanding how labor markets operate and how those models can be applied to better understand the impact of wage floors.

This chapter summarizes the key insights from the literature on minimum wages, with a focus on advancements made over the last two decades. We begin by reviewing the motivation behind the introduction of minimum wages and how these policies have evolved over time and across countries. Next, we provide an overview of the intellectual debate surrounding minimum wages. This topic has been highly controversial among economists and has significantly influenced the field of economics, particularly in economic modeling and research methods.

The relationship between employment and wages is central to testing the neoclassical view of low-wage labor markets. However, the best approach to testing that relationship has been the subject of a long-running debate among marginalist (e.g., Stigler (1946)), institutionalist (e.g., Lester (1946)), and empiricist (e.g., Card and Krueger (1994b)) traditions within the discipline. The minimum wage literature has also played a crucial role in shifting economics from a predominantly theoretical discipline to an empirical and data-driven one.

The “new” minimum wage research began in the early 1990s, focusing on studying the policy’s impact using quasi-experimental research designs. In practice, this has involved the use of various difference-in-differences style estimators and their intellectual offshoots. In this chapter, we take stock of these methods and critically examine them, with a focus on the contributions made since the last *Handbook of Labor Economics* review by (Brown, 1999). The literature has advanced considerably since then.

We discuss the introduction of more sophisticated empirical approaches to studying the impact of minimum wages, such as synthetic control, border discontinuity design, and refined difference-in-differences estimators like the stacked event-study approach. We also address challenges with common approaches,

such as the two-way fixed effects regression, which has been widely used in the literature. Additionally, we cover empirical methods developed outside the U.S. context, particularly those that exploit nationwide variation in minimum wages. Finally, we explore recent advances in understanding the broader impact of the policy on low-wage workers, moving beyond a sole focus on specific groups like teens.

Methodological advances over the past three decades have provided a more nuanced understanding of the effects of minimum wage policies. We begin by offering an up-to-date review of the evidence on employment and wages. Specifically, we summarize estimates of the own-wage employment elasticity, which compares the percentage change in employment to the percentage change in wages—representing the labor demand elasticity in the standard competitive model. We also examine the heterogeneous impacts of the policy across individuals, firms, regions, and depending on the nature of the minimum wage shocks. In this chapter, we use data from U.S. minimum wage events to illustrate some of the empirical challenges in the literature and demonstrate how state-of-the-art approaches can yield reliable estimates for both wages and employment.

However, simply studying the effects of the policy on employment and wages does not fully capture how firms, workers, and markets respond to minimum wage changes. Fortunately, the empirical literature of the 21st century has significantly advanced our understanding of the various margins of adjustment that play a role in response to minimum wage policies. We explore the effects of minimum wages on amenities, other inputs (such as capital and higher-skilled workers), firm entry and exit, output prices and demand, profits, and productivity. While the evidence on some of these margins is still in its early stages, a relatively clear picture has emerged regarding the importance of price pass-through and the productivity-enhancing effects of the policy. We also offer ideas on how these findings could shape future theoretical developments in understanding low-wage labor markets.

Minimum wage policies are a significant tool for combating rising inequality. Considerable attention has been devoted to quantifying how minimum wages shape the wage distribution and their role in the evolution of inequality over time. We review the evidence and methods developed to study the inequality consequences of these policies. Additionally, the minimum wage is often seen as a major redistributive tool outside of the tax and benefit system. In light of this, we discuss the evidence on the distributional implications of the policy.

While the main focus of this review is on understanding the labor market consequences of the policy, there are also broader implications worth considering. An extensive literature has explored the impact of minimum wages on various “downstream” socioeconomic outcomes, such as health, crime, and education. Although our review cannot delve into the details of these areas, we provide a summary of the state-of-the-art findings regarding these indirect consequences of the policy.

This chapter broadens the scope of existing reviews by examining the impact of minimum wages in various national and economic contexts. The previous *Handbook* chapter on minimum wages by [Brown \(1999\)](#) provided an excellent review of the evolution and history of minimum wage policy in the U.S. However, a key feature of 21st-century research is the increasing number of high-quality studies emerging from outside the U.S. This reflects the reality that major policy changes have been instituted globally over the last three decades—including the (re)introduction of minimum wages in places like

China (1993), the U.K. (1999), Hong Kong (2010), and Germany (2015)—alongside growing access to high-quality administrative datasets. Throughout this chapter, we highlight evidence from these novel analyses.

Additionally, a growing number of studies are exploring the implications of minimum wage policies in less developed economies. We conclude our review by summarizing the key findings from this emerging literature and discussing how the policy operates in less developed economic contexts.

2 A brief history of minimum wages

2.1 The rationale for minimum wage policies

The motivation behind minimum wage regulations has varied significantly across countries, regions, and historical periods. The first modern-day minimum wages were enacted in New Zealand in 1894 and in the Australian state of Victoria in 1896 (Starr, 1981). The British Parliament adopted legislation in the same spirit in 1909. Initially, these minimum wages focused only on specific industries (e.g., only four trades for Great Britain: chain making by hand, paper-box making, lace finishing, and wholesale tailoring). The main rationale was to prevent and settle industrial disputes and to eliminate “sweating”—the payment of exceptionally low wages (Webb, 1912; Hammond, 1915; Metcalf, 1999).¹ Proliferation of “sweatshops” in 1890s was a major concern in these countries (Nordlund, 1997). Sweatshops were typically lower-productivity businesses that relied on recruiting cheap labor as their business model, primarily employing women and young workers (including orphans) and paying them substandard wages. More productive companies employing working-class breadwinners were threatened by these business practices and were often supportive of initiatives that sought to reverse these trends.

In the U.S., minimum wage legislation first emerged at the state-level in the early 1910s, beginning in Massachusetts. Concern for workers seen as being the most vulnerable to exploitation by low-paying employers—such as immigrants, women, and children—played a key role in the push for minimum standards through state legislation. However, the United States Supreme Court invalidated most of these laws. Minimum wage legislation was also a part of the National Industrial Recovery Act (NRA) introduced in 1933, but the States Supreme Court later found that legislation unconstitutional.

In 1938, the first federal minimum wage (25 cents per hour, or around \$4.50 in 2023\$) was established under the Fair Labor Standards Act, which also aimed to regulate hours and restrict child labor (Grossman, 1978; Brown, 1999; Fishback and Seltzer, 2021). Although the main objective of the act was to ensure the safety and well-being of workers, the law was also designed to mitigate race-to-the-bottom competition (Newell, 2009). Importantly, the Fair Labor Standards Act had a limited scope and applied only to workers in “inter-state or foreign commerce”, which in practice meant workers in manufacturing and tradable goods sectors. Private sector coverage gradually increased from

¹Although these are the first examples of wage floors in recent history, the concept of a minimum wage goes back much farther. For example, the policy was present in the Hammurabi code (Rositani, 2017).

around 50% in 1938 to 60% in the 1960s (Brown, 1999). In 1967, a significant extension in coverage brought in much of the service workforce, such as restaurants, laundries, and retail sectors, reaching almost 80% of all private sector workers (see Bailey et al. (2021); Derenoncourt and Montialoux (2021)). With further gradual increases, the coverage eventually reached 90% of workers (Brown, 1999).²

Latin American countries were also among the first to introduce minimum wage laws, with many, such as Mexico and Brazil, adopting these laws in the 1930s and 1940s (Grimshaw and Miozzo, 2003). Unlike in the United States, constitutions in these countries played a major positive role in establishing pay standards. Many countries in the region had constitutional provisions recognizing the right of workers to receive wages sufficient for a decent standard of living, and explicitly stating the state’s responsibility for setting minimum wages. (Collier and Collier, 2002). This state intervention in labor relations was also motivated by an effort to make workers look to the state, rather than to unions, to protect their interests. During this period, Latin American labor unions were often associated with radical political ideologies that challenged state authority (Cook, 2010).

After the Second World War, countries in western Europe started adopting minimum wage policies with examples including Luxembourg (1945), France (1950), Spain (1963) and the Netherlands (1969) (Dolado et al., 1996). This expansion occurred during a period when the idea that workers should have the right to protection against extremely low wages was gaining traction (ILO, 2017). However, many European countries with strong union representation, such as Austria, Italy, Denmark, Norway, and Sweden, avoided legislating minimum wages. Instead, these countries opted for collective bargaining, where unions had greater influence in setting labor standards. Unions in these countries were often opposed to minimum wage legislation, fearing that direct state involvement in labor relations could render them redundant (Meyer, 2016). There were also concerns that state-sponsored wage floors could end up becoming wage ceilings, constraining unions’ ability to negotiate better standards. In recent decades, however, the decline in union membership and coverage—especially in the low-wage service sector—has led to a more favorable view of minimum wage regulations among unions in some countries, like Germany, paralleling the introduction of minimum wage legislation.

Minimum wages in Africa were implemented on a significant scale around the 1940s and 1950s. Their pay regulations were closely linked to the colonial powers and were modeled after the wage councils in developed countries where standards were set only for specific trades (Starr, 1981; Malan, 1978). In Asia, India introduced its minimum wage policy in 1948 (though today provincial-level minimum wages are the relevant standards), Taiwan in 1956, and Japan in 1959. More recently, other countries in the region followed suit, including South Korea in 1988, China in 1993, and Vietnam in 2006. Singapore stands out as somewhat of an outlier that does not have a minimum wage in place. The government has long held the belief that the efficiency cost of a minimum wage is too high and the key to helping low-wage earners is to incentivize them to “climb up the skill ladder” (Ministry of Manpower in Singapore, 2018).

Figure 1 plots recent trends in minimum wages for several major economies around the world

²Minimum wage exceptions apply under specific circumstances to workers with disabilities, full-time students, youth under age 20 in their first 90 consecutive calendar days of employment, tipped employees, and student-learners.

(Brazil, China, India, Germany, France, Spain, the U.K., and the U.S.). To make minimum wages comparable over time and across countries, we calculate the ratio of the minimum wage to the average wage in each country.³ Unfortunately, average wage data are not always available for past periods, which limits the time frames for which the minimum wage ratio can be calculated. Additionally, we calculate two time series for the United States: one that considers only the federal minimum wage, and another that also accounts for state-level minimum wages.

Among these major economies, the highest minimum wages today are in India and several European countries (France, UK, and Germany), where they stand at around 50-55% of the average wage. The minimum wage levels in Brazil and Spain are slightly lower —around 45% of the average wage. In contrast, in the two largest economies in the world (China and the U.S.), current minimum wage levels are relatively low, at around 20-30% of the mean wage.

Interestingly, recent trends differ considerably between these countries. In India, relative minimum wages started to decline in the mid 1980s from a relatively high initial level. Despite this decline, India’s minimum wage (relative to the average wage) remains among the highest today. Meanwhile, we have seen a dramatic increase minimum wage levels among a number of European countries (Germany, Spain, UK) and Brazil in recent years. In contrast, the two largest economies (China and the U.S.) have experienced large drops in their levels of minimum wages. Notably, China’s minimum wage was around 55% of the average wage in the mid- 90s, but it has since fallen close to 20%. In the U.S., the federal minimum wage was never very high (around 45% of average wage in mid-1970s) but has fallen dramatically since then. It is important to note that in the U.S., this decline is considerably more muted when factoring in state-level minimum wages, which have become more significant over time. Today, with no nominal increase in the federal minimum wage for over 15 years, the federal minimum wage is largely nonbinding in the U.S., making state and local wage floors the only economically relevant ones in the American economy.

2.2 The minimum wage debate

The intellectual controversies surrounding the minimum wage go back more than a century. At the core of this debate is the question of what the appropriate model for the (low-wage) labor market is and whether government regulations can improve labor market outcomes.

The central debate on the minimum wage revolves around its employment consequences. A long-standing tradition in economics argues that a minimum wage set above the competitive level will reduce employment (Mill, 1848; Marshall, 1897; Stigler, 1946). These predictions were formalized within the neoclassical framework (Hicks, 1932). According to this framework, if labor markets are assumed to be 1) competitive with many firms passively accepting prevailing market-level wages, 2) characterized by free entry and exit of firms; and 3) composed of homogeneous workplaces from

³The minimum wage is often expressed relative to the median wage for this type of comparison (Dube, 2019a). Here we use the mean wage instead because it is easier to obtain for several countries. The ratio of the minimum wage to the median (or mean) wage is commonly referred to as the Kaitz index; the original formula also included the coverage rates by sector, but the relevance of that adjustment has greatly diminished as coverage rates today tend to be very high (Kaitz, 1970).

workers' perspective, then equilibrium employment and wages will be fully efficient.

With two inputs in production (capital and low-skilled labor), the effect of wages on labor demand in this framework can be expressed as follows (see the Hicks-Marshall rule of derived demand in Hamermesh (1995)):

$$\frac{\Delta \ln Emp}{\Delta \ln Wage} = -(1 - share_L) * \sigma_{K,L} - share_L * \eta \quad (1)$$

where $share_L$ is the share of labor in production, $\sigma_{K,L}$ is the elasticity of substitution between capital and labor, and η is the elasticity of output demand. The first term represents the substitution effect: since minimum wages make labor more expensive relative to capital, firms will substitute away from labor. The second term reflects the scale effect. In a competitive environment with free entry and exit, firms have a limited ability to absorb the minimum wage by lowering profits, so they raise prices to cover the increased expenses. This price increase leads to a reduction in consumer demand, which, in turn, decreases the scale of production and employment. The key insight from this framework is that introducing a binding minimum wage will have an unambiguously negative effect on labor demand, leading to lower employment.

Economic thinkers have long pointed out the unintended consequences of minimum wage policies (Sidgwick, 1886; Marshall, 1897). The minimum wage is often cited as a prime example of how market interference, even with the best of intentions, can do more harm than good (Stigler, 1946). However, there is also a long tradition of challenging the key prediction of the standard competitive framework. For instance, if firms have considerable market power in setting wages, then there are pre-existing market distortions. In such cases, minimum wages can serve as a second-best policy to alleviate these distortions and move closer to an efficient allocation. A leading example of this is the monopsony framework, where a single firm operates in the labor market. Historically, however, the monopsony model was thought to have limited empirical relevance, as few low-skilled labor markets were characterized by a single dominant employer with substantial wage-setting power (Stigler, 1946; Brown, 1999).⁴

Another critique came from institutionalist economists such as John Dunlop, Clark Kerr, Richard Lester, and Lloyd Reynolds, who challenged the neoclassical (or “marginalist”) view of the labor market. These economists questioned the notion of unfettered profit maximization and deemed as overly simplistic the view that firms always produce at a point where marginal costs equal the marginal product. Instead, they relied on institutional expertise, detailed case studies, and surveys of management and workers to paint a more empirically-minded picture of firm behavior (Freeman, 1989).

A key example is Lester (1946) who surveyed executive officers in various industries. The survey responses overwhelmingly emphasized the significance of current and prospective market demand for the firm's products as the key factor in determining its employment level. In contrast, managers rated

⁴The idea that wage-setting power can emerge in the presence of strong but imperfect competition began gaining attention in the 1980s. The search literature developed around the notion that while firms compete, search frictions create some job-specific surplus once workers and firms meet. Imperfect substitution between workplaces can also create wage-setting power, much like imperfect substitution between goods in output markets gives firms price-setting power (Card et al., 2018).

the importance of wages as quite low (perhaps surprisingly so). The correlation between concerns about the level of wages and the labor share of production was also quite weak, which goes against the prediction of the neoclassical framework as illustrated in the formula discussed earlier.⁵ Furthermore, [Lester \(1946\)](#) pointed out that even though the introduction of federal minimum wages reduced the North-South wage differentials in the manufacturing sector, it did not translate into an employment reduction at Southern firms (if anything, employment increased).

This approach to studying labor demand and the effect of wage shocks faced vehement opposition from neoclassical economists on various grounds ([Machlup, 1946](#); [Stigler, 1946](#); [Friedman, 1953](#)). First, it was argued that Lester’s survey employed an overly simplistic version of the neoclassical model. The model could be easily extended by incorporating factors such as non-pecuniary benefits and more dynamic decision-making processes on production capacity. Once these modifications are taken into account, the mapping between the neoclassical theory and the survey questions becomes less clear. Second, [Machlup \(1946\)](#) and [Friedman \(1953\)](#) argued that, similar to how an automobile driver who overtakes a truck, or a professional pool player who strikes a ball, cannot describe the calculations and the physics underlying their behaviors, business executives cannot explain what drives their decisions based on exact mathematical rules. Instead, good business executives, much like a good driver or a good pool player, can make the right decision by sensing profit-maximizing behavior. And if they did not know how to do that—the argument went—they would be replaced by more adept decision-makers. Based on this “as if” reasoning, [Friedman \(1953\)](#) argued that as long as a model has good predictive power, the realism of its assumptions does not matter.

By the 1960s, the neoclassical view of the labor market became the dominant approach among economists. Furthermore, the advancement of the human capital theory during the 1960s also led to the presumption that wages simply reflect skills, knowledge, and ability, leaving little scope for the government to effectively intervene in the labor market ([Osterman, 2011](#)). This perspective was further supported by a series of empirical studies showing a negative link between employment and wages using time series and plant-level data ([Hamermesh, 1995](#)). These findings underscored the predictive power of the neoclassical approach, and led to a consensus regarding the negative implications of minimum wages.

However, this consensus began to dissolve with the emergence of a series of more credible empirical evidence in the early 1990s. Riding the wave of the credibility revolution, a set of papers studied the impact of the policy by applying new and more credible empirical techniques, exploiting state-level variation in minimum wages, and applying difference-in-differences style estimators. Many of these papers found no ([Katz and Murphy, 1992](#)) or even positive employment effects ([Card, 1992a,b](#); [Card and Krueger, 1994b](#)), while others confirmed the negative effect of the policy ([Neumark and Wascher, 1992](#)). The previously established consensus broke, and the debate on the employment effects of the policy (re)started.

The initial reaction to the positive—or lack of negative—employment effects in response to the

⁵[Lester \(1946\)](#) argued that most businesses face a declining marginal cost between 70% and 100% of their production capacity, so wage determination is not driven by the marginal cost being equal to marginal benefit. Furthermore, the capacity of production is determined by market demand, and that production cost considerations play only a limited role in that.

policy was not particularly welcoming.⁶ The vehement opposition to the findings of [Card and Krueger \(1994b\)](#) reflect that a fundamental issue was at stake ([Leonard, 2000](#)). The lack of disemployment effects of the policy challenged a basic prediction of the neoclassical view of the labor market, and suggested that the workhorse model of labor markets might be missing some crucial mechanisms through which labor markets adapt to the policy change.

During the last three decades, there have been three major developments. First, the idea that core theories in economics can be tested and possibly rejected has gained greater acceptance within the discipline. As a consequence, the debate on the policy has shifted to a more empirical focus. Second, an extensive literature has emerged on the employment and wage effects of the policy, which we will review in the next section. A fruitful corollary is that this literature has become an arena in which various new empirical techniques have been developed and tested. Finally, new and richer models of low-wage labor markets have emerged (or were rediscovered). This also provided a theoretical basis to studying the impact of the policy in a richer environment, given the ambiguous predictions of the policy on employment and wages. To assess the empirical relevance of these new models, the empirical literature has started to explore outcomes beyond employment, which we discuss in Section 4.

3 The wage and employment effects of minimum wages

3.1 Wages, employment and labor demand

At the risk of understatement, there is an extensive literature studying the impact of minimum wages on employment. In this section, we will provide a guided tour of this literature using the very useful construct of the own-wage employment elasticity (OWE). This elasticity scales the observed (percentage) change in employment by the observed (percentage) change in wages. Therefore, it shows how responsive employment is to a change in wages paid to workers induced by the minimum wage policy—which is equivalent to the elasticity of labor demand in the standard neoclassical framework. However, in some models—such as ones with monopsony power—the labor demand curve is not well defined (or is not binding), even as the OWE remains an economically meaningful measure. For this reason, we use the more catholic employment elasticity terminology in order to avoid imposing model-based interpretations of the empirical evidence.

When quantifying employment responses, the main alternative to using the employment elasticity with respect to own wages is to use the employment elasticity with respect to the minimum wage. The latter approach is common in the literature (e.g., [Doucouliagos and Stanley, 2009](#); [Neumark and Shirley, 2022](#)). However, the OWE offers several important advantages.

First, scaling employment changes by wage changes makes the estimates more comparable across

⁶A good example of the critical response comes from an opinion piece penned by the Nobel laureate James Buchanan and published in the *Wall Street Journal*, where he wrote: “Just as no physicist would claim that ‘water runs uphill,’ no self-respecting economist would claim that increases in the minimum wage increase employment. Such a claim, if seriously advanced, becomes equivalent to a denial that there is even minimal scientific content in economics, and that, in consequence, economists can do nothing but write as advocates for ideological interests.” ([Buchanan, 1996](#))

studies and economic contexts. This is important because the same change in the minimum wage may be much more binding for certain economies, demographic groups, industries, or time periods. For this reason, we expect the effect of the minimum wage on employment to depend on the fraction of workers affected by the policy. For example, we anticipate that minimum wages will raise an average teen’s wage more than an average prime-age (25 to 54 year old) worker’s wage, since teens are relatively more likely to be low-wage workers. Consequently, even if there were *the same* (negative) employment effect of the minimum wage on all affected workers, the minimum wage employment elasticity for teens would be larger in magnitude than that for prime-age workers. In contrast, scaling employment changes by wage changes accounts for the variation in the bite of the policy across different groups, making the comparison of elasticities across groups much more meaningful.

Second, by focusing on the relative magnitudes and signs of the employment and wage changes, we are better positioned to calculate the policy-relevant trade-off. Policy-makers are interested in the consequences of the policy on workers’ wage and employment, but that is difficult to infer from the employment elasticity with respect to the minimum wage. For example, if a 10% increase in the minimum wage leads to a 1% reduction employment, it could mean an own-wage employment elasticity of -1 if the minimum wage increases wages by 1%. This would be considered high. However, the 1% reduction of employment could be also consistent with an own-wage elasticity of -0.1 if a 10% increase in the minimum wage leads to a 10% increase in wages. In the latter case, job loss would be considered small by most, and might be welfare maximizing (see [Lee and Saez \(2012a\)](#)).⁷

Third, as we explained in the previous section, the minimum wage debate centers around the relationship between employment and wage changes, and their alignment with the prediction of the neoclassical framework. In theoretical discussions, it is often assumed that the link between minimum wage and actual wages is one-to-one, but that assumption rarely holds up in real-world data. Therefore, an additional step is required to shed light on the link between wages and the minimum wage. Formally, this could be conceptualized as a two-step instrumental variable (IV) procedure. Suppose we are interested in estimating the causal effect of wage changes on employment, formally:

$$\Delta \ln emp = \beta \Delta \ln wage + \epsilon$$

We can use the change in the minimum wage policy denoted as ΔMW , which may be either a binary or continuous measure, as an instrument for the change in wages, and implement the following two-step procedure.

$$\begin{aligned} \text{1st stage:} \quad & \Delta \ln wage = \alpha \Delta MW + \varepsilon \\ \text{2nd stage:} \quad & \Delta \ln emp = \beta^{IV} \overline{\Delta \ln wage} + u \end{aligned}$$

where $\overline{\ln wage}$ is the predicted wage change ($\alpha \Delta MW$) and β^{IV} identifies the OWE. In the first stage, we estimate the link between the minimum wage and the actual wage earned by a worker. The

⁷The usual argument in favor of using the employment elasticity with respect to the minimum wage is that the policy-maker can change the minimum wage, but not the wages directly. However, without knowing the sign and magnitude of the wage effects, the policymaker can only evaluate the potential employment costs of the policy without knowing the benefits coming from higher wages.

second stage estimate then provides the link between the wage and employment. The reduced-form relationship between the instrument and the employment is given by the following equation.

$$\Delta \ln emp = \gamma \Delta MW + v$$

If ΔMW is the change in log minimum wage, then γ identifies the employment elasticity with respect to the minimum wage; if it is a binary policy change measure then γ is the reduced form policy effect on employment. This formulation clarifies that the OWE can be thought of as an indirect least squares estimator:

$$\beta^{IV} = \frac{\gamma}{\alpha}$$

With a binary ΔMW , β^{IV} is also the Wald estimate of the LATE (e.g. [Angrist et al., 2000](#)).

$$\beta^{IV} = \frac{\gamma}{\alpha} = \frac{E(\ln emp | \Delta MW = 1) - E(\ln emp | \Delta MW = 0)}{E(\ln wage | \Delta MW = 1) - E(\ln wage | \Delta MW = 0)} \quad (2)$$

Looking at the OWE through the lens of an IV approach means that the considerations of the validity of the instrument, the role of the compliers, and the measurement error in wages all apply here. For example, whether the minimum wage shock is a valid instrument or not depends on the following assumptions: 1) relevance—there is a strong enough first stage, i.e., there is a clear association between the minimum wage change and change in the average wage of the relevant group,⁸ and 2) exclusion restriction—the minimum wage change only affects employment through wages.

Furthermore, we can also gain insights from the IV literature by recognizing β as a Local Average Treatment Effect (LATE). Consider the case of a binary instrument ΔMW representing a policy change. Here, the treatment is an increase in log wages for “compliers” induced by a minimum wage policy event, and the outcome is log employment.

Note that if we add a set of individuals completely unaffected by the policy to the sample, both the first-stage (α) and the reduced form relationship (γ) weaken, but the ratio of α and γ still stays the same, since the complier share enters (multiplicatively) in both the numerator and the denominator of equation (2). This shows that as long as the relevance requirement is satisfied, changing the share of affected workers (say by adding more “never-takers” who are unaffected by minimum wages to the sample) does not alter the fact that β^{IV} is still a valid LATE for the complier sub-population. In contrast, the employment elasticity with respect to the minimum wage (i.e., γ divided by the change in log minimum wage in the case of a binary ΔMW) will be sensitive to adding nontreated workers to the analysis.

While recognizing the advantages of the OWE, it is worth noting some important limitations. First, while measuring the impact of the policy on wages helps scale and interpret the employment

⁸Note that the IV interpretation suggests that the first stage estimate of α needs to be strong enough to avoid a weak instrument bias when estimating the OWE, which is a higher standard than statistical significance at conventional levels.

effects, such a scaling can be imperfect. For instance, changes in non-pecuniary benefits should also be taken into account, both from the policy perspective and the economic debate around the policy. Policymakers presumably care about workers’ welfare, and not only about their wages *per se*. A significant decrease in non-wage benefits could fundamentally alter the trade-off between employment and the total (wage and non-wage) compensation of workers. Second, the neoclassical debate focuses on the link between labor cost and employment. However, if firms respond to a minimum wage shock by cutting non-wage benefits, then the change in effective labor costs will be limited. In that case, even if the employment elasticity with respect to *labor costs* is large in magnitude, the employment elasticity with respect to *wages* could be small. We will discuss the evidence on non-wage amenities in Section 4.1.2.

3.2 Empirical methods to study the impact of minimum wage policies

Roadmap. This section covers the main empirical strategies developed to assess the impact of minimum wages on employment and wages (although many of the methodological lessons are relevant for studying other outcomes as well). Until the early 1990s, the predominant empirical strategy to estimate the impact of minimum wages relied on time series analysis, which struggled to isolate causal effects due to its reliance on extrapolating trends without suitable control groups. The limitation of that literature has been discussed in previous reviews, and the presence of serious publication bias has been documented (see Chapter 6 in [Card and Krueger \(1995b\)](#)). The early 1990s witnessed the emergence of the “new minimum wage literature”—which developed and applied various difference-in-differences style estimation strategies. Needless to say, relying on an explicit control group—instead of only extrapolating past trends—to form a counterfactual was a major innovation in the literature. Still, the choice of control groups is not always straightforward; as we will see, such choices have sometimes been contentious in the literature. Furthermore, the literature is sometimes unclear on exactly what economic object is identified by the proposed difference-in-differences empirical strategy.

3.2.1 Exploiting local variation in the level of minimum wages

The effect of minimum wage policies can be studied most directly by examining local variations in the policy. Since the emergence of the “new minimum wage” literature, the United States has been a fertile ground for research as there is considerable variation in the level of statutory minimums across states and over time. More recently, a new literature has emerged that studies the effect of the policy exploiting variation at an even more granular level (cities or counties).

Case studies. A classic approach in the literature evaluates the impact of a single policy change. The best-known example of such a case study in the modern literature is [Card and Krueger \(1994b\)](#), which evaluated the impact of a 19% increase in the minimum wage that went into effect in New Jersey on April 1st, 1992. Card and Krueger use a first-differenced regression, where the change in employment at fast-food chain restaurants in New Jersey is compared to the change in similar restaurants from neighboring counties in Pennsylvania. More recently, this case study approach has

re-emerged in the context of studying city-level minimum wage changes: [Jardim et al. \(2022\)](#) study the impact of the policy in Seattle and [Karabarbounis et al. \(2022\)](#) in Minneapolis by applying the synthetic control method.

By focusing on one specific event, researchers can often rely on more detailed data sources and implement a clearer identification strategy. For example, [Card and Krueger \(1994b\)](#) collect survey data from fast food chains, as detailed firm-level data on employment and wages were not available in the early 1990s. Subsequently, in responding to the critique by [Neumark and Wascher \(2000\)](#), they also access state-level administrative payroll records from New Jersey and Pennsylvania to further improve data quality ([Card and Krueger, 2000](#)). In a similar vein, [Jardim et al. \(2022\)](#) and [Karabarbounis et al. \(2022\)](#) use administrative data on hourly wages that are only available for a few states in the United States. By focusing on a specific event, we may be able to better understand the context of the minimum wage change and consider the biases induced by shocks that coincided with the policy implementation. For instance, in the context of Seattle’s minimum wage ordinance, there has been an extensive discussion on how Seattle’s exceptionally strong labor market boom may contaminate the estimated employment effects ([Dube and Lindner, 2021](#)).

TWFE-log(MW). At the same time, an individual case study has the obvious limitation that it is just that—a single case. Statistical inference for single cases is challenging, and findings from cases have limited generalizability. Being able to go beyond individual cases, therefore, has obvious appeal.

Parallel to the case-study approach, an extensive literature emerged applying a two-way fixed effect empirical strategy that involves unit- and time-specific fixed effects to control for unobserved factors (e.g., [Neumark and Wascher, 1992](#); [Orrenius and Zavodny, 2008](#); [Allegretto et al., 2011](#); [Meer and West, 2016](#); [Neumark et al., 2014](#)). This is often thought to be a generalization of the case-study approach which pools various individual cases; but as we will see later, this is not quite right, as this approach introduces additional (sometimes opaque) assumptions.

Usually, the following regression is estimated in some form, which we will refer as the TWFE-log(MW):

$$y_{st} = \beta \ln(MW_{st}) + X_{st}\Lambda + \gamma_s + \delta_t + \nu_{st} \quad (3)$$

where y_{st} is typically the employment rate or average wages in state s at time t , in levels or in logs, and $\ln(MW_{st})$ is the log of the binding minimum wage (greater of state or federal minimums). As we shall see below, in most applications, the outcome variables are restricted to certain demographic groups such as teens or young workers, as those groups are the most exposed to the policy. The regression includes state fixed effects (γ_s), time fixed effects (δ_t), and typically time-varying, state-specific control variables (X_{st}).

The inclusion of state and time effects implies that within-state variation in the level of minimum wages over time is used for identification. In principle, this allows researchers to better isolate the causal impact of minimum wage changes by comparing changes within the same state rather than across different states with potentially confounding differences. A key advantage of this design is that

it pools together many minimum wage changes. Peculiar shocks coinciding with one particular event are averaged out, and so those shocks have less influence on the estimates. The design also leads to an estimate for an “average” minimum wage shock with greater external validity than the findings based on specific events. Additionally, since the TWFE is just a panel regression, standard inference is available when variation across many minimum wage changes is exploited.

However, it turns out that this design also has some serious problems when used to study minimum wages. One issue is the sensitivity of results to how unobserved time-invariant heterogeneity is controlled for. For example, [Neumark and Wascher \(1992\)](#) show that the estimated negative impact in the TWFE specification becomes positive in a first-differenced (FD) specification. [Cengiz et al. \(2019\)](#) (see Appendix G) also find similar differences between the TWFE and FD estimates for the 1979 and 2016 period (though not after 2000). Such differences are surprising since the two models should converge to the same estimated coefficients in large samples (with many states and time periods). [Cengiz et al. \(2019\)](#) linked it to violations of parallel trends in the early part of the sample. The large discrepancy also highlights that specifications need to be carefully scrutinized and additional falsification tests are required to pick the better estimation model.⁹

Second, a contentious issue emerged about what to control for in the regression equation (3). The original study by [Neumark and Wascher \(1992\)](#) on the impact of the minimum wage on teens includes estimates with and without the proportion of teens in school. The relationship between minimum wage and employment is positive without controls, whereas the sign flips and becomes negative when controlling for that variable. Controlling for the proportion of teens in school might be simply a “bad control” because it picks up a key mechanism through which the minimum wage affects teen labor markets. In any case, the sensitivity of the estimates could also reflect that minimum wage shocks are implemented parallel to other policy changes or in an unusual economic environment. In the presence of such shocks, understanding how those shocks affect low-wage labor markets is essential, and simply controlling for those shocks might not be enough. As we will also see below (Section 3.2.2), controlling for covariates in a difference-in-differences context requires care when the effect of treatment is heterogeneous.

In the follow-up literature, the debate about controlling for time-varying observable factors largely moved to how to control for the unobserved time-varying differences across states, e.g., by including state-specific time trends ([Neumark et al., 2014](#); [Allegretto et al., 2011](#)). [Manning \(2021\)](#) reviewed this debate and estimated seven different TWFE specifications using the Current Population Survey data between 1979 and 2019. He found the employment estimates to be sensitive to the inclusion of state-specific time trends: the estimated teen employment falls without time trends and increases when variation in state time effects is taken into account. Concern about violations of parallel trends led to the use of other ways to form counterfactuals such as synthetic control, factor models, or border

⁹The TWFE and FD estimators are the same if $T = 2$. If $T > 2$, TWFE and FD will not be identical, but both are unbiased and consistent according to standard assumptions. One source of discrepancy could be the presence of lagged effects, which the FD can miss. However, the use of distributed lags does little to resolve the discrepancy, as demonstrated in [Cengiz et al. \(2019\)](#). When TWFE and FD specifications yield different results, it is difficult to determine which one is more accurate. As a statistical matter, if T is large, and N is small, FD (especially with lags) is likely preferable, as TWFE is sensitive to small violations in assumptions (we will see more on this). If T is small, N is large: relative performance will depend on the autocorrelation of the error term (for more on this, see [Wooldridge \(2012\)](#), page 490).

discontinuity (Dube et al., 2010; Neumark et al., 2014; Totty, 2017). However, the TWFE panel regressions were somewhat of a black box, making it difficult to diagnose the problems.

In addition to these issues, a growing number of studies have highlighted the fragility of TWFE estimates under staggered adoption and heterogeneous treatment effects (De Chaisemartin and d’Haultfoeuille, 2020; Callaway and Sant’Anna, 2021; Sun and Abraham, 2021; Goodman-Bacon, 2021). Although the TWFE estimates an average of the treatment effects, the weights are sometimes negative, creating possibly serious bias. In order to address these concerns, Cengiz et al. (2019) proposed an event study design. As we shall see, an event-based approach also helps shed light on exactly what variation is being used for identification, allowing us to better see where violations of parallel trends occur.

Event study. A pooled event study is the proper generalization of the individual case study approach pioneered by Card and Krueger (1994b). This can be seen transparently using the stacked event study approach proposed by Cengiz et al. (2019). They study all sizable minimum wage changes between 1979 and 2016.¹⁰ For each prominent state-level minimum wage change, h , they create an event-specific dataset, which includes observations from the treated state $s^*(h)$ around the policy change (e.g., between three years before, \underline{T} , and five years after \overline{T} the date of the policy change). The event-specific dataset also includes observations from all the “clean control” states that did not experience any major changes in minimum wages over the event window. Therefore, just like for an individual case study, a single event-specific dataset includes one treated state, where the counterfactual outcome is obtained by averaging across all untreated (control) states. Here τ represents time relative to the event date, so $\tau = 0$ is when the minimum wage is raised in state $s^*(h)$. Subsequently, all the event-specific data are stacked. Once this stacked dataset is created, a number of approaches can be applied.

First, one can assess the impact of the policy across all events, by running the following regression:

$$y_{hs\tau} = \beta \times \mathbb{1}[\tau \geq 0] \times D_{hs} + \Omega_{hs\tau}\Lambda + \gamma_{hs} + \delta_{ht} + \nu_{hst} \quad (4)$$

where $y_{hs\tau}$ is the outcome in the dataset belonging to event h in state s and at event time τ . The regression includes γ_{hs} , event-specific state fixed effects, and τ_{ht} , event-specific time effects, ensuring that all the identification comes from within an event. The treatment variable equals one for the treated state and zero otherwise, $D_{hs} = \mathbb{1}[s^*(h) = s]$. Hence, the coefficients β show the change in outcome in the treated states relative to the change in outcome in the control states. The Ω_{hst} are covariates. For example, Cengiz et al. (2019) control for small minimum wage changes that took place in the control states.¹¹ Finally, standard errors should be clustered at the state level; this automatically takes into account that some observations from a state may show up more than once in the stacked dataset.

¹⁰Cengiz et al. (2019) define prominent state minimum wage changes as being more than \$0.25/hour and affecting more than 2 percent of the workforce; there were 138 such events. Note that some of these events are part of multiple year phase-ins of large minimum wage changes. In the our empirical exercise below, we combine these multiple-year phase-in events into one, reaching 60 major events.

¹¹Cengiz et al. (2022) offer a refinement to this approach that discards units with small changes from the control group. This is preferred (when feasible) because it requires fewer assumptions.

It is worth highlighting that this event study regression aligns events by event time (τ) and not by calendar time (t), so it is equivalent to a setting where events happen all at once. Moreover, by dropping all states with any events within the 8-year event window from the “clean” control set, we guard against bias due to heterogeneous treatment effects. Therefore, the negative weighting issues from staggered implementation (see, e.g., [De Chaisemartin and d’Haultfoeuille \(2020\)](#), [Goodman-Bacon \(2021\)](#), [Borusyak et al. \(2024\)](#)) do not arise here. The proposed regression estimation provides us with a (positively) weighted average impact of the minimum wage across events. These weights are proportional to the variance of the treatment; therefore, the simple stacking approach produces a variance-weighted ATT (VWATT) ([Gardner et al., 2024](#); [Dube et al., 2023](#)). As shown in [Dube et al. \(2023\)](#), one can obtain exactly the same estimate as the stacked approach by using the local projections DiD (LP-DiD) regression in the original panel (i.e., not stacked) dataset. LP-DiD simply regresses:

$$\overline{\Delta y_{st}} = \beta \times D_{st} + \delta_t + \nu_{st} \quad (5)$$

where $\overline{\Delta y_{st}} = \frac{1}{6} \sum_{k=0}^5 (y_{s,t+k} - y_{s,t-1})$ is a post- minus pre-treatment difference in outcome averaged over the post-treatment window, and $D_{st} = 1$ equal to one if state s is newly treated at time t .¹² The estimation sample only includes observations that are either newly treated on date t (i.e., $D_{st} = 1$) or “clean control units” as defined above. Below, when we show empirical results for the event study analysis, we implement it using equation (5).¹³

Second, being clear about the exact treatment dates, the event window, and the admissible control units (as is the case with the stacked dataset) also allows us to naturally obtain event-by-event estimates. This can be done by separately estimating a β_h for each event, which is equivalent to replacing β with β_h in equation (4), or changing the sample to include only one treated unit at a time in equation (5). However, given that each β_h is identified from a single treated unit, the usual approach of clustering the standard errors at the state level does not produce the correct statistical inference. This can be addressed by using the [Ferman and Pinto \(2019\)](#) standard errors when calculating event-specific estimates. The event-by-event estimates can then be used to transparently study whether the impact of the policy is heterogeneous and/or driven by a few outlier events. One can also explore the impact of different weighting schemes when aggregating events (e.g., instead of weighting by the sample size, variance, etc., equal weights can be applied).

Third, this event study approach allows us to cleanly estimate the changes in outcome around the timing of the policy change by running an event-study type regression:

$$y_{hs\tau} = \sum_{k=\underline{\tau}}^{k=\overline{\tau}} \beta_k \mathbb{1}[\tau = k] \times D_{sh} + \Omega_{hs\tau} \Lambda + \gamma_{hs} + \tau_{h\tau} + \nu_{hs\tau} \quad (6)$$

¹²We use the following convention when defining the treatment variable D : if D has a time subscript, it equals to one only on the date the state first receives treatment, otherwise it equals to zero. If the treatment variable has no time dimension, as in equation 4, then it is only a function of the state, s , and is equal to either zero or one in all periods.

¹³As also shown in [Dube et al. \(2023\)](#), it is straightforward to re-weight the observations to produce an (equally weighted) estimate of the ATT; this is available in the Stata and R code of `lpdid`. The equally weighted version of the LP-DiD is numerically equivalent to the DiD estimate of [Callaway and Sant’Anna \(2021\)](#). We also demonstrate this equal weighting below in our empirical implementation.

Analogously, separate regressions could be estimated for each k using equation 5 with long-differenced outcomes to recover the same dynamic responses:

$$y_{s,t+k} - y_{s,t-1} = \beta_k \times D_{st} + \delta_t + \nu_{st} \quad (7)$$

where the estimation sample only uses observations that are newly treated or are clean controls.

Looking at the coefficients before the reform allows us to test for pre-existing trends, which is the standard approach to assess the parallel trend assumption—the treated and control states would have evolved similarly in absence of the minimum wage change. If there are concerns about the presence of pre-existing trends, one can balance the treatment and control by matching on covariates, or by implementing a synthetic control approach (e.g., see [Allegretto et al. \(2017\)](#) for such an application studying the impact of city-level minimum wages). Alternatively, matching and synthetic control methods can be combined as in [Kellogg et al. \(2021\)](#).

Until relatively recently, distributed lag versions of TWFE were thought to identify dynamic responses to treatment using a difference-in-differences design, similar to equations (6) or (7). However, this is not the case, and the dynamics implied by a distributed lag TWFE model can be quite misleading due to issues like negative weighting ([Sun and Abraham, 2021](#)). This further underscores the importance of using a proper difference-in-differences event study design instead of TWFE-based strategies.

Event-study vs. TWFE-log(MW): an empirical exploration. A key difference between the event-study and the TWFE-log(MW) specification is that event studies focus on employment changes within the event window, while the TWFE-log(MW) compares employment changes across different time-periods, even if these periods are far apart. The TWFE assumes that the parallel trends assumption holds across the full sample; this is a stronger assumption than those made by an event-study design, which assumes parallel trends holds within the event window. These properties are important to understand for why TWFE-log(MW) produced sensitive results, especially for samples that include the 1980s and 1990s. Below we present evidence to underscore this issue and shed some light on the parallel trends assumption.

Much of the variation used for identifying minimum wage effects in the U.S. comes from the comparison between two groups of states: 35 states that have raised their minimum wages at some point between 1980 and 2019 (we will call them the “ever-treated” group), versus 15 that did not (the “never-treated” or the “control” group). To be clear, there are differences in timing in when these 35 states raised their minimums. However, in practice they mostly did so in 3 waves: a small wave in the late 1980s, and two much larger waves in the 2000s. These waves emerged as a natural response to prolonged periods in which the federal minimum wage remained stagnant. As a result, most of the minimum wage “events” happened in the 1998-2006 and post-2013 periods, with a smaller set of events in the late 1980s and a handful in the early 1990s including the famous 1992 New Jersey increase studied by [Card and Krueger \(1994b\)](#). Panel A of Figure 2 plots the nominal minimum wages in the two groups of states, showing the clear patterns of these three waves.

Since much of the identifying variation comes from comparing these two groups of states, a natural question is how similar these groups are. As it turns out, the assumption that the “ever-treated” states have followed parallel trends as compared to the “never-treated” states seems much more plausible since the late 1990s. In contrast, for the period between early 1980 to mid-1990s, labor markets in these two groups of states greatly diverged. We can see this in Panels B and D of Figure 2 which plot the overall average weekly earnings (in logs) and per-capita employment for the two groups of states.¹⁴ As compared to the “never-treated” states, the “ever-treated” states’ employment rate rose sharply in the 1980s, and then fell back in the 1990s. What is striking is that this happened even though the minimum wage barely differed between these two groups over this period (see Panel A). This pattern strongly suggests a violation of parallel trends in the potential outcomes for employment, because: (1) in general we do not expect *overall* outcomes to be move very much from minimum wages given the small share of workers affected by the policy; (2) there was not much change in the relative minimum wage between these two groups; (3) but even so, average earnings rose more strongly in “ever treated” states in the late 1980s.¹⁵ In contrast, after the mid-1990s, overall employment and earnings trends were much more alike between the two groups of states. As the figure shows, this is also precisely the period in which we see major differences in the minimum wage between the “ever treated” and “control” states.

Of course, most minimum wage analysis is not based on overall employment (with a few exceptions like Meer and West (2016)). Turning to the low-wage labor market, Panels C and E of Figure 2 plot log earnings and employment in the restaurant sector, a much studied group.¹⁶ Here we see that restaurant employment fell greatly in the “ever treated” states between the early 1980s and the late 1990s, even though the relative minimum wages were virtually identical in 1980 and 1995, consistent with a violation of parallel trends. More generally, there is very little relationship between the timing of the minimum wage changes and the changes in employment.¹⁷

In contrast, over the post-1998 period, we see that the gap in restaurant employment as well as overall employment to be highly stable between the two groups of states—even as minimum wages and average restaurant earnings varied greatly. The post-2013 period, in particular, led to a very durable gap in wages between the “ever-treated” and “never-treated” states, but without any visible gap in restaurant employment. Concretely, between 2000 and 2019, average restaurant earnings increased by 0.07 log points more in the “ever-treated” states, while log employment increased by 0.02 log points more.

We can use these state-by-year data to shed light on how estimates based on the popular TWFE-log(MW) regression with long panels are sensitive to violations of parallel trends in the 1980s

¹⁴This is based on annual data from Quarterly Census of Employment and Wages (QCEW).

¹⁵The regional labor market differences were recognized at the time. For example, Freeman (1990) analyzed the effects of this boom in states such as California, New York, and Massachusetts (dubbed the “Massachusetts Miracle”) on young Black workers.

¹⁶The data from the QCEW for restaurants is available based on NAICS codes only after 1990; SIC-based data is available through 2000. For restaurants, the differences in employment levels between SIC and NAICS is modest. We construct a state-level NAICS prediction for the pre-1990 period as described in Appendix D.

¹⁷While the federal minimum wage fell in real terms over the 1980s, note that the average restaurant earnings were similar in the “ever-treated” and “never-treated” states in the early 1980s; this means that the “bite” of the federal minimum decline did not differ very much across these two groups; and hence the relative reduction in restaurant employment is unlikely to reflect federal changes.

and 1990s. To start, we present event study estimates for restaurant earnings and employment from Equation (5) in Table 1, Panel A, column 1.¹⁸ These are based on 60 combined events (which combines successive minimum wage increases into a single event). The details of the event construction are provided in Appendix E; Appendix Figure A1 plots the events by state and time; the number of control units, the duration of the post-treatment windows, and the size of the increase by the year of the event are reported in Appendix Table A1. Averaged over a maximum of 6 years following the initial increase, treated states saw a rise in minimum wages of 0.26 log points, a rise in restaurant earnings of 0.03 log points, and a fall in employment of -0.001 log points (statistically indistinguishable from zero). This yields a small OWE of -0.03 (s.e.=0.15).¹⁹ Focusing on the 46 events after 1998 yields an OWE of 0.15 (s.e.=0.14). (See Appendix Figure A4 for the event study plots.) The other columns present outcomes for other groups of workers than in restaurants, including for teens, as well as a demographic prediction-based group (“high recall”) that captures most minimum wage workers. As we can see, the event-based OWE estimates for these groups are small or positive, falling between -0.07 and 0.34.

In contrast, when we estimate static TWFE-log(MW) regressions in the 1980-2019 sample, we obtain estimates of 0.28 for restaurant wage and -0.15 for restaurant employment, leading to an OWE of -0.52 (s.e.=0.35). Importantly, when we estimate distributed lag versions of this regression (with 5 annual lags and 2 leads), the implied dynamic responses from the distributed lags regression also look very different from the event study estimates, as shown in Appendix Figure A2. In contrast, as Table 1 also shows, the TWFE estimates in the post-1998 period are more similar to the event study estimates, with an OWE of 0.5 (s.e.=0.41). For the 1980-2019 sample, the “long run” effect on restaurant employment is large, negative, and statistically significant (panel B of Appendix Figure A2); but for the post-1998 sample, it’s very close to zero (Panel D). Other low-wage groups such as teens (Table 1, column 4) follow a similar pattern, with the TWFE estimates from the 1980-2019 sample indicating larger dis-employment than either event study estimated or TWFE estimates from 1998 and later.

Why do the TWFE estimates differ so much from event study estimates when including data from 1980s and 1990s? After all, there is not much geographic variation in the minimum wage during this earlier period, as we have already seen. The reason is because inclusion of the early period affects the estimation of the state fixed effect—and hence the “baseline” employment—in somewhat non-transparent ways. As we add older pre-treatment observations, the regression adds comparisons between the post-treatment period and the newly-added older observations (e.g., from the 1980s).

As a demonstration, we can show how adding older data affects the estimates in an otherwise clean event study-type design. To begin, we consider the 2010-2019 period with a sharp increase in the minimum wage. Here we take the same 35 “ever-treated” states ($D_s = 1$) and 15 “never-treated”

¹⁸The one difference from Equation 5 is that we have 6 years as a *maximum* post-period, which for certain events may be truncated due to either federal minimum wage increases or end-of-sample considerations. See Appendix E for details.

¹⁹Recall that these estimates based on equation (5) represent a variance-weighted ATT. Instead, if we equally weight each event by re-weighting, the OWE is 0.02 (s.e.=0.16) as shown in Appendix Table A3, panel A, column 5.

states ($D_s = 0$) and consider a regression:

$$y_{st} = \beta_\tau \times Post_t \times D_s + \gamma_s + \tau_t + \nu_{st} \quad (8)$$

Here $Post_t = 0$ if the year is 2010-2013, while $Post_t = 1$ if year is 2014-2019. 20 of the 24 states experiencing a prominent increase in minimum wage in the 2010-2019 period first did so in 2014 or later, which is why we use 2014 as the cutoff point for $Post$. This formulation is very close to the event study estimates we saw from Panel A, column 1 of Table 1, except that to keep things even simpler, we are using an early event-start date as the common event-start date for all 35 states to avoid any staggered adoption issue. For this reason, we call this a “quasi event study”. The wage and employment estimates can be thought of as being for the “intent to treat” since some of the post periods are not treated; but the OWE estimates can be interpreted just as before. As first shown in Column 1 of the Appendix Table A2, the log minimum wage gap between the two groups rises by 0.15 points (goes up from 0.05 to 0.20) in the $Post$ period, a large increase. Consequently, log restaurant earnings increase by 0.04, while log employment increases by 0.005, with an OWE of 0.13 (s.e.=0.27), close to the post-1998 event study estimate of 0.15 (s.e.=0.14) in Table 1.

Now consider the same regression, but expand the Pre period to include all years between 1980-2013 when the gap in log minimum wage between the “ever treated” and “never treated” groups was smaller than the original pre-treatment (2010-2013) gap of 0.05. This adds in years 1980-1988, 1990-2000, and 2009 to the Pre period. Doing so does not substantially change the treatment effect on log minimum wages (estimate changes from 0.15 to 0.17). However, now restaurant wage effects are somewhat larger and employment effects are much more negative, leading to an OWE of -0.76 (s.e.=0.36). The reason for this is obvious from the inspection of Figure 3: the “ever-treated”-“never-treated” gap in restaurant employment was much larger in the 1980s and early 1990s. Simply expanding the Pre period in the event study by adding data from the 1980s and 1990s leads to an OWE much closer to the TWFE estimate of -0.52 from the 1980-2019 period (Table 1. This exercise also offers intuition behind why the TWFE regression elasticity becomes much more negative from adding data from 1980 to 1999—a period with relatively few minimum wage events. Namely, it affects the baseline employment through the estimation of the state fixed effects.

Although in this case Figure 3 transparently indicates a likely failure in the parallel trends assumption during the 1980s and 1990s, a distributed lag version of the TWFE regression does not necessarily clarify that point. This is related to the fact that the weights put on different events by the TWFE model can be possibly negative; this, combined with some violations in parallel trends, can lead to potentially spurious impulse response estimates in the distributed lag TWFE model. For example, the implied dynamics of restaurant employment around the event date according to the TWFE model (see Appendix Figure A2) look very different from those in the event study (see Appendix Figure A4).

Importantly, the problem with the TWFE model is not primarily a matter of using continuous versus discrete treatment measures. For example, take exactly the same 60 events D_{st} used in the event study analysis, and cumulate those over time within the state as follows: $\tilde{D}_{st} = \sum_{k \leq t} D_{sk}$. Now estimate a TWFE regression (static or distributed lag) using \tilde{D}_{st} using the 1980-2019 sample. This TWFE-binary regression generates similarly spurious effects as the TWFE-log(MW) regression, and

the associated impulse response looks very different from the event study estimates we saw before, with an OWE of -1.43 (see Column 7 of Appendix Table A3, and Appendix Figure A3).²⁰ This is because both regressions use all possible long difference comparisons over the full 40 year window, as opposed to the event study estimates from equations (4) or (5) which make comparisons only within the actual event window.

To further clarify the nature of the bias in the TWFE estimates, in Appendix B, we conduct a Monte Carlo simulation with two treatment cohorts and staggered adoption. The simulation shows that when there are violations of the parallel trends assumption early on in the sample, TWFE distributed lag models can produce highly misleading impulse responses, even when the true treatment effects are constant and equal to zero. Importantly, in our simulations, these violations occur well before the treatment events and outside of the event window. As a result, while the TWFE model’s stricter parallel trends assumption (that it holds across the full sample) leads to biased conclusions, modern difference-in-differences event study approaches—such as Callaway and Sant’Anna (2021), Cengiz et al. (2019) or Dube et al. (2023)—yield unbiased estimates and correctly show a null effect.

Overall, our take-away from the recent literature is that an event-based difference-in-differences design (which has become increasingly popular) is a much more promising approach and should be the default among researchers using a DiD approach. Estimates from a TWFE panel regression with log minimum wage have proved opaque and fragile, and we think they are unlikely to yield convincing evidence going forward. This should also be kept in mind when reviewing the older employment evidence presented in Section 3.3, much of which has been based on the TWFE model. It is also good to recognize that much of the debate around the sensitivity of results to specifications (as recounted in Manning (2021) took place because the early sample with substantial state-level variation (i.e., 1980s and 1990s) had some serious violations in parallel trends, making identification genuinely more difficult. Thankfully, these issues have become much less pronounced over the past 25 years.

Finally, this discussion has focused on the U.S., which is unique in having about half the country with no change in minimum wages for over 15 years, while some parts have enacted very large increases. At the same time, the issues around event-based estimates and the TWFE are broadly relevant also for other contexts with sub-national minimum wage variation, including Canada, China, and India.

3.2.2 Other considerations for the comparability of treatment and control groups

In this section, we discuss further considerations about adding controls to the regression and choosing the treated and control units in the regression.

Controlling for covariates. It is common to include covariates such as state-level unemployment rate (Neumark et al., 2014; Allegretto et al., 2011), housing price index (Clemens and Wither, 2019), state-level income per capita (Clemens and Strain, 2021), or pre-treatment sectoral and demographic characteristics of areas (Dube and Lindner, 2021) in the estimating equations. Some of these are contemporaneous (time-varying) controls, while others are pre-determined.

²⁰For an example of a recent paper which implements such a TWFE-binary regression with distributed lags over a long horizon, see Jha et al. (2024).

In general, as common as it is, the inclusion of variables that could be directly affected by the policy (e.g. state unemployment rate) should be avoided. It is better to construct covariates of state-level business conditions that exclude low-wage workers (for example, [Cengiz et al. \(2022\)](#) use the employment rate of those who are unlikely to be minimum wage workers as a control in one specification).

However, even when a variable is not directly affected by minimum wage treatment, there are subtleties about how best to control for it in estimation (see the discussion in [Roth et al. \(2023\)](#)). A common approach is to add time-invariant (e.g., pre-determined) covariates (X_j) interacted with time effects. However, this can bias estimates when the effect of the minimum wage varies by the level of X_j , because the $X - by - t$ interactions would constitute “bad controls,” as they are also mediators. However, this problem can be addressed by “regression adjustment,” where the coefficients on the $X - by - t$ interactions are estimated using the clean control sample only.²¹ Alternatively, researchers can use propensity score reweighting of control units to match the averages of X in treatment units.

Time-varying covariates can be handled in an analogous fashion. If the effect of the minimum wage treatment is heterogeneous by a time-varying control X_{jt} , simply including X_{jt} as a covariate in the regression poses a problem for similar reasons as described above. For example, if the minimum wage effect on employment differs during periods of high aggregate unemployment, then the inclusion of aggregate unemployment (or unemployment rate of individuals unlikely to be minimum wage workers) as a covariate can lead to a biased estimate, since the unemployment variable is also a mediator. Regression adjustment is one option. As an illustration, column 6 of Table [A3](#) reports estimates from equation [5](#), but uses regression adjustment to additionally control for the change in employment rate of a “low-probability” group (whose construction using demographic predictors is explained later in this section). One can also use a propensity score re-weighting approach (see [Caetano et al. \(2022\)](#)).

Border discontinuity. Since [Card and Krueger \(1994b\)](#)’s comparison of the border region of Pennsylvania to New Jersey, the minimum wage literature has sometimes considered nearby areas to be more reliable control groups. [Dube et al. \(2010\)](#) generalized this approach by considering all contiguous border county pairs (BCP) and estimating a regression with pair-specific time effects. This essentially only considers variation within each BCP, washing out variation between pairs. The analysis in [Dube et al. \(2010\)](#) was based on using log minimum wage and a modified TWFE model, TWFE-BCP (where the time effects are interacted with BCPs). However, the findings—strong earnings effects and small employment effects in the restaurant sector—are similar when we implement the BCP design using a stacked event study regression (see Appendix G of [Cengiz et al. \(2019\)](#)). Appendix G of [Cengiz et al. \(2019\)](#) also shows that the TWFE-BCP estimates are not sensitive to the inclusion of the earlier data from 1990-1997, in sharp contrast to the classic TWFE. Taken together, these results are consistent with the idea that the looking across the border can help mitigate the bias from violation of parallel trends in the longer panel with data from the 1990s.²² More recent work using the border discontinuity

²¹In STATA this can be implemented using `teffects,ra` command. For staggered DiD designs, something analogous is done within new DiD packages such as by [Callaway and Sant’Anna \(2021\)](#), [Borusyak et al. \(2024\)](#), or [Dube et al. \(2023\)](#).

²²[Jha et al. \(2024\)](#) find that using a TWFE-CZ design and cross-state commuting zones when the data stretches back to 1990 yields more negative estimates for restaurant employment than from a TWFE-BCP design. However, [Dube et al. \(2024\)](#) show that these TWFE-CZ results are driven by the same violation in parallel trends in the 1990s that we have discussed above; and use of clean event studies suggests small employment effects for either BCP or CZ design.

design has found broadly similar results as using a standard DiD (Coviello et al., 2022). This is not surprising given the lack of pronounced differences across states raising or not raising minimum wages after 1998, as shown above.²³

However, there are disagreements about whether the border discontinuity design is superior to other strategies, including comparing all counties within the U.S. (Neumark et al., 2014; Jha et al., 2024). On many observable dimensions, BCP’s appear to be more similar to each other than a randomly chosen pair of two counties from the United States (see, e.g., Dube et al. (2016)). There is also indication from other, non-minimum wage settings that BCPs can help mitigate endogeneity problems in policies (for example, Boone et al. (2021) show this in the context of unemployment benefit extensions during the Great Recession.)

At the same time, as we see, there are several possible limitations to this design. In some cases, neighboring areas are particularly different from areas farther away. One example is estimating the effects of city-wide minimum wages. Most cities are surrounded by areas that are less urban and possibly less comparable than other cities farther away. This suggests that the border design is better suited for studying state-level policies than city-level ones. The second concern relates to the external validity. Border counties experiencing minimum wage changes in some cases may be quite different from interior counties experiencing the same event, so the causal effects found there may not generalize (in the western United States, most border areas tend to be more sparsely populated).

The third—and possibly the most important—issue is that there may be cross-border spillovers from the treated area that can contaminate the control area, violating the stable unit treatment value assumption (SUTVA). One possible mechanism is the flows of workers across the border. There is some evidence that this might play a role, at least when looking very close to the border, from (Jardim et al., 2017). At the same time, this is a case study of a city minimum wage, which is probably not a good setting for this design to begin with. In contrast, Coviello et al. (2022) do not find that cross-border worker flows respond to minimum wage increases in the retail sector. Given the uncertainty, empirical spillover tests are therefore important to implement when using the border discontinuity design. Relatedly, estimates that are more robust to cross-border spillovers can be used, such as using “control rings,” motivated by the “donut holes” in regression discontinuity designs. That is, one can exclude areas that are very close to the border where cross-border spillover may be more likely. (Most spillover estimates die out farther away from border (e.g., Jardim et al., 2017)). This could retain the benefits of proximity while mitigating some of the possible bias arising from spillovers (for examples of this, see Boone et al. (2021) and Jha et al. (2024)).

The substantive stake from whether to use TWFE-BCP versus TWFE-CZ is minor, since we have already seen the issues with TWFE models and data from earlier years. Nonetheless, this does point to the practical issue that results can sometimes be sensitive to “bandwidth choice” or the definition of local labor markets.

²³The border discontinuity design has also been used to study minimum wage policies outside of the U.S., like the study of Indian minimum wages by Soundararajan (2019).

3.2.3 Methods to estimate the overall effect of the policy

So far in this section, we have focused on constructing the counterfactual for minimum wage treatment. However, there is also the question of how to define the treatment group. The key choices here revolve around: (1) whether to define the group narrowly or broadly, and (2) what measures to use in defining the group.

As we will see in our review of the evidence on employment effects in section 3.3.2, most studies in the existing literature consider relatively narrow groups, driven largely by the need to find a group where the policy is clearly binding. For example, a large part of the literature focuses on teens because it is easier to detect a wage effect for this group than for broader demographic groups, such as all workers, or all workers without a college degree. This is especially true during periods when the bite of the minimum wage has been low, making it binding for a relatively small share of the workforce. However, the external validity of estimates based on teens is questionable, particularly since teens constitute a steadily declining share of the low-wage workforce.

Another approach has been to study narrow industries like restaurants, which employ a sizable share of the minimum wage workforce. As we saw in Section 3.2.1, (Table 1, Panel A, column 1) there is a clear increase in restaurant earnings after minimum wage events. However, many minimum wage workers are not in the lowest-wage sectors. To assess the impact on those workers, a different approach is needed.

One option is to look at the entire population or workforce. This may be feasible when studying large, highly binding, changes. For example, the 1967 increase in coverage and level of the U.S. minimum wage was substantial enough to have a detectable effect on the average wage. This allowed researchers to estimate an OWE for all working age (16-64) or prime age (25-54) individuals (Bailey et al., 2021; Derenoncourt and Montialoux, 2021). Studying recent U.S. minimum wages in low-wage areas, Godoey and Reich (2021) is also able to obtain a clear earnings estimate for workers with high school or less education.

However, studying such broad groups of workers is not feasible for all minimum wage increases studied in the literature, especially when studying changes at a relatively low level of the minimum wage. For example, while the event study estimates for overall earnings are statistically significant for the full period (Table 1, Panel A, column 2), they are not for events before 1998 (wage estimate of 0.017 (s.e.=0.10) not reported in the table). The existing literature has addressed this issue in three different ways: frequency distribution and bunching; demographic prediction and machine learning; and incumbent workers.

Bunching in the frequency distribution. When the minimum wage is raised to, say, \$12 an hour, and there is full compliance, jobs that would pay below \$12 absent the policy would now “disappear” from that part of the frequency distribution. The number of these “missing jobs” (missing from the part of the distribution below \$12) as a share of the overall workforce gives us an indication of the “bite” of the policy change, or the fraction of workers directly affected by the minimum wage increase. Some of these missing jobs are in fact upgraded—now paying at or slightly above

\$12—creating a “bunching” or “excess jobs” in the frequency distribution (here excess is relative to the counterfactual frequency). In contrast, some of the missing jobs may be destroyed if they are no longer profitable, and do not show up in the frequency distribution at all. Yet other jobs paying at or slightly above \$12 may come into existence if, for example, an employer forced to pay a higher wage is better able to retain or recruit workers—thereby adding to the excess job count. This implies that by subtracting the missing jobs count from the excess jobs count provides us with an estimate for the overall effect of minimum wage policy on the number of low-wage jobs.

The key to this exercise is estimating the counterfactual distribution absent the minimum wage increase. An early approach by [Meyer and Wise \(1983\)](#) used parametric assumptions about the wage distribution to infer the employment effect; but these proved fragile [Dickens et al. \(1998\)](#). Instead, [Cengiz et al. \(2019\)](#) proposed estimating how the entire wage frequency distribution (that is, jobs per capita by wage bin) changes in response to a minimum wage increase using a DiD design. Figure 4 shows the bin-by-bin estimates using this approach applied to 138 increases in state-level minimum wages during the 1979-2016 period. There is clearly a reduction in jobs that pay less than the new minimum wage (bin labeled “\$0”); at the same time, there is an increase in jobs that pay at or within a few dollars of the new minimum wage. When these are added up (as in the running sum plotted as a dashed line), the total effect on low-wage jobs is found to be close to zero (e.g., at +\$2 or higher).

There are a number of advantages to this approach. First, by locally evaluating excess and missing jobs around the minimum wage, we can study the impact for a wide variety of workers. Subject to sample size issues, it is in principle possible to get a “first stage” wage effect for low-wage workers even from groups with high average wages (such as college graduates). Second, beyond a narrow range above the new minimum, it is highly unlikely to have a major causal impact on the upper tail of the frequency distribution. Therefore, the lack of such upper-tail effects serves as a useful check on the research design. For example, Figure 4 shows no such upper tail effects when studying state minimum wage increases. In contrast, the presence of such upper tail effects provides a warning against mistaken interpretation; an example of this comes from evaluating city-side minimum wages (e.g., Seattle), as discussed in [Dube and Lindner \(2021\)](#). Third, this approach allows joint estimation of the impact on low-wage jobs as well as on wages earned by those workers, including possible spillover effects. Importantly, the spillover effects estimated using this method do not assume away possible disemployment effect, in contrast to approaches using the wage density only. Finally, excluding changes in the upper tail can also improve the precision of the employment estimates. A growing number of papers have implemented a frequency-distribution-based approach to estimate the overall impact on jobs (e.g., [Derenoncourt and Montialoux, 2021](#); [Jardim et al., 2022](#); [Wursten and Reich, 2023](#); [Azar et al., 2023](#)).

Demographic prediction-based approach. While the frequency distribution approach estimates the overall impact on low-wage employment, it is less useful for studying other outcomes, including labor force participation, job finding, or non-labor market outcomes such as health. An alternative approach uses demographic predictors to construct a likely low-wage group. An early example was [Card and Krueger \(1995a\)](#), who used a probit model, CPS-based demographic variables, and a judiciously chosen set of higher order interaction terms to predict low-wage workers. Their “high probability” group successfully captured a somewhat broader set of workers than just using teens or

other age/education based groupings.

More recent work by [Cengiz et al. \(2022\)](#) builds on this in two ways. First, they use modern machine learning methods to optimally choose predictors, partitions, and interactions to classify workers. Second, they show how one can construct a broader “high recall” set of workers that in the U.S. can capture the vast majority of likely minimum wage workers while still obtaining a clear wage effect. They also show that optimally constructed groups are able to capture much larger shares of true minimum wage workers (a high “recall rate”) while maintaining the same share of workers classified as minimum wage workers who truly are so (the “precision rate”) than *ad hoc* groupings often used in the literature such as teens, individuals under age 30 without a high school degree, etc.

We can illustrate this approach using the event study design described previously. Appendix [F](#) describes the construction of a “high recall” sample covering 75 percent of minimum wage workers using CPS-based demographic predictors and a gradient boosting algorithm. For this “high recall” group, in Table [1](#), Panel A, column 5, we see that there is a clear effect of the minimum wage events on log wages (an increase of 0.02), allowing for a meaningful assessment of the employment effect. The “high probability” sample (Column 3) shows a larger wage effect (0.03), hence focusing on lower-skilled workers; this group allows better assessment of possible labor-labor substitution. Neither group shows a statistically significant or sizable employment effect; the OWE’s are -0.07 for the “high probability” groups of 0.34 for the “high recall” group. Finally, the low-probability group (column 6) offers a useful falsification test: reassuringly, there is no statistically significant wage or employment effect for this group. Alternatively, the employment rate of the “low-probability” group can be used as a control in the event study design through regression adjustment. This is implemented in the Appendix Table [A3](#), column 6; comparing columns 5 and 6, we find that in this case, the additional control has very little impact on the wage or employment estimates.

As a practical matter, when implementing this approach, it is useful to use a model to directly predict who is a low-wage worker—as opposed to predict an individual’s wage, and then classify workers based on their predicted wage. With the latter approach, the model predicts the mean wage conditional on covariates and misses out on who might be making low wages conditional on covariates.

Incumbent worker based approach. Finally, a number of papers have constructed treatment groups by considering the wages that workers were earning before the minimum wage. This approach focuses on “incumbent workers” who were earning low-wages. For example, if the minimum wage is rising to \$15 per hour, we can select workers who were initially earning at or below that level (or perhaps below a slightly higher threshold, such as \$17). We could then compare outcomes for similarly earning workers in jurisdictions that raised the minimum wage versus jurisdictions that did not. This approach requires panel data that allows us to track the outcomes of workers over time; examples of papers using this approach include [Clemens and Wither \(2019\)](#); [Jardim et al. \(2022\)](#); [Hampton and Totty \(2023\)](#). We consider a related design of comparing these workers to others earning slightly higher wages in Section [3.2.4](#). We also discuss some issues with this empirical design there.

3.2.4 Exploiting nation-wide variation in the level of minimum wages

Many countries set minimum wages at the national level only, meaning that there is no local variation. An extensive literature leverages major changes in such nation-wide minimums to study the effect of the policy, most notably in Germany, Portugal, the U.K., as well as the US. Here, we review the key methodologies used in such a setting.

Most studies of these changes in national minimum wage exploit the variation in exposure to minimum wage between groups. These groups could be defined on the basis of locations, individuals, firms, or industries. The key idea is to compare the evolution of employment and wages for the more exposed (“treated”) groups with the less exposed (“control”) groups. The definition of a group varies across applications, so we begin with a general discussion followed by group-specific aspects.

Studies of nation-wide minimum wages typically estimate the following class of difference-in-differences regression:

$$y_{gt} = \gamma_g + \beta BITE_g \times Post_t + \delta_t + \varepsilon_{gt} \quad (9)$$

where y_{gt} is the outcome variable (usually wages—to assess the first stage—and employment) in group g at time t , where δ_t and γ_g are group and time effects, respectively, and $Post_t$ is equal to one after the reform. The key variable, $BITE_g$, corresponds to the exposure based on the pre-reform period. For example, $BITE_g$ may represent the average wage or the fraction of minimum wage workers in group g before the policy change. Alternatively, it could represent the expected change in average wages within group g if there were full compliance with the policy, sometimes referred to as the “gap” that must be closed to comply with the policy. In the U.S. context, some studies have leveraged pre-policy variation in statutory minimums across jurisdictions to study the effect of federal minimum wage changes (e.g., [Clemens and Wither \(2019\)](#), [Derenoncourt and Montialoux \(2021\)](#)). In many applications, $BITE_g$ is calculated based on multiple pre-policy years to avoid issues related to mean reversion. Furthermore, some applications explore non-linear relationships between exposure ($BITE_g$) and the outcome (y_{gt}), or estimate the relationship non-parametrically.

The standard assumption in these types of DiD estimators is that of parallel trends: that changes in potential outcomes for highly exposed and less exposed groups would be the same in the absence of the policy. While this assumption is not directly testable, it is possible to check whether it holds during the years prior to the policy by studying pre-existing trends.

It is also important to clarify what this exposure design identifies and under what conditions. For example, one must assess whether the spillover effects from more-treated to less-treated units could create problems. In general, such spillover effects would imply a violation of the SUTVA, thereby biasing the estimates. As a practical matter, a serious violation of the SUTVA is unlikely in most applications given the low share of minimum wage workers in the economy. Even in the presence of significant spillovers, any effect on the outcomes of untreated groups will be limited if, for example, that group represents a much larger share of the population than the treated group. Nonetheless, in some situations where the analysis sample of both the treated and control groups is restricted to highly exposed industries or locations, a violation of the SUTVA could become a more important

issue.

Another (related) concern about identification arises if workers are able to move across groups—for example, from highly exposed locations to less exposed ones, or from highly exposed firms to less exposed firms. In such cases, again, the resulting violation of the SUTVA means that the relative change in group-level employment would not represent the true causal effects of the policy.²⁴

Given these considerations, it is also useful to analyze the time series evolution of the outcomes for the least exposed or untreated groups around minimum wage hikes. A significant change in their outcomes could reflect some spillover effects and the violation of the SUTVA. Of course, such changes might simply reflect the presence of aggregate time effects, which the DiD empirical strategy addresses under the parallel trends assumption. However, being transparent about the outcomes for the untreated groups helps us better understand the economic environment in which the minimum wage is instituted. It also allows us to assess whether the shocks affecting the untreated groups are likely to have a similar impact on the treated groups (i.e., the parallel trend assumption holds).

In addition to parallel trends, the identification of national minimum wage effects also relies on assumptions about the correlation between heterogeneous treatment effects and the exposure measure $BITE_g$. Consider a scenario in which the responsiveness of low-exposure groups to the minimum wage differs from that of high-exposure groups. For instance, when exploiting regional variation, it is possible that low-wage locations have higher labor market concentration, resulting in less negative employment effects in these regions—say exactly zero. In contrast, high-wage locations might be less concentrated, leading to strongly negative employment effects there. In this case, a regression estimate from the comparison of low- and high-wage regions would suggest a strong *positive* relationship between employment change and exposure, even though the *average* causal impact of the policy is negative.

A similar issue arises when exploiting variation across industries (see page 20 of [Krueger \(1994\)](#) for a discussion of this issue). Imagine an industry A (e.g., local services) that is highly exposed to the minimum wage hike but faces an inelastic output demand—so there is no change in employment. In contrast, industry B (e.g., manufacturing) is higher-paying with lower exposure but faces a very elastic output demand; therefore, the employment effects there are negative. In this hypothetical example, the average causal effect of the policy on employment is negative. However, when employment changes are compared across industries, there is a positive association between exposure and employment change.

These examples highlight the key issue with the exposure design in the presence of treatment effect heterogeneity. To see the problem more formally, suppose that the true data-generating process has the following form:

$$y_{gt} = \gamma_g + \beta_s BITE_g \times Post_t + \delta_t + \varepsilon_{gt} \quad (10)$$

where g reflects all the groups (e.g. 4-digit industries), while s reflects a higher group level aggregation

²⁴The presence of such worker mobility indicates that the untreated groups are also affected by the minimum wage, implying a violation of SUTVA. However, estimates can be adjusted to account for the effect on untreated firms. If the treated population represents a small share of the total population, this adjustment will have only a minor effect on the estimates.

(e.g., s could be 1-digit industries). The above equation can be rewritten as:

$$y_{gt} = \gamma_g + \beta BITE_g \times Post_t + \delta_t + (\beta_s - \beta) BITE_g \times Post_t + \varepsilon_{gt} \quad (11)$$

If we estimate the above equation by simply regressing $BITE_g$ on y_{gt} , then the error term is going to be $u_{gt} = (\beta - \beta_s) BITE_g \times Post_t + \varepsilon_{gt}$. We would obtain a consistent estimate of β only if u_{gt} is uncorrelated with $BITE_g$. Note that this can hold only if the bite ($BITE_g$) is independent of the treatment effect ($\beta - \beta_s$) otherwise, we have a classic endogeneity problem. This issue has also received attention in other contexts (e.g., impact of medicare on healthcare expenditure). Consistent with the previous two stylized examples, [Sun and Shapiro \(2022\)](#) formally show that when effects are heterogeneous across groups, the above estimator can perform very poorly and may produce an estimate that falls outside the range of individual treatment effects.

The above derivation also provides a hint on how to deal with heterogeneity. One solution proposed by [Sun and Shapiro \(2022\)](#) is to parameterize this heterogeneity. In our case, this would mean estimating equation (10) and allowing β to vary across the broader group s . A different approach would include time-varying s -specific fixed effects, ψ_{st} , in the regression. Given that $(\beta_s - \beta) BITE_g \times Post_t$ is st -specific, this would address the endogeneity issue. The inclusion of such fixed effects also allows us to control for s -specific trends over time.

This second approach has a close connection to the “grouping” estimator of [Blundell et al. \(1998\)](#). In many applications, groups are created by splitting the data based on two or more dimensions. For example, groups can be created based on location and age. The following regression estimates the effect of the policy

$$y_{rat} = \gamma_{ra} + \beta BITE_{ra} \times Post_t + \theta_{rt} + \psi_{at} + \varepsilon_{rat} \quad (12)$$

where y_{rat} is the outcome in location r of a worker with age a at time t , γ_{ra} are age-by-location group fixed effects, θ_{rt} are location-time effects, and ψ_{at} are age-time effects. This grouping regression design only exploits changes within location and within education; therefore, any heterogeneity along these dimensions will not lead to a bias in the estimates. This design also controls for location or age-specific trends.²⁵

A key limitation of these approaches is that they do not allow heterogeneity within s (e.g., at the g -level). As a result, researchers need some prior knowledge about the potential dimensions of heterogeneity in order to implement these methods effectively. In the absence of such prior knowledge, one can estimate the effect of the policy while allowing for arbitrary heterogeneity if there are some groups at the g -level that are not exposed to the policy. In this case, it is recommended to implement a fuzzy difference-in-differences design that conducts pairwise comparisons between unexposed units and exposed units [De Chaisemartin and d’Haultfoeuille \(2018\)](#). In case of the grouping estimator described above, this approach will only identify the effect of the policy for those ages and locations where unexposed units exist.

Regional bite. In practice, a popular approach in the literature is to exploit local variation in

²⁵Sometimes this grouping estimator is implemented in a first-differenced form (see [Manning \(2021\)](#)): $\Delta y_{rat} = \beta BITE_{ra} + \theta_r + \psi_a + \varepsilon_{ra}$.

exposure to the minimum wage. [Card \(1992b\)](#)’s seminal paper studied the impact of the 1990 increase in the federal minimum wage on teenage workers by comparing low-, medium-, and high-wage states. The basic idea behind this approach is that the same nation-wide minimum wage will be more binding in low-wage locations like Louisiana than in high-wage locations like California.

Follow-up work has applied this approach to various contexts such as the introduction of the U.K. national minimum wage (e.g. [Stewart \(2002\)](#); [Dolton et al. \(2012\)](#)), the German national minimum wage ([Caliendo et al., 2018](#)), the 2007-2009 U.S. minimum wage increase (see [Clemens and Wither \(2019\)](#)), and the 1966 expansion of the U.S. minimum wage coverage ([Bailey et al., 2021](#)). All of these studies estimate a regression similar to equation 9, where the group, g , is defined by the region.

In many applications, the variation across locations is not large enough to produce first-stage wage effects when the outcomes y_{gt} are calculated for the total workforce. One solution to this problem is to restrict the sample to specific groups, such as lower-skilled workers or teens. In this case, it is important that $BITE_g$ is calculated for the same sub-sample of workers.

Another way to address the lack of first-stage effects is to take advantage of additional exposure variation by creating subgroups within locations based on worker characteristics, such as age or education (see, e.g., [Kertesi and Köllő, 2003](#); [Dickens et al., 1999](#); [Dube, 2019a](#)). However, most applications do not often include education-time or age-time fixed effects in the regression. Notable exceptions are [Manning \(2021\)](#) and [Dube \(2019a\)](#). Given the issues regarding heterogeneous treatment effects, when exploiting sub-regional variation in the bite, it is advisable to use (or at least report) estimates with the inclusion of time-varying age, education, and location effects.

A recent paper by [Giupponi et al. \(2024\)](#) further develops the idea of using low-wage locations as a control group for high-wage locations. Their method combines the regional approach with the frequency distribution approach discussed in Section 3.2.3. They estimate the effect of the minimum wage throughout the entire skill distribution (instead of using the change in high-skilled workers employment as a control for low-skilled ones). They compare the change in the number of workers employed in a given wage bin at low-wage (treated) locations to the change in the number of jobs with similarly skill requirements at high-wage (control) locations. Once the effect—corresponding to the difference between high- and low-wage locations—is estimated for all job types (i.e., the full frequency distribution), they can zoom in and study the employment and wage changes of lower-skilled workers. Therefore, instead of using higher-skilled workers as a control for lower-skilled workers, they compare change in low-skill employment across high- and low-wage locations; this still allows them to use employment changes of higher-skilled workers for falsification.

A subtle but important aspect that is often overlooked is understanding the characteristics of “compliers” whose wages are shifted by treatment $BITE_g$. These are the workers who would be below the minimum wage at certain low-wage locations but above the minimum wage at high-wage locations. Any identification strategy using regional variation necessarily estimates the average treatment effect only for this “complier” population, and not for those who are always treated—since the latter group earns below the minimum wage at every location. Estimating an overall average treatment effect on the treated—including the always-treated—would require some extrapolation from the complier

groups.

It is also important to keep in mind that the estimated differences between high- and low- wage locations could reflect changes in worker mobility following the minimum wage. Most studies that directly test migration responses do not find any change in internal migration in response to the minimum wage (see Section 4.1.5).

Firm-level exposure. Another approach in the literature exploits firm- or establishment-level variation in the exposure to the minimum wages (e.g. [Card and Krueger \(1994b\)](#); [Machin et al. \(2003\)](#); [Draca et al. \(2011\)](#); [Harasztsi and Lindner \(2019\)](#)). To apply this approach, panel data on firms are needed so that pre-policy exposure of firms can be calculated.

A key advantage of this design is that it allows us to examine the margins of adjustment of businesses that are the most strongly impacted by the minimum wage. Since these firms play a key role in absorbing the minimum wage shocks, this design can provide valuable insights into understanding the mechanisms through which the policy influences low-wage labor markets. Here, we make two points about interpreting the evidence from this cross-firm design.

First, it is important to emphasize that the estimated effects reflect the relative change in employment in highly exposed firms, which is not the same as the employment effects on workers. If workers in highly exposed firms are laid off but find jobs at less exposed firms, then the firm-level exposure design will overestimate the negative consequences of the policy on workers. Recent findings by [Dustmann et al. \(2022\)](#) and [Rao and Risch \(2024\)](#) provide evidence for such a reallocation mechanism in the German and U.S. contexts.²⁶ Therefore, we recommend that firm-level estimates of employment be supplemented with worker- or local-level estimates that guard against such biases.

Second, the firm-level analysis inherently focuses on incumbent firms that existed before the policy change, as minimum wage exposure can be only calculated for these firms. This limits the “representativeness” of the analysis, although newly entering firms typically represent a small share of the workforce, at least for a few years after the policy change. Still, when implementing this design, it is advisable to complement the analysis with evidence on the impact on firm-level entry rates to assess the importance of this concern. In Section 4.1.4, we review the evidence on firm entry, and show that although the effect on entry rate is inconclusive, most studies suggest only a limited change in firm entry following minimum wage hikes.

Worker-level exposure. Some studies exploit exposure to the minimum wage by pre-policy wages. The key idea is to compare the employment trajectories of workers earning below the new minimum wage to the those of workers earning slightly above it (see, e.g., [Currie and Fallick, 1996](#); [Abowd et al., 2000](#)). Implementing this approach requires panel data of workers. This method is similar to the “incumbent workers” design (described above), which compares initially low-wage workers across treated and untreated areas. But different from that design, the worker-level exposure approach can be used in setting where there is only a national level minimum wage change.

²⁶The presence of reallocation leads to a violation of the SUTVA. However, in a context where a small share of the total workforce is in highly exposed firms, the change in outcomes for untreated firms will be limited.

This approach focuses on the wage and employment implications of individuals who are most strongly affected by the minimum wage. The estimates provide worker-level responses to the policy (instead of firm- or local-level ones) and speak more directly to the welfare consequences. However, it is important to remember that simple comparisons of high- and low-wage workers’ trajectories can be misleading. Workers at the bottom of the wage distribution typically experience greater wage growth and less stable employment than those higher in the wage distribution, even in the absence of any policy change (Ashenfelter and Card, 1981).²⁷ To address this issue, pre-policy years are often used to estimate and control for “placebo effects” in the employment trajectories of low- and high-income individuals (see Clemens and Wither (2019); Dickens et al. (1999); Dustmann et al. (2022)).

Finally, the standard design inherently focuses on incumbent workers with pre-policy wages, often missing the responses among those who are unemployed or out of labor force. To study the unemployed, researchers need to “predict” what their wage would be if they were employed. This prediction can be made in various ways (see our discussion Section 3.2.3). In practice, predictions are typically based on demographics (see Currie and Fallick (1996)) or on wages from pre-unemployment spells (see Clemens and Wither (2019)).

3.2.5 Exploiting minimum wage exemptions

In many contexts, there are minimum wage exemptions for small firms, long-term unemployed, or younger workers. A strand of the literature leverages this variation, which inherently creates “treated” (non-exempt) and control (exempt) workers. The most common approach exploits age cut-offs in the level of the minimum wage (see e.g., Kabátek (2021)), but cut-offs in the duration of unemployment that define who is considered long-term unemployed are also used (Umkehrer and Berge (2020)).

These papers effectively demonstrate the employment and wage differences around the cut-off by applying a regression discontinuity design (see, e.g. Kreiner et al. (2020)). However, it is unclear what is identified by comparing the outcomes of very similar workers who are subject to different wage-floors. The fact that a firm might choose the cheaper worker when faced with two similar workers earning different wages tells us more about the substitutability of these two workers and labor market frictions that might dampen this substitution. It is difficult to extrapolate these findings and derive implications for minimum wage policies that apply to all individuals. This discussion points more broadly to the importance of clarifying what theoretical objects a particular research design allows us to identify.

A second related issue is that, in practice, firms in many countries do not seem to utilize these sub-minimum wages, leading to no first-stage wage differences around the cut-off. Although it remains an open question when and why firms do not utilize sub-minimum wages, recent evidence suggests that it could be related to fairness norms (see Giupponi et al. (2024)). The presence of such norms suggests that variation by age group might reveal the influence of such norms, rather than the impact

²⁷The higher wage growth at the bottom could reflect mean-reversion: many workers at the bottom might have had a temporary negative shock and then revert back the following year. Additionally, it could reflect that wages at the bottom can only go up (not down). In terms of employment, the lower stability reflects that minimum wage jobs tend to have high turnover.

of increasing wage floors across the board *per se*.

3.3 Review of the evidence on employment effects

Having laid out various methodological considerations when estimating the effects of minimum wage on employment, we now turn to reviewing the body of empirical evidence. Typically, minimum wage reviews focus on the elasticity of employment with respect to the minimum wage; even worse, sometimes they mix that elasticity with the own-wage elasticity (OWE). [Brown et al. \(1982\)](#) found that the employment elasticity with respect to the minimum wage ranges between -0.1 and -0.3 for teens. The focus on teens allows for a more apples-to-apples comparison of the elasticities, but at the cost of reduced external validity. [Neumark and Wascher \(2008\)](#) and [Neumark and Shirley \(2022\)](#), mixing various elasticities, concluded that the average elasticity is negative. On the other hand, [Doucouliagos and Stanley \(2009\)](#) and [Belman and Wolfson \(2014\)](#) focus on the elasticity of employment with respect to the minimum wage, suggesting that the overall impact of minimum wages on employment is minimal. While these reviews are somewhat informative, we discussed the problems with comparing estimated elasticities of employment with respect to the minimum wage across studies in Section [3.1](#).

Recent articles, such as [Harasztosi and Lindner \(2019\)](#) and [Dube \(2019a\)](#), review the evidence on OWE by considering studies that report both first-stage wage effects as well as employment effects. A more comprehensive list of papers has been compiled by [Dube and Zipperer \(2024\)](#), who have created an online repository of minimum wage studies with both employment and wage estimates.²⁸ These are studies published starting in 1992 that use either “quasi-experimental” or “experimental” variation, and study statutory minimum wages (as opposed to floors set by collective bargaining, private platforms, etc.). For studies that meet these criteria, [Dube and Zipperer \(2024\)](#) consider the database to be close to comprehensive for the U.S., U.K., Germany, and Canada. The database contains a single OWE estimate from each study, based on an assessment of the authors’ preferred empirical specification; when this was unclear, [Dube and Zipperer \(2024\)](#) reached out to the authors for guidance on model or sample selection. When there were estimates for a range of groups, the database includes the estimate for the broadest group of low-wage workers presented in the paper. In some instances, the database includes an average of multiple estimates; this is driven by either the presence of multiple low-wage groups or specifications, or because of the authors’ preferences, as communicated directly. As of the time of writing, the database (version 1.0) contains 88 studies, 72 of which were published in academic journals. In this chapter, we will focus on these 72 studies only.

When the authors report the OWE estimate and standard errors (as was the case for 28 of the studies), these are the estimates included in the database. When the authors do not report the OWE, the estimates in the database are obtained by dividing the employment effect by the wage effect, as described above. Following [Harasztosi and Lindner \(2019\)](#), confidence intervals are constructed using the delta method and assuming the independence of employment and wage estimates.

When considering the magnitudes of the OWE, it is useful to think about the welfare consequences.

²⁸The repository is available at: <https://economic.github.io/owe/>

In this respect, a key ingredient in the welfare analysis of minimum wages is the elasticity of total pre-tax earnings of all low-wage workers (including those without a job) with respect to the policy. When the OWE is -0.1, it means that around 10 percent of the potential increase in earnings is forfeited due to job losses. Conversely, when the OWE is -0.9, it indicates a loss of 90 percent of that potential earnings growth.

Taking these considerations into account, we follow [Dube and Zipperer \(2024\)](#) and categorize the OWE's as follows. We consider an OWE that is less negative than -0.4 to represent either a “small negative” or “positive” effect on jobs—as they imply that the total wage bill collected by workers increased by at least 60% from the policy. We characterize an OWE between -0.4 and -0.8 as having a “medium negative” impact on employment. Finally, we take an OWE more negative than -0.8 to signify a “large negative” impact on jobs—as the disemployment erases more than 80 percent of the potential earnings gains. Like all such categorizations, ours inherently reflects some degree of subjectivity. Others may choose to make their own assessment using the underlying OWE data.

3.3.1 Estimates across all studies

We begin by describing the central tendencies in the OWE estimates, drawing from the discussion in [Dube and Zipperer \(2024\)](#). Spanning more than thirty years of research and a wide range of affected workers, the estimates from the OWE repository suggest that the overall impact of minimum wages on employment is small. The median OWE estimate from the analysis of 72 published studies is -0.13, suggesting a minor impact of the minimum wage on jobs. This estimate implies that the total earnings of low-wage workers rise by 87% of what one would expect if there were no job losses due to the policy. Put differently, employment reductions offset only about 13% of the potential earnings gains. While the mean OWE estimate of -0.25 is somewhat larger in magnitude, it is still consistent with fairly modest job losses and relatively large earnings gains. This conclusion remains when we restrict attention to the 57 US studies: for that group, the median and mean OWE are -0.11 and -0.22, respectively.

In addition to the central tendencies, we can also look at the size distribution of estimates using our rubric: more negative than -0.8 as “large negative”, those between -0.4 and -0.8 as “medium negative”, and those more positive than -0.4 as “small negative or positive”. Figure 5 shows the distribution of OWE for the 70 published studies. 51 (or around 71%) of the published studies have positive or small negative estimates; in contrast, 21 studies (or around 29%) have estimates that would imply a large or medium negative effect. The overall evidence base suggests that to date, minimum wage policies have raised wages of low-wage groups much more than they have reduced their employment—thereby effectively increasing total earnings for low-wage workers.

3.3.2 Estimates for broad groups

Most (51 out of 72) studies considered so far are from narrow subgroups. These include teens, restaurant workers, nurses, and grocery workers, to name a few common groups. However, estimates

from these groups may not be very informative about the overall impact of the policy on all affected workers (the average effect of treatment on the treated, in the language of program evaluation).²⁹

Given the concerns about external validity of such narrow groups, it is useful to separately consider papers that provide estimates for broad groups of low-wage workers that better represent the overall impact of the policy. [Dube and Zipperer \(2024\)](#) also classify the papers in the OWE database as “broad” when they are likely to capture the majority of workers affected by the policy. This includes, for example, [Cengiz et al. \(2019\)](#), which provides an estimate of the impact of US minimum wages on the total number of low-wage jobs as described above. [Derenoncourt and Montialoux \(2021\)](#), which captures all adults aged 25 to 55 years, and [Neumark et al. \(2004\)](#), which captures all incumbent low-wage workers. Finally, OWE estimates which are focused on all workers with at most a high school degree, as in [Monras \(2019\)](#), are considered broad as well.

Based on this classification, just under one third of published studies (or 21 of 72) can be classified as “broad.” (Interestingly, the share of “broad” studies has risen considerably over the years, from zero in the 1990s to 48 percent in the 2020s, as shown in Appendix Figure A8.) Figure 6 plots the distribution of the OWE estimates from these studies and also compares it to the distribution of the other 47 studies of narrower groups. Broad-group studies tend to have estimates closer to zero, with a median OWE estimate of 0.02. 19 of the studies (or 90%) have OWE estimates that are either positive or small negative. The contrast with narrower groups is also illuminating. The narrow-group studies are also consistent with small employment effects, but the median OWE estimate for those 51 studies is somewhat more negative at -0.17. One possible explanation for a slightly more negative effect found among narrower groups is that these studies are more likely to reflect labor-labor substitution or possibly reallocation across industries, although many other explanations exist.

3.3.3 Heterogeneity of employment effects

We can further disaggregate the 51 estimates from narrow group studies. Of these, 21 were about restaurants or retail sectors, while 15 were for teens. There has also been a notable decline in teen share of studies, falling from around 50 percent in the 1990s to 4 percent in the 2020s, which is consistent with declining importance of teens among minimum wage workers. In contrast, the share of restaurant / retail studies has remained fairly stable, going from 33 percent in the 1990s to 26 percent in the 2020s (see Appendix Figure A8).

When we focus on restaurants or retail sectors, we obtain a median OWE of -0.09 and a mean of -0.17. So, the low-wage sector studies yield fairly similar estimates as the overall evidence base, though a bit less negative. In contrast, the 15 teen studies had a median OWE of -0.17 and a mean of -0.25. In other words, the estimates for teens appear to be a bit more negative, potentially indicating more labor-labor substitution, even as the mean OWE of -0.25 still suggests any teen disemployment effect has been small.

Of course, there are other confounders when comparing these studies, including the methods and

²⁹For example, Congressional Budget Office has assumed different OWE estimates for teens and adults in its analysis of prospective minimum wage increases [Congressional Budget Office \(2019\)](#).

quality of the study. Ideally, when assessing heterogeneity, we would account for such differences. As a simple gauge, it turns out that the OWE estimates have tended to become less negative over time, possibly due to falling publication bias or due to improved data and methods. For example, studies published before 2010 had a median estimate of -0.40, while studies published since then had a median of -0.04. As it turns out, slightly fewer teen studies have been published since 2010 than before (8/15); in contrast, the vast majority (41/57) of the non-teen studies were published after 2010. When we look within these two time periods, there is no systematic indication that teen studies are more negative; on average, the gap in the median OWE between teen and non-teen studies was positive prior to 2010, and negative afterwards, with an overall (within-period) gap in the median OWE close to zero (based on a median regression of the OWE on teen and period indicators).

Age. Since so many aspects may differ across studies of teens and non-teens, a better assessment of heterogeneity can come from studies that have explicitly provided estimates for teens or young workers along with an estimate for overall low-wage workers (or a broader demographic group). Here we consider papers that provide wage and employment estimates for both groups. Using event-based approaches, both [Cengiz et al. \(2019\)](#) and [Cengiz et al. \(2022\)](#) find similarly small/positive and statistically insignificant OWE's for teens (0.36 and 0.40, respectively), as well as overall low-wage workforce (0.41 and 0.11, respectively). These mirror our event-based findings in Table 1: teen and high recall OWEs are 0.12 and 0.34, respectively.³⁰ [Hampton and Totty \(2023\)](#) provides employment elasticities using an event-based approach for incumbent low-wage workers by age bins: 16-21, 25-34, 35-44, 45-54, 55-61, and 62-70. The average OWE is 0.07, and for all but the last group, OWEs fall between -0.02 and 0.40, and employment estimates are indistinguishable from zero. For the oldest (retiring-age) group, they find a positive and statistically significant employment effect with OWE of 0.82 (which is consistent with delayed retirement, a result also found in [Borgschulte and Cho \(2020\)](#).) [Manning \(2021\)](#) provides estimates for teens as well as individuals aged 20-24 year using various TWFE-logMW specifications. Without controls for state-specific trends or regional controls, the OWEs for teens and older workers are both around -1.2; and both are highly sensitive to inclusion of controls for heterogeneity, for reasons we have already discussed. For the UK, [Giupponi et al. \(2024\)](#) report an overall OWE of -0.20 for 25-64 year olds; adding younger workers and considering 16-64 year olds produces an estimate of -0.10, suggesting broadly neutral/positive effects for younger workers. Our general take-away is that there is little difference in teen and non-teen OWEs *within studies*; so the between-study differences in estimates uncovered above likely reflects methodological and/or sample differences.

Most studies examine the heterogeneous impact of the policy across broad age groups. One notable exception is [Giuliano \(2013\)](#), who finds that in response to a federal minimum wage increase, a large retailer substitutes *towards* teens overall, but substitutes *away from* teens from low-income households toward teens from higher-income households. This is consistent with the presence of substitution when we look at narrow groups, which we will return to in section 4.1.3, but it could also reflect an increased labor supply of teens from higher-income families.

Race. Besides age, another notable within-study heterogeneity to consider is race. Using the

³⁰Our estimates in this chapter are slightly different from [Cengiz et al. \(2022\)](#) due to our combining multiple phases of events, and using slightly different event definitions.

1960s expansion in the US minimum wage, two recent studies arrived at slightly different conclusions. [Derenoncourt and Montialoux \(2021\)](#) find similar null effects on employment overall and for Black workers in particular, with OWEs of 0.06 and 0.15, respectively. [Bailey et al. \(2021\)](#) also finds a small OWE of -0.14 for workers overall, but a somewhat larger magnitude -0.29 for Black men. At the same time, all these estimates would be considered “small” using our heuristic. In more recent studies that have specifically report wage and employment effects by race, [Cengiz et al. \(2019\)](#) finds an overall OWE of 0.41, as compared to -0.09 for Black/Hispanic workers; however, the latter is fairly imprecise. In [Cengiz et al. \(2022\)](#) for the high recall group overall the OWE is 0.11, while for Black/Hispanic workers it is found to be -0.52; again, however, the latter estimate is imprecise, making the contrast difficult to ascertain. Finally, [Wursten and Reich \(2023\)](#) implement a stacked event study analysis with prominent minimum wage changes (more than 5 percent) between 1979-2019, considering low-wage workers with a high school or less education (as well as workers in the restaurant sector). They find OWEs between -0.01 and 0.04 for white, Black and Hispanic workers.

Overall, there is no clear evidence on systematic differences by race. One obvious challenge is sample size: the samples for Black or Hispanic workers are much smaller, which makes detecting (true) heterogeneity by racial subgroups difficult.

Composition of jobs. The minimum wage can also affect the composition of jobs in the economy. Here we review the evidence on industry composition, while we discuss the impact on occupation/job tasks in Section 4.1.3, since those evidence are often interpreted in the context of capital-labor substitution and automation. We find a strong indication of heterogeneous impact of the policy by sector. [Cengiz et al. \(2019\)](#) and [Gopalan et al. \(2021\)](#) find negative effect in the tradable/manufacturing sector, while there is no discernible employment effect in the local service/non-tradable sector. It is worth keeping in mind that in the U.S. (and other high-income countries), most low-wage workers are in local service jobs, and so the overall policy impact will be dominated by the impact in that sector. In line with that, the across studies median OWE from the retail and restaurant sectors (-0.09) is very similar to the overall median OWE of 0.13. Similarly, [Harasztosi and Lindner \(2019\)](#) find a significant reduction in employment in the tradable sector, but not in the local service sector in Hungary. A larger share of low-wage workers in Hungary are employed in the tradable sector, so those effects have more important overall implications there. Still, the significant reduction in employment in one sector does necessarily lead to a reduction of overall employment in the presence of reallocation of workers to other sectors. [Harasztosi and Lindner \(2019\)](#) find a more muted disemployment effect in their worker level (as opposed to firm-level) analysis—consistent with the presence of such reallocation.

Short vs. medium run estimates. One concern often raised about the relatively small employment estimates is that they reflect short-run effects that do not capture the dynamics (e.g., firm exit) that take longer to take effect. Given the nature of minimum wage variation in the U.S., very long-run estimates are difficult to obtain. At least until recently, most state-level variations did not last for more than a decade. Some researchers have used filtering of the minimum wage measure to concentrate on “low frequency” variation by decomposing minimum wage variation to a permanent and a transitory components (e.g., [Baker et al. \(1999\)](#)). In practice, the [Baker et al. \(1999\)](#) approach is similar to adding lagged minimum wages as explanatory variables (distributed lags models), but it is not a particularly transparent way of understanding long run impacts.

However, what is true is that quite a few studies estimate what we would call “medium run” estimates (e.g., 4-7 years out effects). Historically, these were estimated based on distributed lags TWFE-type models (e.g., (Allegretto et al., 2011; Dube et al., 2010; Neumark et al., 2014; Jha et al., 2024, e.g.)). Some of these TWFE-type estimates show larger longer-run estimates, but are highly sensitive to specification, as demonstrated already.

Numerous recent work tends to report event-based estimates up to at least 4 years (sometimes up to 7 years) after a policy change (Azar et al., 2023; Bailey et al., 2021; Cengiz et al., 2019; Harasztosi and Lindner, 2019; Clemens and Strain, 2021; Derenoncourt and Montialoux, 2021; Godoey and Reich, 2021; Godøy et al., 2024; Monras, 2019; Ruffini, 2022; Wursten and Reich, 2023, e.g.). Monras (2019) finds that the employment effect grows steadily more negative between event years 0 and 3.³¹ Clemens and Strain (2021) also find an indication of a more negative medium run effects, but only for the 6 “large increase” states (more on this below). However, the rest of the papers do not suggest that medium run employment effects are substantially more negative than shorter run effects. This comes out of the 1967 expansions (studied by Derenoncourt et al. (2021) and (Bailey et al., 2021)); state-level changes after 1979 (studied by Cengiz et al. (2019), Cengiz et al. (2022), Godoey and Reich (2021), Rao and Risch (2024), Vergara (2023)) or studying the impact of doubling the minimum wage from Hungary Harasztosi and Lindner (2019).

An added piece of evidence comes from the 60 events we consider in this chapter. For at least 35 of these events, we can calculate a 4 year out effect, and for 17 we can calculate a 6 year out effect, as shown in Appendix Figure A4 for restaurants, and Appendix Figure A5 for the high-recall group. In both cases the 6-year out impact on wages is larger than at other time horizons. However, the 6-year out employment estimates are small, positive in sign, and do not indicate a more negative effect on employment than at shorter horizons.

Finally, using region-by-demographic variation, the 20-year effect of the 1998 introduction of the UK National Minimum Wage was also found to be close to zero (see the working paper version of Manning (2021), and Dube (2019a) which reports an OWE of -0.04 (s.e.=0.21)).

High versus low minimum wages. Naturally, there is a lot of interest in understanding possible non-linearities in the employment effect at higher levels of minimum wages. This is especially so given the increasing experimentation by governments on more ambitious wage standards.

Cengiz et al. (2019) use event-by-event estimate to assess heterogeneity by the bite of minimum wage (as measured by the Kaitz index) at the state level, and finds no heterogeneity for events up to 2016. Godoey and Reich (2021) estimate impact for broad group of low-wage workers at the sub-state level and does not find larger job losses in more binding areas. Clemens and Strain (2021) consider heterogeneity by size of increases in the post-2013 period using an event study design. Considering all 27 events, their preferred specification find a small overall OWE of -0.26 for their “low-skill” group (individuals ages 25 and under without a high school degree); for the six states (and D.C.) raising their minimums by \$2.50/hour or more, they find an OWE of -1.01, while for the states raising less

³¹This finding is critically dependent on the use of a linear pre-treatment trend adjustment, which is different from the other event-study approaches discussed here.

the OWE estimate is 0.46.³² (For a different group—those under 21—the two OWEs are much closer at -0.41 and -0.03.) However, it is unclear how sensitive their findings are to the particular cutoff, or use of specific set of time-varying covariates (such as the state GDP).

Additionally, we consider estimates from our 60 events by the size of the minimum wage change. In Appendix Table A3, column 4, we report the estimates from the top tercile (third) of the events in terms of the size of minimum wage increase. We find a larger wage effect for the restaurant sample, but the OWE of 0.23 (restaurant) and 0.41 (high recall) are not meaningfully different from their counterparts based on all events in column 1 (-0.03 and 0.34, respectively). But as we mentioned, it is quite useful to plot the full set of event-by-event estimates to transparently show the patterns. This is what we do in Figure 7 for restaurant sample, and Appendix Figure A7 for the high-recall sample; where the estimates are all sorted by the size of the minimum wage increase associated with each event. There is an overall positive association between the wage estimates and the log minimum wage change, with correlation coefficients of 0.43 for restaurants, and 0.27 for high recall. However, there is no evidence of a negative association between employment estimates and the minimum wage change; the correlation coefficients are 0.17 and 0.12 for restaurants and high recall, respectively.

Outside of the U.S. context, [Harasztosi and Lindner \(2019\)](#) study the impact of doubling the minimum wage in Hungary—raising from the current U.S. median to minimum wage ratio to the level of France—in two years. They estimate an OWE of -0.22 in the medium term, as we discussed above, mainly coming from the manufacturing sector. The sharp increase in the UK National Living Wage wage between 2016 and 2019 also offers another example, and the OWE from the UK in this period were all small (e.g., [Giupponi et al., 2024](#)).

Going forward, we think further research is needed to carefully assess possible “non-linearities” from minimum wage increases that are both large and are at high levels. The challenge is to consider the different dimensions of what a “bigger” minimum wage means (size of change, pace of change, ending level), while at the same time recognizing how results on thresholds may be fragile, especially when considering different dimensions of the policy. For this reason, we think it is valuable for researchers to transparently show event-specific estimates along various relevant dimensions, such as the minimum wage level, size of the increase, bite of the policy, etc. This can mitigate some of the challenges arising from false discovery when searching for heterogeneity by thresholds.

Market concentration. Models of imperfect competition in the labor market imply more positive employment effects when the labor market is more monopsonistic. [Azar et al. \(2023\)](#) assess this by considering heterogeneous effects of minimum wage hikes during the 2010-2016 period, focusing on the retail sector. Using a similar bunching approach as [Cengiz et al. \(2019\)](#), they show that in non-concentrated markets ($\text{HHI} < 0.25$), there is a negative (but statistically insignificant) overall OWE of around -0.50. In contrast, for concentrated markets ($\text{HHI} \geq 0.25$) there is a sizable positive and statistically significant OWE of 1.8. Interestingly, this heterogeneity appears to be the clearest in the retail labor market; for restaurants, the differences are much more muted, though that may in part reflect less accurate concentration measurement. In a similar vein, [Wiltshire \(2022\)](#) finds a more

³²Hourly wages increased 18% in the six states, raising their minimums by \$2.50/hour following the minimum wage, while 11% in other treated states. Therefore, these OWE estimates imply a very high degree of non-linearity in employment responses.

positive employment effect from the 1996-1997 federal minimum wage increase in local labor markets that had become more concentrated from prior Walmart Supercenter entry. Finally, and broadly along the same line, [Okudaira et al. \(2019\)](#) find less negative employment effects from minimum wage increases in Japan on plants with greater estimated monopsonistic markdown. There is similar evidence from developing country contexts as well: employment effects (OWEs) are found to be more positive in more concentrated local labor markets with greater concentration in India ([Soundararajan, 2019](#)).

Interactions with earnings subsidies. Minimum wages are often presented as a substitute for earnings subsidies such as the Earned Income Tax Credit (EITC). However, there may also be important interactions between these policies. From a theoretical perspective, while the EITC is thought to encourage labor supply—a desired outcome—this in turn also push down pre-tax wages, an unintended consequence ([Rothstein, 2010](#)). In principle, the minimum wage can mitigate such a downward push in pay (e.g., [Lee and Saez \(2012b\)](#)). The empirical evidence on the EITC and minimum wage interaction is quite limited. [Neumark and Wascher \(2011\)](#) use a TWFE-logMW model and interactions with state minimum wage and EITC policy over the 1997-2006 period. They find that the earnings effects from EITC for single mothers are more positive in the presence of a higher minimum wage, though the earnings and employment effects are more negative for some childless adults. Given the problems with the TWFE estimates (compounded when dealing with two separate policies), we think it would be worth re-visiting this topic using a clean event-based design.

Summary. So what can we learn from the heterogeneity of employment effects found in the literature so far? First, the literature suggests limited heterogeneity by age group, especially when we look within studies. The evidence on differences by race is somewhat under-powered due to small sample problems, but the differences found tend to be modest. There is more evidence on the changing composition of jobs—including by sector or occupation—but these are unlikely to be contributing to an overall decline in employment. At the same time, more work needs to be done to understand heterogeneity by policy environment and market structure, especially in light of existing work suggesting relatively more positive employment effects in concentrated markets. Finally, most event study estimates do not show significantly larger medium-term effects for typical minimum wage increases observed to date. However, the variation across the size or level of minimum wages remains poorly understood. Further rigorous research on non-linearities in employment effects, and the identification of potential turning points (beyond which disemployment is more pronounced) would be highly valuable.

3.4 Effect on total hours

Much of the literature has focused on headcount employment. However, some papers directly study the adjustment in hours conditional on employment. [Zavodny \(2000\)](#) considers teen employment for the 1979-1993 period and finds no meaningful effect on hours. [Couch and Wittenburg \(2001\)](#) find large negative hours elasticities, but these are based on specifications without any year fixed effects (i.e., pooled time-series regressions). [Neumark et al. \(2004\)](#) finds some reductions in hours using a short panel of workers using household CPS data from 1979-1997. Using a TWFE model, [Sabia \(2009\)](#) finds

no effect on hours conditional on employment for the retail sector during the 1979-2004 period. [Dube et al. \(2007\)](#) study the 2004 San Francisco minimum wage increase using a restaurant survey and found no impact on hours; both the FTE and the headcount estimates were close to zero. [Allegretto et al. \(2011\)](#) find that estimates for teen hours are sensitive to regional and trend controls during the 1990-2009 period, mirroring employment results.

More recently, [Cengiz et al. \(2019\)](#) use an event-based approach and changes in the frequency distribution, considering the effect on headcount, as well as full-time equivalent (FTE) employment for all affected workers. The OWE for FTE employment is 0.60 while for headcount it is 0.41, both being statistically indistinguishable from zero and indicating little overall impact on hours. [Cengiz et al. \(2022\)](#) use a similar event-based design for demographic probability groups and find a small (but statistically significant) increase in full-time share for the narrower (lower wage) high-probability group, but a null effect on the broader high-recall group, with no change in overtime shares. Using private payroll data from a sample of businesses, [Gopalan et al. \(2021\)](#) do not find any significant reduction in weekly hours conditional on employment either among incumbent workers, or affected establishments. Studying the increase in the 2014 Seattle minimum wage using administrative data on hours and wages, [Jardim et al. \(2022\)](#) find a large reduction in total hours worked in jobs paying below \$19, although, as discussed previously, these results likely reflect more hours being paid above \$19. Following incumbent workers, they find a statistically significant reduction in hours. However, the implied magnitudes are small: the average overall FTE-based OWE across specifications is around -0.2. Using SIPP data to track incumbent workers and events over the 1984-2014 period, [Hampton and Totty \(2023\)](#) find a small own-wage elasticity of hours of around -0.2 for all affected incumbent workers. Based on event-based estimates from 2005-2017 and ACS data, [Godoe and Reich \(2021\)](#) find no reduction in weekly hours, either overall or in local areas where the minimum wage was highly binding.

Overall, most U.S. evidence—especially more recent, higher quality estimates—do not suggest large adjustments to hours of work conditional on employment. This is particularly true for low-wage workers generally; for particular subsets of workers, there are some estimates in both negative and positive directions.

Outside the U.S., a handful of papers have examined the impact on hours adjustments. The evidence is mixed at best. In the UK, [Stewart and Swaffield \(2008\)](#) finds a reduction in hours, while [Connolly and Gregory \(2002\)](#) find no effect. In Ireland, [McGuinness and Redmond \(2018\)](#) find a reduction in hours, especially for workers on temporary contracts. In Germany, [Bossler and Gerner \(2020\)](#) and [Buraudel et al. \(2020\)](#) find a reduction in contractual hours. However, [Buraudel et al. \(2020\)](#) shows that this does not translate into a significant reduction in actual hours worked. This suggests that some of the responses in hours may reflect reporting rather than a substantial change in behavior.

4 Margins of adjustment

The previous section reviewed the extensive empirical literature on how minimum wages affect employment and wages. In this section, we broaden the scope and consider the impact on low-wage labor markets more generally, with a focus on unpacking the various margins of adjustments. The outcomes considered in this section are those that can shed light on how firms actually respond to the cost shock resulting from higher minimum wages.

4.1 Review the evidence on various margins of adjustment

There is a long list of potential channels through which minimum wages affect the low-wage labor market. Here we list the most important ones:

1. Increased wages
2. Non-compliance of the policy (i.e., firms pay below the minimum wage)
3. Reduction in non-wage compensation or amenities (e.g., fringe benefits or working conditions)
4. Change in employment (firm exit/entry; substitution to other inputs)
5. Change in population through in/out-migration of workers, or change in labor force participation
6. Wage retrenchment of higher-skilled workers
7. Pass-through to consumers by raising output prices
8. Pass-through to suppliers by lowering input prices (e.g., rent, prices of intermediate goods and services)
9. Pass-through to firm-owners by lowering profits
10. Change in worker turnover and cost reduction coming with lower training costs
11. Improved productivity (increased effort, reorganization, greater allocative efficiency)

Among these mechanisms, the first five are closely related to the previous section. The evidence on non-compliance (Channel 2), change in non-wage amenities (Channel 3), capital-labor substitution and firm entry and exit (Channel 4), migration and participation responses (Channel 5) refine the measurement of the employment responses and the change in worker's compensation. For example, the presence of a substantial cut in non-wage amenities could imply that firm-level labor costs are less affected than suggested by the first-stage wage estimates. Furthermore, estimates on firm dynamics can shed light on sources of employment changes and inform modeling choices.

In contrast, channels 6 through 11 shed light on how the change in labor cost is absorbed—if not through curtailing low-wage employment. A rise in the costs associated with low-wage workers could be

covered by: a reduction in the expenses for high-wage workers (Channel 6.); an increase in total revenue through higher output prices—which in turn is aided by insensitive consumer demand (Channel 7.); lower rents or prices of other inputs (Channel 8.); lower profits (Channel 9.); decreased turnover and the implied savings on training and recruitment costs (Channel 10.); and improved productivity (Channel 11.). Therefore, these channels provide us with critical evidence on the incidence of the policy.

Next, we review up-to-date evidence on all of these channels. For each adjustment margin, we list the relevant studies along with our summaries in Tables 2-5. For some of these margins (like price, turnover, employment) we have relatively well-developed literature with many available studies. In these cases, we restrict our attention to published papers, as they have been peer-reviewed and are finalized. However, for some margins, we only have a few (or no) published papers to consider. Furthermore, sometimes this limited evidence comes from different national contexts (e.g., U.S., U.K., Germany, Hungary, China), making it difficult to compare the results. In such cases, we expand our discussion to include relevant working papers in order to paint as complete a picture as possible. We also apply our subjective assessments on the quality and credibility of papers to draw some (at times tentative) conclusions whenever possible. Our hope is that future research will fill in the gaps in these areas, allowing us to reach more definitive conclusions.

4.1.1 Non-compliance

As we discussed in Section 2.2, to date, the core goal of the minimum wage literature has been to understand the link between increased labor costs and the change in employment. Indeed, the lack of a sizable overall employment effect—as documented in Section 3.3—is notable only if minimum wage policies significantly raise labor costs. If firms simply do not comply with the policy, or if they are able to offset policy-induced wage changes by sharply reducing non-wage benefits, there would be no reason to necessarily expect a substantial change in employment.

What are some ways in which non-compliance could occur? These would, in general, depend on the institutional setting, and here we list some possibilities. First, and most simply, firms may pay their regular employees below the mandated floor. Second, higher minimum wages could create incentives for employers to shift some of their employment to the informal sector characterized by greater non-compliance. This is especially relevant in the context of low- and middle-income countries with a higher degree of informality; therefore, we discuss this form of non-compliance in Section 6. Third, firms could also shift their hiring toward undocumented immigrants, for whom minimum wage policies may be less likely to be enforced. Finally, another important margin is shifting work to the gig economy (independent contractors) to get around minimum wage and other employment regulations. Importantly, a growing number of jurisdictions (including California, Massachusetts, and New York) have used litigation and/or legislation to enact wage floors for some gig workers, such as ride-share or delivery drivers (Jacobs et al., 2024). Given the growing importance of such jobs in the economy, this channel could be more important in the future (Boeri et al., 2020).

The extent of non-compliance is primarily an empirical question. However, it is worth noting

that compliance might not be driven solely by legal enforcement. In fact, a recent paper by [Stansbury \(2024\)](#) documents that employers' legal incentives to comply with the minimum wage laws in the U.S. and U.K. are rather limited. At the same time, compliance could still end up being wide-spread if the minimum wage also shapes fairness perceptions, raises reservation wages, or shifts wage-setting norms.

Studying non-compliance is inherently difficult since it is not directly observed in most data. A common approach pioneered by [Ashenfelter and Smith \(1979\)](#) proxy non-compliance with the share of workers who report earning wages below the statutory floor in surveys such as the CPS. This approach usually indicates a nontrivial level of non-compliance. [Ashenfelter and Smith \(1979\)](#) find 20-40% of workers earn below the statutory minimum wage; more recently, [Bernhardt et al. \(2013\)](#) find that around 20% of low-wage workers experience some form of wage-and-hour violation in the U.S. Notably, the incidence of minimum wage violations is lower than the violations related to other types of workplace regulations such as overtime, meal breaks, or off-the-clock working hours. [Caliendo et al. \(2019\)](#) find a similar extent of non-compliance (around 7%-25%) in Germany. Studies have also found that the share of workers below the minimum wage seems to increase with the level of minimum wages ([Goraus-Tańska and Lewandowski, 2019](#); [Clemens and Strain, 2022](#)). [Clemens and Strain \(2022\)](#) estimate that underpayment reflects around 16-20% of the realized wage gains.

The key problem with these estimates is that separating actual underpayment from measurement error in reported wages is inherently difficult. For example, even if there is no true under-payment, but a constant share (γ) of workers misreport their wage as being below or above the true level with equal probability, at least $\frac{\gamma}{2} \times P(MW)$ of workers would report earning below the minimum wage, where $P(MW)$ is the share of workers truly earning exactly the minimum. Moreover, one would find that this measured share rises with the level of the minimum wage, which naturally raises $P(MW)$.³³ This illustrates the challenges in reliably quantifying how non-compliance mediates the impact of minimum wage policy on earnings. Furthermore, some of the actual payments below the minimum wages could reflect exemptions (e.g., tipped workers in most states in the U.S.) and not real non-compliance with the policy. Most studies deal with this issue by focusing on non-exempt workers, but that itself is measured noisily in surveys. Therefore, this approach is likely to overestimate the extent of non-compliance.

Regardless of the precise *level* of non-compliance, it is unlikely that it very substantially erodes the *change* in earnings resulting from a minimum wage increase. One way to think about this is through an analogy with speed limits on highways. There is substantial non-compliance with speed limits, as many drivers speed. At the same time, when the speed limit drops from 65 mph to 55 mph, even most non-compliers adjust their speed downward to avoid being far out of compliance. As a result, the average driving speed can fall substantially even as there is a non-trivial amount of noncompliance at both 55 and 65 mph regimes. Applying this logic to the low-wage labor market would suggest that minimum wages can substantially raise wages even without strict compliance.

This logic is consistent with the empirical fact that we observe clear wage increases among

³³Using this same setup, [Autor et al. \(2016\)](#) show that measurement error can also inflate the measured spillovers in wages (in a manner analogous to overstating non-compliance).

low-wage workers following minimum wage hikes in both household surveys and administrative data. This wage increase is also present for undocumented immigrants (see [Orrenius and Zavodny \(2008\)](#)) and in the informal sector (see [Section 6](#)). These findings indicate that firms raise workers' wages following minimum wage increases, regardless of the exact extent of compliance with the law. Finally, it is worth noting that while greater compliance might lead to a larger wage increase (something that workers would care about), it would be of secondary importance for estimating the own-wage employment elasticity—as long as there is a clear increase in the average wage earned by the relevant low-wage group (i.e., a strong first stage).

4.1.2 Amenities

Firms can respond to the minimum wage by reducing non-wage amenities, such as fringe benefits. In the extreme case, if companies could fully offset the increase in wages by spending less on amenities, there would be no net change in labor costs ([Clemens, 2021](#)). Alternatively, in the presence of some market imperfections (for example, monopsony), the provision of amenity could even increase when the minimum wage increases, especially when amenities and wages complement each other in workers' preferences [Dube et al. \(2022\)](#).

Unfortunately, direct evidence on amenities is rather limited, especially since we rarely observe all relevant workplace characteristics. [Simon and Kaestner \(2004\)](#) use various measures from the National Longitudinal Survey of Youth (NLSY) and the Current Population Survey (CPS) to assess this channel. Exploiting state-level variation in the level of minimum wages and applying a TWFE design, they find no discernible change in fringe benefits (employer provision of health insurance, pension coverage, dental insurance, vacation pay, and training/educational benefits) and working conditions (shift work, irregular shifts, and workplace safety) of low-wage workers. An alternative approach considers the impact on firm-level expenditures on fringe benefits. [Card and Krueger \(1995b\)](#) and [Brown \(1999\)](#) review the earlier literature, finding a very small or no reduction at all in these expenses. More recently, [Harasztosi and Lindner \(2019\)](#) find no indication of a decrease in fringe benefits in response to a large increase in the minimum wage in Hungary.

In addition to this broader evidence, the literature paid special attention to two specific elements: employer-provided health insurance and on-the-job training.

Employer-provided health insurance. A particularly important fringe benefit in the U.S. context is employer-provided health insurance. [Marks \(2011\)](#), [Meiselbach and Abraham \(2023\)](#) and [Clemens et al. \(2018\)](#) study the relationship between health insurance and the minimum wages by exploiting various survey data (Current Population Survey, Medical Expenditure Panel Survey, American Community Survey). They document a small reduction in providing health insurance for low-skilled employers, especially at small firms. [Clemens et al. \(2018\)](#) also provide a back-of-the-envelope estimate implying that roughly 9-16% of the wage increase of low wage workers is offset via reductions in insurance. This suggests that workers benefit substantially from the policy even after any fringe benefit offset is taken into account.

On-the-job Training. How minimum wages impact on-the-job training is unclear *a priori*. In the standard competitive framework, general training will be financed by employers in the form of lower wages. In response to the minimum wage, employers will offer less on-the-job training, which could partly offset the positive impact of the policy. However, in the presence of imperfect competition in the labor market, the impact on training is uncertain. [Acemoglu and Pischke \(2003\)](#) show that when firms earn some rent on workers, it might be more profitable for employers to respond to a minimum wage by increasing worker productivity through additional training rather than laying them off.

The early empirical literature on training often confirmed the prediction of the neoclassical model on the negative consequences of minimum wages on training ([Shiller, 1994](#); [Neumark and Wascher, 2001b](#)), although not always ([Grossberg and Sicilian, 1999](#)). [Acemoglu and Pischke \(2003\)](#) provide a critical review of this literature and conduct their own empirical analysis. Contrary to the prediction of the neoclassical model, they find no discernible effect in on-the-job training in the U.S. context. Outside of the U.S. context, [Hara \(2017\)](#) finds a significant reduction in firm-provided training, while [Arulampalam et al. \(2004\)](#) and [Bellmann et al. \(2017\)](#) find limited effects on-the-job training in the U.K. and Germany, respectively.

Indirect evidence on amenities. While direct evidence on the impact on amenities is limited, there is a sizable body of indirect evidence suggesting that the offset in non-wage amenities is likely to play a limited role in the adjustment to the minimum wage. First, studies consistently find that worker turnover and quit rates decrease in response to the minimum wage (see Section 4.1.10 for details)—a sign that workers like their jobs more when the minimum wage goes up. Consistent with this interpretation, [Holzer et al. \(1991\)](#) show that minimum wage jobs attract more job applicants than jobs that pay either slightly more or slightly less than the minimum wage. Second, overall job satisfaction increases among low-wage workers following the minimum wage ([Bossler and Broszeit, 2017](#); [Güral and Ayaita, 2020](#)), suggesting that the wage increases are not substantially offset by reductions in amenities that workers value. Finally, a large amenity offset would be inconsistent with the significant adjustments along other margins such as output prices that we discuss below. After all, there would be no need (or scope) for raising prices if firms were to compensate most of the wage increase through cutting non-wage compensation.

4.1.3 Substitution with other inputs

In Section 3.3, we concluded that the evidence suggests that the employment and hours responses are likely to be limited in many contexts. A closely related literature directly inspects the channels through which employment may be affected. In the standard neoclassical framework, one such channel is the substitution between low-wage workers and other inputs in production, including higher-skilled workers and capital (see, e.g., equation (1)). Here we review that evidence to better understand the mechanisms, while also considering their implications for low-wage employment overall (as opposed to employment at particular firms, occupations, etc.).

Capital-labor substitution. Based on firm-level evidence from Hungary, firms exposed to minimum wages seem to increase their investment (and hence the capital stock) relative to the

unexposed firms ([Harasztosi and Lindner, 2019](#)). [Geng et al. \(2022\)](#) and [Hau et al. \(2020\)](#) also find an increase in investment and the capital stock in the Chinese manufacturing sector. [Fan et al. \(2021\)](#) find a firm-level increase in robot adoption in China in certain periods but not in others. In both countries, some reduction in employment at highly-exposed firms is found as well (relative to less-exposed ones), which is in line with the prediction of the standard theory: the higher labor cost forces firms to substitute labor with capital, e.g., through increased automation (see equation 1). However, it is important to note that these firm-level effects may not reflect the market-level impact on capital stock, since less exposed firms could reduce their investment in response to the minimum wage. Unfortunately, there is little direct evidence on market-level responses in the stock of capital.

Several papers study this capital adjustment indirectly through the lens of a changing occupational structure. Here, the key question has been whether minimum wages lead to a reduction in automatable/routine jobs. Both [Lordan and Neumark \(2018\)](#) and [Aaronson and Phelan \(2019\)](#) find that the minimum wage reduces the employment of the lowest-wage workers in routine automatable occupations. However, the two studies disagree on the overall employment implications of these results. [Lordan and Neumark \(2018\)](#) argue that the reduction in automatable jobs contributes to a lower overall employment level of low-wage workers, while [Aaronson and Phelan \(2019\)](#) find that workers laid off from routine jobs are able to move into other expanding low-wage occupations. The broader evidence on employment reviewed in Section 3.3 appears to be more consistent with the interpretation of [Aaronson and Phelan \(2019\)](#).

Minimum wages may also reduce automation in certain scenarios. [Downey \(2021\)](#) studies the impact of minimum wages on higher-paying, automatable, jobs. He argues that higher minimum wages decrease the incentive to adopt de-skilling automation technologies (that replace more expensive middle-skill jobs with cheaper lower-skilled ones). His empirical analysis finds that an increase in the minimum wage significantly decreases mid-skill IT employment while increasing routine jobs, consistent with lower automation.

Labor-labor substitution. Besides substituting labor with capital, firms can also replace lower-skilled workers with higher-skilled ones. Such substitution would show up as heterogeneous employment impact of the policy by skill groups, e.g. younger minimum wage workers could be replaced by older, more experienced ones. We reviewed the evidence on the heterogeneous impact of the policy in Section 3.3.3 where we find little systematic indication of such heterogeneity along key demographic dimensions.

Some papers also study the impact of minimum wages on firms' skill demand as evidenced by the demographic composition within high-impact jobs or firms. [Clemens et al. \(2021\)](#) find that workers in low-wage occupations are more likely to be older and have a high school degree following minimum wage changes. They also document an increase in skill requirement in low-wage occupations using Burning Glass data on job postings. [Horton \(Forthcoming\)](#) implements a field experiment on an online platform and shows that imposing a minimum wage on a random subset of firms led those firms to hire more productive workers in an online platform. These pieces of evidence suggest that firms may shift their demand towards more productive minimum wage workers when forced to raise their pay. However, whether such firm-level changes in skill demand lead to greater unemployment in general

for the least productive workers is a distinct question. For example, displaced workers from narrow categories of jobs may be absorbed elsewhere.

A separate strand of the literature exploits discontinuities in the level of minimum wages by age to estimate substitution. As we discussed in Section 3.2.5 the evidence from such exemptions should be carefully interpreted since they do not directly speak to the impact of broad-based minimum wage shocks—the focus of this chapter. Even if there is significant labor-labor substitution when similar workers are allowed to be paid differently, such substitution may not be possible in response to an across-the-board minimum wage.

The evidence based on such age exemptions is mixed. In many contexts, firms apparently do not utilize sub-minimum wages to begin with, and pay the same wages to younger workers even when it is possible to pay them less. This is true in the U.S. (see [Katz and Krueger \(1992\)](#); [Card et al. \(1994\)](#) who find limited utilization, while [Neumark and Wascher \(1994\)](#) find more utilization in certain states); in the U.K. ([Giupponi and Machin, 2022b](#)); and in Finland ([Böckerman and Uusitalo, 2009](#)). In other countries, we do see greater utilization of age-specific floors, leading to reduced wages for younger workers and sometimes to higher employment (see [Kreiner et al. \(2020\)](#) in Denmark; [Kabátek \(2021\)](#) in the Netherlands; [Portugal and Cardoso \(2006\)](#) and [Pereira \(2003\)](#) in Portugal). These estimates suggest that in some cases, applying the minimum wage policy selectively to different groups could create sizable employment differences and inequality between treated and untreated workers.

4.1.4 Firm-entry and exit

In some models, exit and entry of firms play an important role in mediating the response to wage floors, as the increased cost of labor drives out low-productivity firms. In terms of firm exit, [Draca et al. \(2011\)](#), [Harasztosi and Lindner \(2019\)](#), [Mayneris et al. \(2018\)](#), and [Dustmann et al. \(2022\)](#) find that exposed firms exit at a faster rate following a minimum wage hike. In the U.S., [Rohlin \(2011\)](#) finds no change in the exit rate, while [Aaronson et al. \(2018\)](#) and [Chava et al. \(2023\)](#) find a significant increase. Furthermore, [Luca and Luca \(2019\)](#) document an elevated exit rate for restaurants, especially among establishments with very low consumer ratings to begin with. This latter finding is consistent with a “cleansing” effect of the minimum wage, where the most inefficient firms exit the market.

What is the contribution of firm exit to overall employment? This critically depends on whether workers at (likely low-productivity) exiting firms are able to find jobs elsewhere or are more likely to remain without employment. As [Dustmann et al. \(2022\)](#) show in the context of a model with monopsonistic competition (see their Online Appendix D), even if some small, inefficient firms exit the market, workers may move to more productive, surviving companies.³⁴ In this case, the employment effect of the policy would remain limited, while overall productivity improves. Therefore, the results on firm exit cannot be interpreted as direct evidence of loss in net employment. Empirically, findings from Germany documented by [Dustmann et al. \(2022\)](#) as well as U.S. evidence by [Rao and Risch \(2024\)](#) are consistent with such productive reallocation. In a somewhat different context in South

³⁴Imperfect competition in the product market can also lead to a similar type of re-allocation ([Rao and Risch, 2024](#), see, e.g.,).

Africa, [Bassier \(2022\)](#) also finds a similar reallocation of employment from lower productivity to higher productivity firms in response to a higher sectoral wage floor set by collective bargaining.

The effect on firm entry has also received attention in the literature, as it is unclear *a priori* what to expect in this margin. The standard neoclassical channel suggests that the increased labor cost would dampen firm entry. However, in the presence of some friction in the adaptation of technology, this prediction can be reversed. [Aaronson et al. \(2018\)](#) study the impact of the minimum wage in a putty-clay model. In this framework, newly entering firms can freely choose their input mix, allowing them to substitute costly low-wage workers with capital. Once this choice is made, however, firms operate using a fixed input mix: so incumbent firms' capital-labor substitution possibilities are restricted.

This framework predicts that short-term employment responses to the minimum wage will be muted, reflecting only scale effects, while the larger effects from substitution will come through firm dynamics. Incumbent firms with *ex post* suboptimal technology will exit slowly, creating entry opportunities for new firms that are able to choose a more efficient mix of inputs. Over time, obsolete labor-intensive technologies will disappear and will be replaced with capital-intensive technologies—leading to reduced employment among low-wage workers. Of course, this is only one theoretical possibility; different types of friction would suggest different compositions of firm entrants. For example, in a model with monopsonistic competition, higher-productivity firms may be more likely to enter following a minimum wage rise. It is worth noting that the muted short-term response followed by a more amplified long-term one is not a pattern that we see in much of the data (see Section 3.3.3 on evidence on short- and medium-term employment effects).

Moving on to direct empirical findings on firm entry, [Aaronson et al. \(2018\)](#) find an increase in entry rate in response to minimum wages, [Harasztosi and Lindner \(2019\)](#) find no change, while [Rohlin \(2011\)](#), [Draca et al. \(2011\)](#), and [Luca and Luca \(2019\)](#) find reduced entry. Therefore, the evidence is mixed but broadly suggests a fall in the entry rate.

Overall, the weight of evidence on firm dynamics suggests that the number of firms probably falls following a minimum wage increase (as the exit rate increases, while the entry rate likely does not rise—at least not strongly). In line with this, [Orazem and Mattila \(2002\)](#) and [Dustmann et al. \(2022\)](#) find a decrease in the number of firms resulting from a higher wage floor.

4.1.5 Migration and participation

Minimum wages can also shape the local working population through various mechanisms, including changes in labor force participation and migration responses.

Migration. The relationship between minimum wage levels and migration patterns is theoretically ambiguous. If jobs become more scarce as a result of the policy, one potential adjustment mechanism is for low-skilled workers to relocate to higher wage areas, where the wage floor is less binding. On the other hand, absent a sizable reduction in job finding rates, a higher minimum wage can also act as a magnet for low-skilled workers. For instance, in the monopsonistic competition framework,

where many firms are supply-constrained, employers are willing to hire more workers at higher wages—something that could induce migration to locations with higher minimums.

Several studies have analyzed the effect of minimum wages on the composition of the local population. An early study by [Castillo-Freeman and Freeman \(1992\)](#) examined the impact of a substantial increase in the minimum wage in Puerto Rico and identified the out-migration of low-skilled workers to the mainland United States as a significant adjustment mechanism. However, we should note that [Krueger \(1994\)](#) points to the fragility of the empirical designs applied in [Castillo-Freeman and Freeman \(1992\)](#). More recently, [Monras \(2019\)](#) studied the internal migration pattern in the U.S. context by examining the impact of 441 state- and federal-level minimum wage changes instituted between 1985 and 2012. He finds that minimum wages reduced both the low-skilled employment and the low-skilled population share. The main analysis is complicated by the lack of parallel trends, which might be related to selective timing of federal minimum wage increases (which are parts of the events used in his analysis). [Minton and Wheaton \(2023\)](#) assess the change in the low-skilled population (instead of the change in the share of low-skilled workers as in [Monras \(2019\)](#) and [Castillo-Freeman and Freeman \(1992\)](#)) using the prediction probability approach. They rely on the American Community Survey, which provides more precise information on population across locations than the Current Population Survey used by [Monras \(2019\)](#). [Minton and Wheaton \(2023\)](#) find an increase in net migration of low-skilled workers towards higher minimum wage states—a finding that is in stark contrast with the earlier U.S. literature.

Some papers also studied the impact on migration in the European context. Since there is no variation in statutory minimums within countries, nation-wide minimum wages could affect internal migration by inducing individuals to move away from (or toward) locations with a lower average wage and more binding policy. However, [Giupponi et al. \(2024\)](#) and [Ahlfeldt et al. \(2018\)](#) find no indication of a significant response of internal migration to minimum wages in the U.K. and in Germany, respectively.

A closely related literature studies the impact of the minimum wage on immigrants' location decision. The evidence on this aspect of the policy is somewhat mixed. [Orrenius and Zavodny \(2008\)](#) and [Cadena \(2014\)](#) find that low-skilled immigrants have been discouraged from settling in states with higher minimum wages. In contrast, [Giulietti \(2014\)](#) finds that the federal minimum wage changes in 1996-1997 and 2007-2009 induced a sizeable flow of low-skilled immigrants to more affected states. [Boffy-Ramirez \(2013\)](#) suggest that the impact of the policy depends on the amount of time immigrants have spent in the country: they find that a higher minimum wage attracts migrants who have already been in the U.S. for 2 to 4 years. However, for immigrants who have been in the U.S. for less than 2 years, or longer than 4 years, the minimum wage effect is indistinguishable from zero.

Overall, the evidence suggests that changes in low-skilled population at locations that raise minimum wages are unlikely to be large. The evidence on immigrant population is more mixed, but interestingly, the more mobile immigrants seem to favor places with higher minimum wages. This latter finding is consistent with the prediction of the monopsonistic competition framework.

Participation. Even if the size of the local population does not change, there could be a change in

participation rate in response to the minimum wage. The evidence on participation decisions is mixed. [Wessels \(2005\)](#) and [Luna-Alpizar \(2019\)](#) report a decrease in participation among young individuals, and [Lavecchia \(2020\)](#) finds a negative (but imprecise) effect among low-skilled individuals. Conversely, [Laws \(2018\)](#) and [Agan and Makowsky \(2021\)](#) find that an increase in the minimum wage leads to greater labor force participation. [Luna-Alpizar \(2019\)](#) finds an increase in participation for prime-age workers with high school education, [Borgschulte and Cho \(2020\)](#) and [Hampton and Totty \(2023\)](#) find positive effects on individuals near retirement age, while [Godøy et al. \(2024\)](#) show similar patterns for parents of young children. Finally, [Cengiz et al. \(2022\)](#), using the prediction-based approach, find no evidence of changes in participation overall, and find no indication of significant heterogeneity across age, race, and education.

In addition to the change in the participation rate, the effective labor force could also increase through increased effort by the unemployed to find a job. The evidence on the impact of minimum wages on search effort is limited and mixed. A few studies document small responses. [Laws \(2018\)](#) finds no clear evidence that minimum wages affect search effort, with some suggestive evidence that the unemployed may decrease it in response to the policy. [Adams et al. \(2022\)](#) finds a very short-lived positive effect only during the month of the minimum wage change, which vanishes soon afterwards. In contrast, recent evidence by [Piqueras \(2023\)](#) indicates that the policy has a significant long-term impact on the search effort of low-wage unemployed.

4.1.6 Wage retrenchment of higher-skilled workers

Another channel through which businesses can absorb higher labor costs for low-wage workers is by reducing pay for higher-wage employees. Is this theoretically possible? To the extent that wages reflect rent sharing, higher pay for some could lead to lower pay for others by reducing the surplus being bargained over. This would imply a *negative* wage spillover or ripple effect higher up in the pay distribution, especially within firms. We will review the evidence on spillovers and wage inequality in Section 5. However, the evidence is clear that there are *positive* (but limited) wage spillovers higher up in the distribution. This is also true within firms ([Hirsch et al., 2015](#); [Gopalan et al., 2021](#); [Dube et al., 2019](#)).³⁵ [Yao et al. \(2023\)](#) also find a positive relationship between executive compensation and the minimum wage hikes in China. In conclusion, the available empirical evidence indicates that the increase in the minimum wage for low-wage workers is not financed by a reduction in the wages of their better paid counterparts.

4.1.7 Output prices and consumers reactions

Output prices. There is an extensive literature studying the effect of the minimum wage on output prices. Table 4 lists the published papers at the time of writing. The evidence on prices comes from various sectors of the economy: restaurants, fast-food, manufacturing, and retail. The evidence almost

³⁵[Clemens et al. \(2018\)](#) documents a decline in employer-sponsored health insurance in occupations moderately higher up the wage distribution. They calculate that these amenity offsets for above-minimum wage workers represent a larger share of the observed wage increase than analogous offsets from minimum wage workers. Therefore, wage spillover estimates can overstate the spillovers in compensation, and hence the increase in workers' utility.

uniformly suggests that output prices increase following a minimum wage change. The magnitude of the price increase is also noteworthy. In the last column of Table 4, we report the pass-through rates: the size of the change in the unit price of the output relative to the change in unit labor costs. It is worth noting that the latter is typically imputed; as a consequence, the evidence on the pass-through rate should be taken with a grain of salt. That caveat notwithstanding, most papers find full pass-through. In other words, the unit labor cost increase is fully covered by the unit price increase; and in some cases, we see even more than full pass-through.³⁶

The evidence therefore shows that there is a clear and sizable increase in output prices following minimum wage increases. However, the corresponding changes in the quantity and quality of the products are less well understood. In the retail sector, [Leung \(2021\)](#) and [Alonso \(2022\)](#) find an *increase* in quantities following the changes in the minimum wage, while [Renkin et al. \(2022\)](#) find no effect on quantities (their point estimate is negative and somewhat noisy). In the case of hotel service, [Agarwal et al. \(2024\)](#) find no price pass-through, but at the same time a significant reduction in hotel occupancy following minimum wage changes. There are other puzzling aspects of their results, as they find a decline in output prices for unbranded and upscale hotels accompanied by a large reduction in occupancy rate—suggesting that other demand shocks may have coincided with the minimum wage changes they study.

[Cooper et al. \(2020\)](#) find nominal spending increases more than price gains, suggesting that, on net, consumers purchase more food, both at and away from home. For durable goods, they find that cumulative nominal spending increases roughly in line with prices when the minimum wage rises, suggesting no change in quantity. [Harasztosi and Lindner \(2019\)](#) evaluate the revenue response to the minimum wage and find a significant increase in the local service sector, but a reduction in the manufacturing sector. This indicates that the quantity response to the output price increase—and hence the scale of production—could vary substantially across sectors. As discussed in Section 3.3.3, researchers seem to find limited dis-employment effect in the local service sector, suggesting that the decline in the scale of production is likely to be limited there. At the same time, there is more evidence of negative employment effects in the tradable/manufacturing sector, in line with the larger reduction in the scale of production there.

The evidence on changes in production quality is also limited, but points in a positive direction. [Ruffini \(2022\)](#) finds significant improvement in patient health and safety (e.g., measured by patient deaths) in nursing homes after workers’ earnings increase due to minimum wage changes. [Brown and Herbst \(2023\)](#) find improved service quality in the child care context. However, there is much less evidence on service quality in major low-wage sectors such as hospitality and retail.

³⁶Here we focus on price pass-through. Note, however, that even if output prices were to increase one-to-one with the change unit cost cost (full pass-through), this does not necessarily imply that all of the costs associated with the minimum wage are borne by consumers. The latter depends on consumer reactions, and revenue responses. For example, [Harasztosi and Lindner \(2019\)](#) find a close to full price pass-through, but only 75% of the total wage bill increase is paid for by consumers. The discrepancy between the two can be explained by quantity responses to the price change.

4.1.8 Input prices and rent

In addition to being passed downstream to customers, the increase in labor costs could also be passed upstream to suppliers. This could be a quantitatively important channel, as 30-75% of firm-level expenses (depending on industry) are related to the purchase of intermediate goods and services.

In many sectors, such as U.S. retail and agriculture, the primary cost of business is rent. One way firms might finance an increase in labor costs is by reducing the rent they pay. In fact, if rental markets are sufficiently competitive and the land and properties used for lower-skill-intensive production cannot easily be repurposed by businesses relying on higher-skilled workers, the main adjustment to the minimum wage could come from lowering rental or land prices. Unfortunately, we are not aware of any direct evidence on rental price adjustments, and future research should focus on assessing the empirical relevance of this channel. However, in some cases, businesses own the land or stores they use. If minimum wage increases are primarily passed through to landowners, profits would decline for these land-owning firms. We review the evidence on profits in the next section.

In addition to adjusting rents, firms can also lower the prices they pay their suppliers. The extent of pass-through to upstream firms depends on those firms' own exposure to the minimum wage. If minimum wage-intensive firms tend to purchase goods from other minimum wage-intensive suppliers, the extent of pass-through will be limited. Therefore, the sorting of suppliers and buyers based on the type of labor they employ is an important consideration. [Demir et al. \(2024\)](#) find strong positive assortativity in wages within supplier-buyer relationships, suggesting the possibility of limited pass-through.

Direct evidence on changes in intermediate goods expenses following minimum wage increases is rather limited. [Harasztosi and Lindner \(2019\)](#) find that total spending on materials increased in the short term, while in the medium term, the effect on materials is smaller and insignificant. This evidence suggests that lowering supplier prices or rent is not a major margin of adjustment.

4.1.9 Profits

Given the stated redistributive goal of minimum wage policies, the impact on profits is a margin of considerable interest. Economists have developed various approaches to estimate the effect of minimum wages on firm profitability. One approach infers profitability from financial market reactions. The evidence from this approach presents a mixed picture. [Card and Krueger \(1995b\)](#) find no changes in the stock market valuation of firms employing low-wage workers around the time of U.S. minimum wage announcements, while [Bell and Machin \(2018\)](#) find significant reactions in the U.K. A key limitation of the financial market event study approach is that initial market reactions reflect investors' beliefs about the policy's impact, which may differ from the actual responses that manifest over time. Moreover, these event study estimates of abnormal returns are meaningful only to the extent that these announcements contain information not already anticipated (and thus priced in) by the market. In this regard, the surprise announcement in the U.K. studied by [Bell and Machin \(2018\)](#) offers a more informative case.

A second strand of this literature evaluates actual firm-level changes in profitability, typically measured as the profit-to-revenue ratio. Using data from the U.K., [Draca et al. \(2011\)](#) find a reduction in profits in response to minimum wage policies, while [Mayneris et al. \(2018\)](#) find no such change in China. A limitation of these estimates is that a decline in the profit-to-revenue ratio could result from either a reduction in profits (the numerator) or an increase in revenue (the denominator), possibly due to higher output prices. To avoid this complication, some studies have considered changes in profits themselves as the outcome. [Harasztosi and Lindner \(2019\)](#) exploit a large (60%) minimum wage increase and find a modest fall in profits, amounting to around 20% of the labor cost increase at highly exposed firms (relative to less-exposed ones). [Drucker et al. \(2021\)](#) apply a similar methodology and find a more substantial reduction in profits in Israel.

Three recent studies exploit state-level variation in minimum wages in the U.S., contrasting with the firm-level variation in exposure to nationwide minimum wage shocks used in the studies mentioned in the previous paragraph. In principle, state-level variation provides a more holistic picture of profitability than cross-firm evidence. [Vergara \(2023\)](#) examine the effect of minimum wages on state-level profit (gross operating surplus) per establishment and find a significant reduction in profits following minimum wage increases. Conversely, [Rao and Risch \(2024\)](#) study the impact of the minimum wage on independent (pass-through) businesses in exposed industries and find no indication of a significant reduction in profits. They document an increased rate of firm exit (which obviously represents a reduction in profits for those firms) but also an increase in profits among surviving firms. Finally, [Coviello et al. \(2022\)](#) investigate the impact of minimum wages on a large U.S. retailer. Depending on the specification, they find either a small and insignificant positive effect or a significant negative effect on profits, with estimates ranging from a 5% increase to a 16% decrease. Their estimates, along with the standard errors, can rule out profit reductions greater than 35%.

To summarize, most (though not all) studies find evidence of some reduction in profits. However, the magnitude of the profit reduction reported in the literature tends to be modest.

4.1.10 Worker turnover and reduction in training costs

A key adjustment channel predicted by various models of dynamic monopsony ([Burdett and Mortensen, 1998](#); [Manning, 2003](#)) is that minimum wages reduce worker turnover (separations). In this context, a higher minimum wage makes relatively lower-paying jobs more attractive, thereby reducing job-to-job separations. A substantial body of research has tested this prediction by empirically evaluating turnover responses.

Table 5 summarizes the key results on turnover from these studies. Of the 14 studies that directly examine the impact on turnover, 12 find a significant reduction in worker turnover rates. In one case, the evidence is mixed, with results depending on the specification, while another study finds no reduction. Therefore, the vast majority of the literature suggests that a reduction in turnover plays an important role in the response to minimum wage increases. This evidence also confirms a key prediction of the dynamic monopsony framework and, more broadly, underscores the role of labor market frictions in understanding the impact of minimum wages.

In addition to changing labor supply to the firm, lower turnover can also be a source of productivity improvement (evidence on overall productivity measures is discussed in more detail below). First, a lower turnover leads to savings in training and recruitment costs [Dube and Reich \(2010\)](#). Second, workers tend to be more productive when they stay longer at the same job and accumulate more experience. More research is needed to better understand the magnitude of these productivity improvements resulting from lower turnover.

Finally, the reduced turnover results suggest that we should be cautious when interpreting evidence on minimum wage effects on vacancy posting (see, e.g., [Kudlyak et al., 2023](#); [Clemens et al., 2021](#)), as these may simply reflect lower separations.³⁷

4.1.11 Productivity

The productivity enhancing effect of minimum wages has long been hypothesized. However, direct evidence on the impact on revenue has been scarce until recently.

Revenue-based productivity. The initial evidence on the productivity channel was based on firm-level revenue measures, such as revenue-based total factor productivity. Notable examples include [Mayneris et al. \(2018\)](#), [Riley and Rosazza Bondibene \(2017\)](#), and [Hau et al. \(2020\)](#), all of whom find a significant increase in revenue-based productivity following minimum wage changes. However, such an increase may simply reflect a rise in output prices rather than an actual increase in the quantity of goods or services produced. This concern is not merely hypothetical; as discussed in Section 4.1.7, the literature on output prices suggests that firms often respond to minimum wage hikes by raising prices. Researchers have been aware of this issue, and some have attempted to adjust for firm-level price changes by controlling for industry-level price indexes. However, these sector-wide measures are unlikely to capture the substantial variation in firm-level exposure within sectors, making the resulting estimates difficult to interpret.

Quantity-based productivity. Conceptually, a better approach would be to use a quantity-based measure of productivity that removes price effects, providing a clearer picture of actual efficiency improvements. In practice, however, such measures are usually not available for all or most firms in the economy. As a result, studies following this approach have provided evidence for a select set of sectors where such quantity-based efficiency calculations are feasible: [Ku \(2022\)](#) in agriculture, [Coviello et al. \(2022\)](#) in retail, and [Hau et al. \(2020\)](#) for exporting manufacturing firms. All these papers find evidence of increased quantity in response to the minimum wage. Additionally, a few recent studies have examined the impact on quality: [Ruffini \(2022\)](#) finds an increase in quality in the nursing home sector, while [Brown and Herbst \(2023\)](#) report improved service quality in the child-care setting.

Reallocation. Finally, as the discussion on entry and exit indicated, minimum wages could also affect productivity by improving the composition of jobs in the economy ([Mayneris et al., 2018](#)).

³⁷Reduced turnover can also mean that fewer workers are employed over the course of the year at a firm, even if there is no change in the number of jobs at any given time. This is particularly relevant for interpreting annual payroll data (see, e.g., [Rao and Risch, 2024](#)).

[Dustmann et al. \(2022\)](#) and [Rao and Risch \(2024\)](#) find evidence for such a channel in Germany and the U.S., respectively. Both studies show that in response to minimum wage increases, small, inefficient firms exited the market, and workers at these firms found employment at more productive firms. However, there could be nuanced welfare consequences from such reallocation. For example, while the exit of low-productivity employers might enhance worker productivity, it could also harm some consumers by reducing variety and diminishing product market competition.

Productivity and employment. How improved productivity affects employment is theoretically ambiguous, as it depends on the price elasticity of output demand. If output demand is highly elastic, increased production can be easily accommodated without changing the output price. On the other hand, in industries where output demand is inelastic, higher per-worker productivity means fewer workers are needed to meet consumer demand. In these sectors, some form of quality improvement that shifts the output demand outward would be necessary to explain limited employment effects.

4.2 Summary of evidence on margins of adjustment

The limited role of employment adjustment, as discussed in Section 3.3, suggests that other margins likely play a more significant role in responding to minimum wage increases. Based on the growing body of evidence reviewed above, we now have a better understanding of how these other margins of adjustment help firms absorb higher labor costs. Most importantly, passing the increased costs on to consumers in the form of higher prices plays a crucial role. This strategy, combined with the seemingly limited decline in consumer demand in response to price changes, covers a substantial share of the cost increase. While the exact proportion remains an open question, [Harasztosi and Lindner \(2019\)](#) suggest that around 80% of the change in labor cost can be attributed to this channel.

The evidence also indicates that improved efficiency is another key mechanism for absorbing higher labor costs. The productivity-enhancing effects can manifest in various ways, including lower training costs and turnover, improved operational efficiencies, and greater worker-level productivity. These factors help absorb some of the labor cost increase without a significant reduction in employment. Meanwhile, adjustments to profits appear to be limited, although the evidence on this margin is still evolving and somewhat tentative. For example, [Harasztosi and Lindner \(2019\)](#) suggest that around 20% of the labor cost increase could be offset by a reduction in profits.

These conclusions—drawn from studies that directly evaluate the causal impact of minimum wages on various margins of adjustment—are remarkably consistent with what business executives themselves report about their reactions to the policy. For example, executives generally perceive limited scope for reducing employment in response to minimum wage increases ([Lester, 1946](#); [Levin-Waldman, 2000](#); [Reich and Laitenen, 2003](#); [Bodnár et al., 2018](#)). At the same time, they almost always emphasize the importance of adjusting output prices and improving management and efficiency. Specifically, in [Lester \(1946\)](#)’s seminal study, executives highlight increasing sales efforts and enhancing productivity as the most relevant margins of adjustment. Similarly, [Reich and Laitenen \(2003\)](#) report survey evidence on managers in San Francisco, showing that the two most important responses to a minimum wage increase are raising prices and improving efficiency. [Hirsch et al. \(2015\)](#) find that restaurant

managers in Alabama and Georgia see a significant role for implementing cost-saving, performance-improving standards in response to the minimum wage. Finally, based on a representative sample of firms in eight Eastern and Central European countries, [Bodnár et al. \(2018\)](#) document that managers identify raising product prices, cutting non-labor costs, and improving efficiency as the most relevant channels of adjustment.

In summarizing the evidence on these channels of adjustment, it is important to acknowledge the main gaps in our knowledge. We have extensive evidence on employment (including substitution and firm dynamics), output prices, and worker turnover. There is also a growing body of work on fringe benefits, profits, and productivity, but further research in these areas would help refine our quantitative understanding of these adjustment margins. In contrast, the evidence on training is more limited and somewhat outdated, and there is very little research on service quality and consumer demand (e.g., overall sales). Additionally, we lack sufficient understanding of how rents or the prices of intermediate goods and services respond to minimum wage increases. These under-studied channels represent particularly important areas for future research.

4.3 Modeling implications and open questions

Evidence on these various margins of adjustment suggests that simple explanations for understanding the impact of minimum wages are insufficient. For instance, the competitive neoclassical model correctly emphasizes the role of price pass-through but fails to predict the lack of employment effects—except under the assumption of completely inelastic output demand. In contrast, the presence of imperfect competition in the labor market offers a natural explanation for the limited employment effects but falls short in accounting for a substantial increase in output prices or a limited pass-through to firm owners in the form of reduced profits.³⁸

The evidence on the productivity-enhancing effects of the policy also suggests that insights from institutionalist economists of the 1940s, as well as efficiency wage models, may capture important aspects of the policy response ([Hirsch et al., 2015](#)). Economic models of low-wage labor markets often presume full optimization within constraints ensured by competition and free entry of entrepreneurs. However, the pattern of productivity improvements following minimum wage hikes indicates that the discipline of market competition may not be strong enough to ensure that firms operate at the production possibilities frontier. Consequently, firms' X-efficiency could play an important role in understanding the impact of minimum wage policies ([Leibenstein, 1966](#)).

Given the current state of the literature, it would be overly ambitious for this review to propose a unifying framework for the low-wage labor market that incorporates all these considerations. Instead, we pose some open questions and puzzles that need to be addressed in future research.

³⁸In the standard monopsony framework, when minimum wages are increased slightly from the optimally set wage, the effect on profit will be zero. That framework also predicts increases in labor supply and lower output prices (see page 2108 in [Brown \(1999\)](#)). Interestingly, monopsonistic competition models, such as [Card et al. \(2018\)](#), predict a rent-sharing type of response to the minimum wage: if all firms are forced to raise wages, firm-level employment will not be affected, but profits will decline.

Imperfect competition. Over the past two decades, there has been considerable improvement in modeling the implications of minimum wages in the presence of various types of labor market frictions. This includes models that characterize minimum wage effects in the presence of search frictions (Van Den Berg, 2003; Flinn, 2006, 2010; Engbom and Moser, 2022; Manning, 2003; Dube et al., 2016; Brochu and Green, 2013), monopsonistic competition (Dickens et al., 1999; Manning, 2003; Haanwinckel, 2023), oligopsonistic competition (Bhaskar and To, 1999; Berger et al., 2022; Azar et al., 2023), or monitoring problems (Rebitzer and Taylor, 1995). Together, these studies provide a rich and flexible approach to analyzing the impact of the minimum wage on labor market outcomes such as employment level, job flows, and the wage distribution.

In many cases, the presence of friction dampens the employment effects of the policy relative to a frictionless world, but these models often predict considerable dis-employment under realistic parameterization. This result is partly related to the imposition of the free-entry and zero-profit assumption, which leaves very little room for rent sharing or profit reduction to occur in the constrained equilibrium. However, most of these models focus solely on labor market interactions and tend to ignore adjustment on other margins. For example, the models often overlook the extent to which firms can pass through the minimum wage to consumers without a change in output, or the productivity-enhancing effects of the policy. Therefore, a natural next step would be to incorporate these dimensions in existing models to enrich our understanding of how imperfections in the labor market interact with imperfections in the product market.

Productivity. To guide modeling choices, we first need a better understanding of the relative importance of various sources contributing to the productivity-enhancing effects of minimum wages: (i) worker effort (efficiency wage) mechanisms, (ii) turnover reduction, (iii) reductions in X-inefficiency or improved management practices within firms, and (iv) reallocation from low- to high-productivity employers, including the exit of low-productivity firms. Additionally, it is crucial to determine whether these effects are driven by the market-wide nature of minimum wage shocks (i.e., all firms being affected) or if the productivity enhancement would also be present in response to firm-level shocks.

Exploring the interaction between labor market imperfections, such as monopsony power, and X-efficiency is a particularly intriguing direction for future research. In markets characterized by imperfect competition, both in labor and product markets, firms making suboptimal choices can still survive. Consistent with this idea, recent evidence suggests that significant productivity improvements are possible through the implementation of better management practices (Bloom et al., 2013), especially when competitive forces are not strong enough (Bloom et al., 2014). Moreover, there is some evidence that substantial productivity gains can occur even with firm-specific pay raises, and in some cases, firms may benefit more from these efficiency gains than they lose from paying higher wages to their workers (Emanuel and Harrington, 2020).

Relatedly, a growing body of evidence shows that firms often fail to set prices and wages optimally, instead opting for relatively “clunky” decision rules. National chains frequently use uniform pricing in both product and labor markets (DellaVigna and Gentzkow, 2019; Hazell et al., 2022). The presence of such “mispricing” creates additional opportunities for productivity improvement from a wage floor. For instance, Coviello et al. (2022) document significant productivity effects and an increase in profits

following local-level minimum wage hikes for certain stores of a U.S. nationwide retailer that employs a uniform compensation structure across stores (except for compliance with local minimum wages) along with nationally uniform product pricing.³⁹

Interestingly, the presence of imperfect competition offers a rationale for the prevalence of such optimization frictions. When wages are a choice variable, modest optimization errors have only second-order consequences on profits (due to the envelope theorem), even as they have a first-order impact on workers (see [Dube et al. \(2020\)](#) for more). More research is needed on the impact of firm-level wage shocks on productivity and how this relationship relates to the pre-existing strength of competition.⁴⁰

Consumer demand. As we have seen, a key channel through which firms respond to the minimum wage is by adjusting output prices. Therefore, understanding how output demand changes in response to minimum wage increases is crucial. The apparent contradiction between limited employment effects and significant output price effects has long been recognized in the literature (see, e.g., [Brown, 1999](#); [Aaronson and French, 2007](#)). As equation (1) in Section 2.2 shows, a price increase typically leads to a decrease in output demand, which in turn reduces the demand for workers.⁴¹

Therefore, the increase in prices and the limited employment effects are consistent with inelastic output demand (see, e.g., [MaCurdy \(2015\)](#)). However, there are several explanations for this inelastic consumer demand, each with different implications for welfare. First, it is possible that consumer demand for products that rely heavily on minimum wage labor is *generally* inelastic. This could be true at the market level even if individual consumers are more price-elastic. For example, in the fast-food industry, a market-level price increase might deter price-sensitive consumers but attract higher-spending consumers who previously avoided longer wait times. As a result, the industry as a whole might experience relatively more inelastic output demand due to a changing consumer base.

Another possibility is that this insensitivity is a *specific* feature of the price increases induced by a higher minimum wage. Below, we outline some leading explanations within this category.

First, an important consideration lies in the interaction between product and labor market imperfections. [Bhaskar and To \(1999\)](#) show that minimum wages can alleviate imperfections in the labor market while exacerbating imperfections in the output market. Specifically, firm exit can increase market concentration, allowing surviving firms to raise prices by more than the increase in labor costs. This prediction aligns with empirical evidence showing more than full pass-through of labor costs in certain contexts (see Table 4). In [Bhaskar and To \(1999\)](#)’s framework, there is a negative relationship

³⁹[Coviello et al. \(2022\)](#) find a positive profit effect and a significantly larger increase in productivity when focusing on a subset of stores near the borders of minimum wage changes. When they evaluate the policy’s impact on a more representative set by including all stores across all states, they find smaller productivity increases and a slight reduction in profit. This divergence is consistent with uniform prices being optimized for nationwide profits but not necessarily for local markets.

⁴⁰Valuable insights could be gained from the case study approach of institutional economists or from a better understanding of how decisions within companies are made—a focus of organizational economics (see, e.g., [Gibbons and Roberts, 2015](#)).

⁴¹One interesting explanation for this phenomenon comes from the putty-clay framework described in Section 4.1.4. [Aaronson et al. \(2018\)](#) shows that the model predicts immediate output price effects but delayed employment reductions. However, much evidence suggests that medium-term employment changes are also limited, indicating that this explanation likely plays a limited role.

between output prices and output demand. However, in the presence of significant substitution between low-skilled labor and other inputs, they demonstrate that minimum wage increases could boost both employment and output prices. Intuitively, the minimum wage encourages firms to hire more workers as they move along the upward-sloping labor supply curve. If the additional workers allow firms to reduce capital usage, then overall output may decrease. Thus, the model predicts higher employment but lower capital in this knife-edge scenario. However, this prediction does not align with the empirical pattern observed in Section 4.1.3, where capital use appears to increase following minimum wage hikes.

Another possibility is that minimum wage shocks are salient enough to enable firms to implement coordinated price changes. Such minimum wage-induced price increases might be perceived as fairer and more justified, leading consumers to tolerate them more. Additionally, minimum wage increases could improve overall product or service quality in the sector, which might attract additional consumer demand.

Finally, a common explanation for the muted output response relies on income effects. This argument suggests that a higher minimum wage increases the purchasing power of low-skilled workers, who are sometimes assumed to consume low-quality, minimum wage-intensive services and goods. This explanation, sometimes referred to as the “Hungry Teenager Theory” (see, e.g., [Kennan \(1995\)](#)), seems less plausible than the other explanations discussed so far. Given that a small share of income is typically spent on low wage-intensive products (e.g., 5-7% across countries), this limits the extent to which added income could fuel demand for such products. Furthermore, existing evidence on consumption responses to minimum wage increases suggests that low-wage workers tend to spend their extra income on durable goods such as vehicles, rather than on fast food (see [Aaronson et al. \(2012\)](#)).

Understanding the source and nature of output demand inelasticity has important welfare implications. For example, if inelastic demand reflects improved quality, the price change may not result in a welfare loss. In contrast, if it indicates a lack of alternatives for consumers to substitute away from low-wage-intensive goods and services, the price effect could represent a pure welfare loss. Gaining a better understanding of the nature of the price response will be an important focus for future research on this topic.

Substitution towards other inputs. Another open question is why the standard substitution channel between lower-skilled workers and other inputs appears to be dampened in the context of the minimum wage. A simple and straightforward explanation could be related to the production function: in the minimum wage sector, the technology may not allow for substitution between lower-skilled workers and other inputs.

A more fruitful approach is to explore the interaction between technology and imperfect competition in the labor market. [Datta and Machin \(2024\)](#) provide an excellent recent example of this, showing that in the presence of imperfect competition, firm-specific wage floors can lead to an increase in the employment of low-skilled workers and a decrease in the demand for high-skilled ones—the opposite prediction from the neoclassical framework. The intuition is that under monopsony, a higher relative

wage for low-skilled workers due to a wage floor has both labor supply and labor demand effects, which influence relative employment in opposite directions. [Datta and Machin \(2024\)](#) also estimate the technology parameter after accounting for firms’ wage-setting behavior and find that, in their setup, capital and labor are gross complements.

5 Inequality, distributional implications, and downstream effects

We have seen a considerable increase in wage inequality since the early 1980s in many countries, including the United States ([Song et al., 2018](#)), United Kingdom ([Giupponi and Machin, 2022a](#)), and Germany ([Dustmann et al., 2009](#); [Card et al., 2013](#)). In some cases—especially the United States—this occurred concurrently with a fall in the real federal minimum wage ([DiNardo et al., 1996](#); [Card and DiNardo, 2002](#)).⁴² More recently, we have also seen a reduction in wage inequality in the U.S., concurrently with a rise in state-level minimum wages in parts of the country ([Autor et al., 2023](#)). Wage inequality also declined in Germany and the U.K. over the past decade, at the same time minimum wages were introduced or expanded in both countries ([Bossler and Schank, 2023](#); [Machin, 2024](#)). These descriptive trends raise the question: have minimum wage policies had a substantial impact on wage inequality?

At a *qualitative* level, the idea that minimum wages reduce wage inequality is uncontroversial. After all, conditional on having a job, a higher minimum wage raises the bottom wage, which will naturally reduce pay dispersion. However, if that were the only effect of the minimum wage policy, its *quantitative* impact on wage inequality would likely be quite modest in most cases, since the share of workers earning exactly the minimum wage is relatively small in most countries. For instance, 7.3 percent of the U.S. workforce earned at or below the federal minimum wage in 2010, just after the last federal increase. This imposes bounds on any direct effect. However, it is also possible that a higher minimum wage leads to a higher wage for those already earning above the new minimum. Such spillovers (or ripple effects) up the pay distribution could considerably increase the scope for minimum wages to affect the wage distribution and inequality. Understanding the size of these spillovers is critical for discerning the policy’s importance in determining inequality.

Theoretically, several factors can explain the presence of spillovers. One common idea is that the minimum wage shrinks the wage gap between front-line and higher-paid workers (such as supervisors), which can give rise to concerns related to fairness or incentives. As a result, companies have to adjust the pay of supervisors and others to maintain a wage hierarchy. Existing evidence suggests that such social comparisons matter at the workplace (e.g., [Dube et al. \(2019\)](#); [Gopalan et al. \(2021\)](#)). In addition to within-firm considerations, there may also be market-based mechanisms for spillovers. In an imperfectly competitive labor market, a higher minimum wage can raise the outside options for workers earning slightly above the minimum (see e.g. [Flinn \(2010\)](#)). If workers at better-paying employers were to lose their jobs and be forced to take a lower-paying position, a higher minimum wage could raise the value workers assign to being unemployed (assuming the probability of finding

⁴²To understand overall trends in inequality, it is important to distinguish between lower- and upper-tail inequality ([Autor et al., 2008](#)). Minimum wage is important to explain lower-tail inequality patterns, but the upper-tail changes are likely to be driven by other forces (e.g. skill bias technological change).

a job is not too severely harmed). This, in turn, can force better-paying employers to adjust their wages upwards, causing a ripple effect higher up the distribution (see [Flinn \(2006\)](#); [Butcher et al. \(2012\)](#); [Engbom and Moser \(2022\)](#)). Spillovers could also emerge from substitution with somewhat higher skilled workers ([Teulings, 2000](#)). Finally, there are more specific factors that can amplify wage spillovers, such as compensating differentials ([Phelan, 2019](#)), or employers' tendency to set pay at round numbers ([Dube et al., 2020](#)).

Empirical evaluations have broadly found evidence consistent with minimum wage spillovers, although the estimated magnitudes sometimes differ across studies. [DiNardo et al. \(1996\)](#) provided an early and useful illustration of how much minimum wages could impact the wage distribution using a novel decomposition method. The modern quasi-experimental literature began with [Lee \(1999\)](#), who used variation in the bite of (both federal and state) minimum wage across states in the U.S. between 1979-1988 to quantify the effect of minimum wages on wage inequality. He found sizable ripple or spillover effects, where the rise in the minimum raised pay for the bottom 40 percent of wage earners. Strikingly, his findings suggested that nearly all of the growth in inequality at the bottom half of the wage distribution during the 1980s could be explained by the erosion of the federal minimum wage through inflation, similar to decomposition-based results of [DiNardo et al. \(1996\)](#).

However, there were some puzzling aspects of [Lee \(1999\)](#)'s findings, which seemed to suggest that a higher minimum wage not only raised pay at the bottom compared to the median, but also at the top. In his specification, the median wage appeared in both the dependent variable (log wage at p -th percentile minus log of median wage) and the independent variable (log of minimum wage minus the log of median wage). This raised a concern of division bias, which would tend to yield a spurious positive relationship between the two variables (Lee himself was clear about this possibility).

Subsequent work by [Autor et al. \(2016\)](#) provides evidence consistent with such division bias. Their evidence, largely based on state-level minimum wage increases over a more recent period, uses an instrumental variables strategy to correct for the division bias. Their IV estimates suggest that a 10% increase in the minimum raises the 10th percentile wage by around 1.5%. Additionally, they find some spillover effects extending up to around the 20th or 25th percentile, beyond which the wage effects are close to zero. They also find that minimum wages played an important role in determining the 50/10 gap—which is a measure of wage inequality in the bottom half of the distribution. But compared to [Lee \(1999\)](#), the spillovers found in the [Autor et al. \(2016\)](#) study were smaller, leading to a smaller impact on inequality. Still, they found that maintaining the minimum wage at the 1979 level in real terms would have prevented somewhere between half and three-quarters of the overall increase in the bottom-half wage inequality, depending on the period in question. Moreover, the minimum wage had a larger effect on inequality for female workers, who tend to be lower paid.

More recent work by [Fortin et al. \(2021\)](#) offers a complementary perspective that focuses on possible heterogeneity in the spillover effects. They point out that the U.S. federal minimum wage in early 1980s was much more binding than state minimum wages studied by [Autor et al. \(2016\)](#). If the size of the spillovers depends on how binding the minimum wage is, this could offer another explanation behind the relatively larger evidence found in [Lee \(1999\)](#). [Fortin et al. \(2021\)](#) develop a hybrid approach that uses both state-level policy variation and the bite of the federal minimum

wage, as well as the exact location in the state-level wage distribution where the minimum wage is binding.⁴³ This allows them to better evaluate the changes in the 1980s when the federal minimum wage eroded greatly, but there were only a handful of state-level changes.⁴⁴ Their design suggests that the spillover effects during the 1980s (studied by Lee (1999)) were indeed larger than those from subsequent state policies studied by Autor et al. (2016). Their findings highlight that there may not be “one true estimate” of minimum wage effects, as these can vary based on factors such as the bite of the policy.

One limitation of much of the evidence on wage spillovers is the inability to distinguish true spillovers from job loss. When the minimum wage rises from \$13 to \$15, if all the jobs under \$15 are destroyed, simply comparing the pre- and post-intervention wage quantiles would show a rise in wages at quantiles initially above \$15. This would happen even if none of the jobs initially paying above \$15 raised their wages. Why? Because the disemployment truncates the wage distribution at \$15, changing the composition of jobs at various quantiles. Most studies in the literature acknowledge this possibility but assume that the disemployment effect is close to zero, thereby limiting any bias. Cengiz et al. (2019) show that by estimating the impact of the policy on the *frequency* distribution of wages, one can jointly estimate the employment and wage effect—thereby accounting for potential disemployment effects when measuring wage spillovers. They find economically meaningful, but limited, spillover effects from state minimum wages that constitute around 40 percent of the overall wage increases from the policy. Additionally, Cengiz et al. (2019) show that the measured wage spillovers are not likely driven by measurement error in survey wages, a possibility raised by Autor et al. (2016). Evidence of spillovers from payroll data used by Gopalan et al. (2021) further supports these conclusions.

Similarly, evidence from the U.K. by Giupponi et al. (2024) based on the frequency distribution approach suggests a moderate-sized spillover effect from the National Living Wage policy. Research from Germany by Bossler and Schank (2023) also indicates that the introduction of the national minimum wage in Germany substantially reduced wage inequality, estimating that about half of the recent reduction in wage inequality can be attributed to the minimum wage.⁴⁵

Overall, the body of evidence consistently suggests that higher minimum wages tend to reduce wage inequality in the bottom half of the wage distribution. However, the magnitudes likely vary by the nature of the policy environment.

Distributional-implications. While the impact of minimum wages on *wage inequality* is relatively straightforward, the impact on *income inequality* and poverty is more complicated. First, most poor families have weak ties to the labor market, which limits the scope for minimum wages to raise bottom incomes. The minimum wage could also, of course, affect the probability of work—possibly heterogeneously by latent family incomes. Second, the lowest-wage workers may not always be in the lowest-income families. Finally, a rise in the minimum wage can have complicated interactions with other safety net and tax/transfer policies. As a consequence, a rise in minimum wage can be partly

⁴³Their approach shares similarities with Cengiz et al. (2019) and Giupponi et al. (2024), though unlike those papers, they are focused on the probability and not frequency distribution of wages—hence abstracting from employment impacts.

⁴⁴One caveat is that the 1980s constitute a difficult period to study due to confounding shocks we discussed before, as shown, for example, in Figure 3.

⁴⁵We review the evidence on spillovers from developing countries in Section 6

offset by reduced public transfers, thereby reducing the rise in post-tax-and-transfer income at the bottom.

The evidence on poverty reduction is mixed. Some U.S. research have found negative effects on state-level poverty rates (Addison and Blackburn, 1999; Dube, 2019b; Godoey and Reich, 2021). Other studies have not found evidence of overall poverty reduction (Burkhauser et al., 2023; Sabia and Nielsen, 2015). Interestingly, Godoey and Reich (2021) finds that while the average minimum wage increase did not reduce overall poverty rates, more binding increases did. There is also evidence that a higher minimum wage reduces public transfers, suggesting a rise in pre-tax household incomes at the bottom ((Dube, 2019b; Vergara, 2023; Reich and West, 2015).

However, poverty is likely too limited a measure for evaluating distributional impacts. We need to better understand how the minimum wage affects both the pre- and post-tax/transfer income distribution, and their components. Work by Neumark et al. (2005) represents an early approach along these lines, studying a short panel of individuals. More recently, Dube (2019b) provides estimates of the policy impact on unconditional income quantiles is an effort in this direction using the TWFE approach and re-centered influence functions. However, more work is needed on this topic, especially using event-based analysis, with attention to possibly heterogeneous impact by how binding the minimum wages are.

Moreover, to the extent that minimum wage workers are both relatively few and somewhat dispersed across the distribution (though more concentrated at the bottom), it is difficult to statistically detect the welfare effects of the policy by studying the entire income distribution. Using a predictive approach to identify families likely to be affected and studying the impact on their incomes could be a fruitful avenue going forward.

Finally, there is even less evidence from outside the U.S., making additional research crucial. One notable exception is (Giupponi et al., 2024), who study the distributional effects of minimum wage policies in the British context. They use a micro-simulation model of the U.K. tax-transfer system and incorporate their estimated employment and wage effects (including spillovers). Assuming there is no substantial heterogeneity in the response to the minimum wage across household income deciles—which aligns with the heterogeneity results discussed in Section ??—this approach provides an accurate picture of the policy’s distributional effects. While this method requires stronger assumptions than standard reduced-form analysis of household income, it offers valuable insights into the interaction between the tax and transfer system and minimum wages. Additionally, their approach can provide more accurate estimates by filtering out fluctuations in household income driven by the tax and benefit system.

Price effects and welfare. So far, we have discussed the effect of minimum wages on family income distribution. To fully understand the impact of minimum wages on family welfare, we also need to consider consumption responses. As discussed in Section 4.1.7, an important margin of adjustment to the minimum wage is increased output prices, which could affect consumer welfare and should be taken into account.

The first important question is whether the price increase reflects a change in product or service

quality. Ideally, we would focus on quality-adjusted price changes, but such data are rarely available. As discussed in Section 4.1.7, our knowledge of quality changes following minimum wage hikes is limited. However, the few studies on this topic suggest some improvements in quality, although these improvements may not fully compensate consumers (Brown and Herbst, 2023).

The second relevant question is the effect of (quality-adjusted) price increases on consumers' budgets. Since minimum wage workers make up a small share of the economy, the increased burden on consumers is relatively modest. MaCurdy (2015) estimates that a 20% increase in the minimum wage raises consumer expenditures by 0.5%, a figure similar to what Harasztosi and Lindner (2019) estimate for Hungary. It's important to note that these small effects are likely to be an upper bound, given that these calculations typically assume no change in product quality.

The third important question is whether exposure to price shocks differs across the household income distribution. Studying overall consumption responses, both MaCurdy (2015) in the U.S. and Harasztosi and Lindner (2019) in Hungary find little evidence of differential exposure to price shocks across household income quantiles, suggesting that the burden of the minimum wage is shared relatively equally across the income distribution. On the other hand, Renkin et al. (2022) use detailed data on grocery store expenditures to analyze how different income groups are affected by economic shocks. Their findings suggest that low-income households experience a disproportionately larger relative impact compared to higher-income households. However, it is unclear whether these findings are specific to the grocery store sector or reflect the more detailed spending data used in their analysis compared to MaCurdy (2015) and Harasztosi and Lindner (2019).

Several studies examine the impact of minimum wage increases on household consumption. Aaronson et al. (2012) find that households with minimum wage workers tend to increase their consumption more than their earnings in the U.S. context. This response is largely driven by a small number of households making significant purchases, such as vehicles. To finance this increased consumption, these households often turn to borrowing. The increased borrowing aligns with findings from other studies, which suggest that minimum wage hikes lead to greater availability of unsecured credit, a reduction in payday loan usage, decreased delinquency and bankruptcy rates, and improved credit scores (Detting and Hsu, 2020; Legal and Young, 2024). Overall, these findings highlight that the financial health of low-income households improves following minimum wage changes, at least in the U.S. context.

Furthermore, Dautović et al. (2024) find similar positive effects in China, where minimum wage increases boosted consumption, particularly in households with children. Similarly, Mansoor and O'Neill (2021) finds that higher minimum wages tend to raise household consumption in India, especially in areas with high compliance with minimum wage policies. These studies further corroborate the positive impact of minimum wage policies on consumption, underscoring their broader benefits.

Indirect effects of minimum wages. Minimum wages can influence many other socioeconomic outcomes through their direct impact on wages, household income, and consumption. These downstream implications of minimum wage policies are particularly context-dependent. For example, the effect of a minimum wage-induced positive income shock on health outcomes will depend on the structure of

the healthcare system and the broader welfare system. As a result, it is unlikely that the downstream effects of minimum wage policies are uniform across different societies and social classes.

With this caveat in mind, it is worth briefly discussing the findings on some key downstream effects of minimum wage policies. [Leigh et al. \(2019\)](#) review evidence on health outcomes as of 2018 and find a beneficial impact of minimum wages on smoking prevalence, but no consistent evidence on other health outcomes. [Neumark \(2024\)](#) reviews 57 papers on the impact of minimum wages on various health outcomes and 6 papers on the impact on crime. He provides a detailed description of each paper along with a subjective assessment of the credibility of the estimates. The findings suggest that minimum wages positively influence infant and child health, diet and obesity, depression and mental health, suicide rates, and family structure, while the evidence on teen and adult health is more mixed.⁴⁶ The evidence on crime is also mixed, with three studies showing a reduction in crime, one finding no effect, and two reporting an increase in nonviolent crime.

The effect of minimum wage policies on educational outcomes and human capital accumulation could have significant welfare implications. The evidence on high school completion is mixed. [Neumark and Wascher \(1995a,b, 2003\)](#) report that higher minimum wages depress school enrollment, while [Campolieti et al. \(2005\)](#) and [Smith \(2021\)](#) find no significant changes in dropout rates. Similarly, [Chaplin et al. \(2003\)](#) find no significant reduction in high school enrollment, particularly for students subject to compulsory schooling laws. This latter finding highlights that the impact of minimum wages on schooling is highly dependent on the education system.

A number of recent studies have examined the impact on post-secondary schooling decisions. [Lee \(2020\)](#) finds a substantial reduction in enrollment rates at community colleges in response to minimum wage increases. Despite sizable declines in enrollment at two-year colleges, [Schanzenbach et al. \(2024\)](#) find no reduction in degree completions. This suggests that the social costs associated with lower enrollment could be small (or even socially beneficial), as workers discouraged from enrolling would not have completed their degrees anyway. In the Canadian context, [Alessandrini and Milla \(2024\)](#) find a somewhat different pattern: an increase in enrollment at community colleges but lower enrollment at universities. This discrepancy with the U.S. evidence could be explained by the fact that community colleges in Canada focus almost exclusively on vocational education, while most two-year institutions in the U.S. also provide academic training. Again, this points to the highly context-dependent aspect of the minimum wage impact on educational outcomes.

Overall, a large body of evidence produced over the past two decades suggests that the impact of minimum wages extends beyond merely raising the living standards of low-wage workers. At the same time, the indirect effects of minimum wage policies are strongly influenced by the broader institutional structure. Gaining a better understanding of these interactions is an important direction for future research.

⁴⁶The health implications of minimum wages seem to align with those of other policies that increase the net income of low-wage workers (see [Godøy and Jacobs \(2021\)](#)).

6 Minimum Wages in developing countries

While much of the research on minimum wage policies has focused on the United States and Europe, minimum wages in developing countries have gained increased attention in recent years. There are several issues specific to the developing country context, which we discuss in this section.

Developing countries often face higher levels of informality in their labor markets, complicating the enforcement and effects of minimum wage laws. Raising the minimum wage could incentivize employers to shift work to the informal sector, potentially harming workers by placing them in jobs that often lack health insurance, social security benefits, or job protection. However, the lack of enforcement in the informal sector does not necessarily mean that firms can simply ignore minimum wages in that sector. As discussed in Section 4.1.1, even in developed economies, enforcement is unlikely to be the primary factor behind compliance (see [Stansbury \(2024\)](#)). Therefore, the influence of minimum wages on the informal sector is an empirical question, which we explore in the following discussion.

The second major difference between developed and developing countries is the sectoral composition of minimum wage workers. In most developed economies, minimum wage workers are typically employed in the local service sector, whereas in many developing economies, they are often found in the tradable sector, exporting to the world market. As discussed in Section 3.3.3, the effect of the minimum wage can vary considerably across industries, with more negative employment responses observed in the tradable sector.

Another important consideration is the presence of more significant imperfections and inefficiencies in the developing country context ([Hsieh and Klenow, 2009](#); [Bloom et al., 2014](#)). The empirical evidence presented in Section 4.1.11 suggests that firms in developed countries often find ways to improve productivity in response to minimum wage increases. We suspect that the scope for efficiency improvements is even more substantial in developing economies, making this margin of adjustment potentially more important.

These considerations highlight that it is *a priori* unclear what to expect regarding the impact of minimum wage policies in developing countries. The empirical evidence on the overall employment effects of minimum wages in developing economies is mixed. Two recent reviews have examined the employment effects in low- and middle-income countries. [Neumark and Munguía Corella \(2021\)](#) review 61 papers on the impact of minimum wage policies in developing contexts and find mixed evidence on employment effects. They observe more consistently negative employment effects in formal sectors, particularly where the minimum wage is more binding and strictly enforced. On the other hand, [Broecke et al. \(2017\)](#) review the impact of minimum wages on employment in 14 major emerging economies and find that minimum wages have minimal impact on employment, with evidence of reporting bias toward statistically significant negative results. More vulnerable groups (e.g., youth and low-skilled workers) are found to be slightly more negatively affected, and there is some evidence that higher minimum wages lead to increased informal employment.

However, it is important to note that both reviews focus on the employment elasticity with

respect to the minimum wage rather than the own-wage elasticity. This distinction makes it difficult to compare estimates across groups, studies, or with the extensive evidence from developed economies that we reviewed in Section 3.3.

Less attention has been given to studying various other margins of adjustment in developing countries. However, a few papers, which we reviewed in Section 4, present findings consistent with those observed in developed countries. For instance, [Lemos \(2006\)](#) finds a positive effect on output prices in Brazil, while [Mayneris et al. \(2018\)](#) and [Huang et al. \(2014\)](#) find evidence of significant productivity improvements in China.

The evidence from developing countries also suggests that minimum wage increases can substantially alleviate inequality. [Bosch and Manacorda \(2010\)](#) find that the erosion of minimum wages in Mexico contributed to a significant increase in wage inequality between the late 1980s and early 2000s, particularly at the bottom of the distribution. In India, research exploiting province-level variation finds that minimum wages have contributed substantially to a decline in wage inequality since the beginning of the century [Khurana et al. \(2023a\)](#). In Brazil, minimum wage increases compressed wages in both the formal and informal sectors ([Lemos et al., 2004](#); [Derenoncourt et al., 2021](#)). Notably, the estimated spillover (ripple) effects of the Brazilian minimum wage policy appear to be significantly larger, amplifying the contribution of minimum wages to reducing inequality. [Engbom and Moser \(2022\)](#) estimate that minimum wages affected pay up to the 90th percentile. However, [Haanwinkel \(2023\)](#), after controlling for other supply and demand shocks, finds significant but more muted effects, extending up to the second decile from the bottom.

Interestingly, minimum wages also affect inequality in the informal sector, even though informal workers are not legally entitled to the minimum wage (see [Lemos et al. \(2004\)](#); [Derenoncourt et al. \(2021\)](#)). This phenomenon, sometimes referred to as the “lighthouse” effect, underscores the broader influence of minimum wage policies beyond their immediate legal scope. While understanding why minimum wages influence wages in sectors with limited enforcement requires further research, these findings suggest that firms cannot easily circumvent minimum wage regulations by shifting to the informal sector.

To summarize, the impact of minimum wage policies in developing countries is influenced by various factors, such as a high degree of informality, the sectoral composition of low-wage workers, and more pronounced market imperfections. At the same time, while the empirical literature is still growing, the evidence gathered so far from developing countries does not suggest that the core impacts differ considerably from those observed in developed economies.

7 Conclusion and future directions

Minimum wage policies have been the subject of extensive research for over a century, and there remains considerable interest in understanding their impact. Given the policy’s fundamental relevance for testing basic economic theories, it has been at the forefront of economic thinking, reflecting (and influencing) the evolution of both theories and empirical methods. At the same time, during this

period the debate has often centered on the question of whether we should have a minimum wage at all.

We reviewed the most important evidence on how minimum wages shape the labor market and other downstream outcomes, focusing on contributions made since the beginning of the 21st century. During this period, a substantial body of empirical evidence has accumulated regarding the employment effects of the policy and its broader economic and social impacts. While the evidence is not unanimous, a reasonable conclusion from the existing literature is that minimum wage policies have had limited direct employment effects while significantly increasing the earnings of low-wage workers—at least at certain levels and in particular economic contexts. We see this as suggesting that minimum wages can be beneficial in many situations and should be considered a key economic tool for intervening in low-wage labor markets and improving economic outcomes.

However, the fact that minimum wages are effective at certain levels and in certain contexts does not imply they always work. Research in the 21st century should focus on determining the appropriate levels of the minimum wage, rather than debating the existence of the policy itself. Nonetheless, answering this question presents a number of challenges that researchers will need to overcome to make progress.

First, we need a better understanding of the heterogeneous impact of the policy across different economic contexts. Identifying turning points where the minimum wage begins to significantly affect employment dynamics is an essential next step. A common theme throughout our chapter is the difficulty in characterizing heterogeneous effects. Estimating the average causal response to minimum wages is challenging enough, and sometimes requires strong assumptions for identification. Obtaining reliable estimates by subgroups (such as different types of workers or policy levels) typically require even stronger assumptions and may lack statistical power. Additionally, due to publication bias towards statistically significant results, the search for heterogeneous treatment effects may be more prone to false discovery. Therefore, research on heterogeneity requires careful and transparent analyses, and a variety of econometric and data challenges need to be addressed.

Second, a particularly important recent advance in the minimum wage literature is the use of a more transparent event study design rather than the relatively opaque TWFE panel regression. Future work should rely on event study designs to provide more precise estimates for each major minimum wage event. Accumulating knowledge about these events will be crucial for properly understand the nature of heterogeneity in minimum wage effects. The literature can also advance by providing more details about the treated and potential control units and how they were selected.

Third, to determine the appropriate levels of the minimum wage, we need to improve our models of low-wage labor markets. The standard competitive model often fails to capture the complexity of real-world labor dynamics, suggesting a need for more nuanced economic frameworks. Building such models requires a better understanding of the various margins of adjustments, especially in areas with relatively thin literature, as identified in Section 4. These economic models should account for consumers' demand responses and the productivity-enhancing effects of the policy, along with incorporating labor market imperfections. However, it is important to avoid past mistakes.

Models should be designed to match the best available evidence, instead of primarily being used for extrapolation. At a fundamental level, core theories of the labor market should be falsifiable, and we should not dismiss empirical findings that are difficult to reconcile with existing models.

Lastly, to set the optimal level of the minimum wage, it is crucial to define what “appropriate” means. This is a complex question, as many of the welfare consequences of the policy may depend on its interactions with other policy instruments and the political feasibility of changing them. Future research should better incorporate the interaction of minimum wages with key taxes, benefits, and government policies into the analysis.

In conclusion, while the study of minimum wage policies has made significant strides, there remains ample room for further research to refine our understanding of their economic impacts. By addressing the areas outlined and advancing methodological approaches, researchers can contribute to more informed policy decisions that balance wage fairness with economic efficiency.

References

- Aaronson, Daniel (2001) “Price Pass-through and the Minimum Wage,” *The Review of Economics and Statistics*, 83 (1), 158–169. [101](#)
- Aaronson, Daniel, Sumit Agarwal, and Eric French (2012) “The spending and debt response to minimum wage hikes,” *American Economic Review*, 102 (7), 3111–3139. [60](#), [65](#)
- Aaronson, Daniel and Eric French (2007) “Product market evidence on the employment effects of the minimum wage,” *Journal of Labor Economics*, 25 (1), 167–200. [59](#)
- Aaronson, Daniel, Eric French, and James MacDonald (2008) “The Minimum Wage, Restaurant Prices, and Labor Market Structure,” *The Journal of Human Resources*, 43 (3), 688–720. [101](#)
- Aaronson, Daniel, Eric French, Isaac Sorkin, and Ted To (2018) “Industry Dynamics and the Minimum Wage: A Putty-Clay Approach,” *International Economic Review*, 59 (1), 51–84. [48](#), [49](#), [59](#), [100](#)
- Aaronson, Daniel and Brian J Phelan (2019) “Wage shocks and the technological substitution of low-wage jobs,” *The Economic Journal*, 129 (617), 1–34. [47](#), [100](#)
- Abowd, John M, Francis Kramarz, Thomas Lemieux, and David N Margolis (2000) “Minimum Wages and Youth Employment in France and the United States,” in *Youth Employment and Joblessness in Advanced Countries*, 427–472: University of Chicago Press. [31](#)
- Acemoglu, Daron and Jörn-Steffen Pischke (2003) “Minimum wages and on-the-job training,” in *Worker Well-Being and public policy*, 22, 159–202: Emerald Group Publishing Limited. [46](#), [99](#)
- Adams, Camilla, Jonathan Meer, and CarlyWill Sloan (2022) “The minimum wage and search effort,” *Economics Letters*, 212, 110288. [51](#), [100](#)
- Addison, John T and McKinley Blackburn (1999) “Minimum wages and poverty,” *ILR Review*, 52 (3), 393–409. [64](#)
- Agan, Amanda Y. and Michael D. Makowsky (2021) “The Minimum Wage, EITC, and Criminal Recidivism,” *Journal of Human Resources*. [51](#), [100](#)
- Agarwal, Sumit, Meghana Ayyagari, and Renáta Kosová (2024) “Minimum Wage Increases and Employer Performance: Role of Employer Heterogeneity,” *Management Science*, 70 (1), 225–254. [52](#), [101](#), [102](#)
- Ahlfeldt, Gabriel M., Duncan Roth, and Tobias Seidel (2018) “The regional effects of Germany’s national minimum wage,” *Economics Letters*, 172, 127–130. [50](#), [100](#)
- Alessandrini, Diana and Joniada Milla (2024) “Minimum-Wage Effects on Human Capital Accumulation: Evidence from Canadian Data,” *Journal of Human Capital*, 18 (2), 000–000. [66](#)
- Allegretto, Sylvia A, Arindrajit Dube, and Michael Reich (2011) “Do minimum wages really reduce teen employment? Accounting for heterogeneity and selectivity in state panel data,” *Industrial Relations: A Journal of Economy and Society*, 50 (2), 205–240. [13](#), [14](#), [21](#), [38](#), [41](#)

- Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer (2017) “Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher,” *ILR Review*, 70 (3), 559–592. [17](#)
- Allegretto, Sylvia and Michael Reich (2018) “Are Local Minimum Wages Absorbed by Price Increases? Estimates from Internet-Based Restaurant Menus,” *ILR Review*, 71 (1), 35–63. [101](#)
- Alonso, Cristian (2022) “Beyond Labor Market Outcomes,” *Journal of Human Resources*, 57 (5), 1690–1714. [52](#), [102](#)
- Angrist, Joshua D, Kathryn Graddy, and Guido W Imbens (2000) “The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish,” *The Review of Economic Studies*, 67 (3), 499–527. [11](#)
- Arulampalam, Wiji, Alison L. Booth, and Mark L. Bryan (2004) “Training and the New Minimum Wage,” *The Economic Journal*, 114 (494), C87–C94. [46](#), [99](#)
- Ashenfelter, Orley and David Card (1981) “Using Longitudinal Data to Measure Minimum Wage Effects,” cep discussion paper, Centre for Economic Performance, LSE. [32](#)
- Ashenfelter, Orley and Štěpán Jurajda (2022) “Minimum Wages, Wages, and Price Pass-Through: The Case of McDonald’s Restaurants,” *Journal of Labor Economics*, 40 (S1), S179–S201. [101](#)
- Ashenfelter, Orley and Robert S Smith (1979) “Compliance with the minimum wage law,” *Journal of Political Economy*, 87 (2), 333–350. [44](#), [99](#)
- Autor, David, Arindrajit Dube, and Annie McGrew (2023) “The unexpected compression: Competition at work in the low wage labor market,” Technical report, National Bureau of Economic Research. [61](#)
- Autor, David H, Lawrence F Katz, and Melissa S Kearney (2008) “Trends in US wage inequality: Revising the revisionists,” *The Review of economics and statistics*, 90 (2), 300–323. [61](#)
- Autor, David H., Alan Manning, and Christopher L. Smith (2016) “The Contribution of the minimum wage to U.S. wage inequality over three decades: a reassessment,” *American Economic Journal: Applied Economics*, 8 (1), 58–99. [44](#), [62](#), [63](#)
- Azar, José, Emiliano Huet-Vaughn, Ioana Marinescu, Bledi Taska, and Till von Wachter (2023) “Minimum Wage Employment Effects and Labour Market Concentration,” *The Review of Economic Studies*, rdad091. [25](#), [38](#), [39](#), [58](#)
- Bailey, Martha J, John DiNardo, and Bryan A Stuart (2021) “The economic impact of a high national minimum wage: Evidence from the 1966 fair labor standards act,” *Journal of labor economics*, 39 (S2), S329–S367. [5](#), [24](#), [30](#), [37](#), [38](#)
- Baker, Michael, Dwayne Benjamin, and Shuchita Stanger (1999) “The highs and lows of the minimum wage effect: A time-series cross-section study of the Canadian law,” *Journal of labor Economics*, 17 (2), 318–350. [37](#)

- Basker, Emek and Muhammad Taimur Khan (2016) “Does the minimum wage bite into fast-food prices?” *Journal of Labor Research*, 37, 129–148. [101](#)
- Bassier, Ihsaan (2022) “Collective bargaining and spillovers in local labor markets.” [49](#)
- Bell, Brian and Stephen Machin (2018) “Minimum wages and firm value,” *Journal of Labor Economics*, 36 (1), 159–195. [53](#), [102](#)
- Bellmann, Lutz, Mario Bossler, Hans-Dieter Gerner, and Olaf Hübler (2017) “Training and minimum wages: first evidence from the introduction of the minimum wage in Germany,” *IZA journal of Labor Economics*, 6, 1–22. [46](#), [99](#)
- Belman, Dale and Paul J Wolfson (2014) “What Does the Minimum Wage Do?”. [33](#)
- Berger, David W, Kyle F Herkenhoff, and Simon Mongey (2022) “Minimum wages, efficiency and welfare,” Technical report, National Bureau of Economic Research. [58](#)
- Bernhardt, Annette, Michael Spiller, and Nik Theodore (2013) “Employers Gone Rogue: Explaining Industry Variation in Violations of Workplace Laws,” *Industrial and Labor Relations Review*, 66, 808. [44](#), [99](#)
- Bhaskar, V. and Ted To (1999) “Minimum Wages for Ronald McDonald Monopsonies: A Theory of Monopsonistic Competition,” *The Economic Journal*, 109 (455), 190–203. [58](#), [59](#)
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts (2013) “Vol. 128 February 2013 Issue,” *The Quarterly Journal of Economics*, 1, 51. [58](#)
- Bloom, Nicholas, Renata Lemos, Raffaella Sadun, Daniela Scur, and John Van Reenen (2014) “JEEA-FBBVA Lecture 2013: The new empirical economics of management,” *Journal of the European Economic Association*, 12 (4), 835–876. [58](#), [67](#)
- Blundell, Richard, Alan Duncan, and Costas Meghir (1998) “Estimating labor supply responses using tax reforms,” *Econometrica*, 827–861. [29](#)
- Böckerman, Petri and Roope Uusitalo (2009) “Minimum Wages and Youth Employment: Evidence from the Finnish Retail Trade Sector,” *British Journal of Industrial Relations*, 47 (2), 388–405. [48](#)
- Bodnár, Katalin, Ludmila Fadejeva, Stefania Iordache et al. (2018) “How do firms adjust to rises in the minimum wage? Survey evidence from Central and Eastern Europe,” *IZA Journal of Labor Policy*, 7, 1–30. [56](#), [57](#)
- Boeri, Tito, Giulia Giupponi, Alan B. Krueger, and Stephen Machin (2020) “Solo Self-Employment and Alternative Work Arrangements: A Cross-Country Perspective on the Changing Composition of Jobs,” *Journal of Economic Perspectives*, 34 (1), 170–95. [43](#)
- Boffy-Ramirez, Ernest (2013) “Minimum wages, earnings, and migration,” *IZA Journal of Migration*, 2, 1–24. [50](#), [100](#)
- Boone, Christopher, Arindrajit Dube, Lucas Goodman, and Ethan Kaplan (2021) “Unemployment insurance generosity and aggregate employment,” *American Economic Journal: Economic Policy*, 13 (2), 58–99. [23](#)

- Borgschulte, Mark and HeePyung Cho (2020) “Minimum wages and retirement,” *ILR Review*, 73 (1), 153–177. [36](#), [51](#), [100](#)
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024) “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *The Review of Economic Studies*. [16](#), [22](#)
- Bosch, Mariano and Marco Manacorda (2010) “Minimum wages and earnings inequality in urban Mexico,” *American Economic Journal: Applied Economics*, 2 (4), 128–149. [68](#)
- Bossler, Mario and Sandra Broszeit (2017) “Do minimum wages increase job satisfaction? Micro-data evidence from the new German minimum wage,” *LABOUR*, 31 (4), 480–493. [46](#)
- Bossler, Mario and Hans-Dieter Gerner (2020) “Employment Effects of the New German Minimum Wage: Evidence from Establishment-Level Microdata,” *ILR Review*, 73 (5), 1070–1094. [41](#)
- Bossler, Mario and Thorsten Schank (2023) “Wage Inequality in Germany after the Minimum Wage Introduction,” *Journal of Labor Economics*, 41 (3), 813–857. [61](#), [63](#)
- Brochu, Pierre and David A Green (2013) “The impact of minimum wages on labour market transitions,” *The Economic Journal*, 123 (573), 1203–1235. [58](#), [102](#)
- Brochu, Pierre, David A Green, Thomas Lemieux, and James Townsend (2023) “The minimum wage, turnover, and the shape of the wage distribution.” [102](#)
- Broecke, Stijn, Alessia Forti, and Marieke Vandeweyer (2017) “The effect of minimum wages on employment in emerging economies: a survey and meta-analysis,” *Oxford Development Studies*, 45 (3), 366–391. [67](#)
- Brown, Charles (1999) “Minimum wages, employment, and the distribution of income,” in Ashenfelter, O. and D. Card eds. *Handbook of Labor Economics*, 3, Part B, 1st edition, Chap. 32, 2101–2163: Elsevier. [2](#), [3](#), [4](#), [5](#), [7](#), [45](#), [57](#), [59](#)
- Brown, Charles, Curtis Gilroy, and Andrew Kohen (1982) “The Effect of the Minimum Wage on Employment and Unemployment.,” *Journal of Economic Literature*, 20 (2). [33](#)
- Brown, Jessica H and Chris M Herbst (2023) “Minimum Wage, Worker Quality, and Consumer Well-Being: Evidence from the Child Care Market.” [52](#), [55](#), [65](#)
- Brummund, Peter and Michael R Strain (2020) “Does Employment Respond Differently to Minimum Wage Increases in the Presence of Inflation Indexing?” *Journal of Human Resources*, 55 (3). [124](#)
- Buchanan, James (1996) “Commentary on the Minimum Wage,” *Wall Street Journal*, April 25, A20. [9](#)
- Buraue, Patrick, Marco Caliendo, Markus M. Grabka, Cosima Obst, Malte Preuss, and Carsten Schröder (2020) “The Impact of the Minimum Wage on Working Hours,” *Jahrbücher für Nationalökonomie und Statistik*, 240 (2-3), 233–267. [41](#)
- Burdett, Kenneth and Dale T Mortensen (1998) “Wage differentials, employer size, and unemployment,” *International Economic Review*, 257–273. [54](#)

- Burkhauser, Richard V, Drew McNichols, and Joseph J Sabia (2023) “Minimum wages and poverty: new evidence from dynamic difference-in-differences estimates,” Technical report, National Bureau of Economic Research. [64](#)
- Butcher, Tim, Richard Dickens, and Alan Manning (2012) “Minimum wages and wage inequality: some theory and an application to the UK.” [62](#)
- Cadena, Brian C. (2014) “Recent immigrants as labor market arbitrageurs: Evidence from the minimum wage,” *Journal of Urban Economics*, 80, 1–12. [50](#), [100](#)
- Caetano, Carolina, Brantly Callaway, Stroud Payne, and Hugo Sant’Anna Rodrigues (2022) “Difference in differences with time-varying covariates,” *arXiv preprint arXiv:2202.02903*. [22](#)
- Caliendo, Marco, Alexandra Fedorets, Malte Preuss, Carsten Schröder, and Linda Wittbrodt (2018) “The short-run employment effects of the German minimum wage reform,” *Labour Economics*, 53, 46–62, European Association of Labour Economists 29th annual conference, St.Gallen, Switzerland, 21-23 September 2017. [30](#)
- Caliendo, Marco, Linda Wittbrodt, and Carsten Schröder (2019) “The Causal Effects of the Minimum Wage Introduction in Germany – An Overview,” *German Economic Review*, 20 (3), 257–292. [44](#), [99](#)
- Callaway, Brantly and Pedro HC Sant’Anna (2021) “Difference-in-differences with multiple time periods,” *Journal of econometrics*, 225 (2), 200–230. [15](#), [16](#), [21](#), [22](#), [114](#), [115](#)
- Campolieti, Michele, Tony Fang, and Morley Gunderson (2005) “How minimum wages affect schooling-employment outcomes in Canada, 1993–1999,” *Journal of Labor Research*, 26 (3), 533–545. [66](#)
- Card, David (1992a) “Do Minimum Wages Reduce Employment? A Case Study of California, 1987–89,” *Industrial and Labor Relations Review*, 46 (1), 38–54. [8](#)
- (1992b) “Using regional variation in wages to measure the effects of the federal minimum wage,” *ILR Review*, 46 (1), 22–37. [8](#), [30](#)
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline (2018) “Firms and Labor Market Inequality: Evidence and Some Theory,” *Journal of Labor Economics*, 36 (S1), S13–S70. [7](#), [57](#)
- Card, David and John E DiNardo (2002) “Skill-biased technological change and rising wage inequality: Some problems and puzzles,” *Journal of labor economics*, 20 (4), 733–783. [61](#)
- Card, David, Jörg Heining, and Patrick Kline (2013) “Workplace heterogeneity and the rise of West German wage inequality,” *The Quarterly journal of economics*, 128 (3), 967–1015. [61](#)
- Card, David, Lawrence F. Katz, and Alan B. Krueger (1994) “Comment on David Neumark and William Wascher, ”Employment Effects of Minimum and Subminimum Wages: Panel Data on State Minimum Wage Laws”,” *Industrial and Labor Relations Review*, 47 (3), 487–497. [48](#)
- Card, David and Alan B. Krueger (1994a) “The effect of the minimum wage on shareholder wealth,” *Working paper, Princeton University, Department of Economics, Industrial Relations Section*. [102](#)

- (1994b) “Minimum Wages and Employment: A Case Study of the New Jersey and Pennsylvania Fast Food Industries,” *American Economic Review*, 84 (4), 772–793. [2](#), [8](#), [9](#), [12](#), [13](#), [15](#), [17](#), [22](#), [31](#)
- (1995a) *Myth and measurement: the new economics of the minimum wage*, New Jersey: Princeton University Press. [25](#), [101](#)
- Card, David and Alan B Krueger (1995b) “Time-series minimum-wage studies: a meta-analysis,” *The American Economic Review*, 85 (2), 238–243. [12](#), [45](#), [53](#)
- (2000) “Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania: reply,” *American Economic Review*, 90 (5), 1397–1420. [13](#)
- Castillo-Freeman, Alida J. and Richard B. Freeman (1992) *When the Minimum Wage Really Bites: The Effect of the U. S. -Level Minimum on Puerto Rico*, 177–212, Chicago: University of Chicago Press. [50](#), [100](#)
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and David Zentler-Munro (2022) “Seeing beyond the trees: Using machine learning to estimate the impact of minimum wages on labor market outcomes,” *Journal of Labor Economics*, 40 (S1), S203–S247. [15](#), [22](#), [26](#), [36](#), [37](#), [38](#), [41](#), [51](#), [100](#), [126](#)
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) “The effect of minimum wages on low-wage jobs,” *The Quarterly Journal of Economics*, 134 (3), 1405–1454. [14](#), [15](#), [21](#), [22](#), [25](#), [35](#), [36](#), [37](#), [38](#), [39](#), [41](#), [63](#), [93](#), [115](#)
- Chaplin, Duncan D, Mark D Turner, and Andreas D Pape (2003) “Minimum wages and school enrollment of teenagers: a look at the 1990’s,” *Economics of Education Review*, 22 (1), 11–21. [66](#)
- Chava, Sudheer, Alexander Oettl, and Manpreet Singh (2023) “Does a One-Size-Fits-All Minimum Wage Cause Financial Stress for Small Businesses?” *Management Science*, 69 (11), 7095–7117. [48](#), [100](#)
- Clemens, Jeffrey (2021) “How Do Firms Respond to Minimum Wage Increases? Understanding the Relevance of Non-employment Margins,” *Journal of Economic Perspectives*, 35 (1), 51–72. [45](#)
- Clemens, Jeffrey, Lisa B Kahn, and Jonathan Meer (2018) “The Minimum Wage, Fringe Benefits, and Worker Welfare,” Working Paper 24635, National Bureau of Economic Research. [45](#), [51](#), [99](#)
- Clemens, Jeffrey, Lisa B. Kahn, and Jonathan Meer (2021) “Dropouts Need Not Apply? The Minimum Wage and Skill Upgrading,” *Journal of Labor Economics*, 39 (S1), [10.1086/711490](#). [47](#), [55](#), [100](#)
- Clemens, Jeffrey and Michael R Strain (2021) “The heterogeneous effects of large and small minimum wage changes: evidence over the short and medium run using a pre-analysis plan,” Technical report, National Bureau of Economic Research. [21](#), [38](#)
- (2022) “Understanding “wage theft”: Evasion and avoidance responses to minimum wage increases,” *Labour Economics*, 79, 102285. [44](#), [99](#)

- Clemens, Jeffrey and Michael Wither (2019) “The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers,” *Journal of Public Economics*, 170, 53–67. 21, 26, 27, 30, 32
- Collier, Ruth Berins and David Collier (2002) *Shaping the political arena: critical junctures, the labor movement, and regime dynamics in Latin America*: University of Notre Dame Press. 5
- Congressional Budget Office, (CBO) (2019) “The Effects on Employment and Family Income of Increasing the Federal Minimum Wage,” <https://www.cbo.gov/system/files/2019-07/CBO-55410-MinimumWage2019.pdf>, Accessed: 2024-07-30. 35
- Connolly, Sara and Mary Gregory (2002) “The National Minimum Wage and Hours of Work: Implications for Low Paid Women,” *Oxford Bulletin of Economics and Statistics*, 64 (supplement), 607–631. 41
- Cook, Maria Lorena (2010) *Politics of labor reform in Latin America: Between flexibility and rights*: Penn State Press. 5
- Cooper, Daniel, Maria Jose Luengo-Prado, and Jonathan A. Parker (2020) “The Local Aggregate Effects of Minimum Wage Increases,” *Journal of Money, Credit and Banking*, 52 (1), 5–35, <https://doi.org/10.1111/jmcb.12684>. 52, 101, 102
- Couch, Kenneth A. and David C. Wittenburg (2001) “The Response of Hours of Work to Increases in the Minimum Wage,” *Southern Economic Journal*, 68 (1), 171–177. 40
- Coviello, Decio, Erika Deserranno, and Nicola Persico (2022) “Minimum wage and individual worker productivity: Evidence from a large US retailer,” *Journal of Political Economy*, 130 (9), 2315–2360. 23, 54, 55, 58, 59, 102
- Currie, Janet and Bruce C. Fallick (1996) “The Minimum Wage and the Employment of Youth Evidence from the NLSY,” *The Journal of Human Resources*, 31 (2), 404–428. 31, 32
- Datta, Nikhil and Stephen Machin (2024) “Government contracting and living wages; minimum wages,” Technical report, Centre for Economic Performance, LSE. 60, 61
- Dautović, Ernest, Harald Hau, and Yi Huang (2024) “Consumption Response to Minimum Wages: Evidence from Chinese Households,” *The Review of Economics and Statistics*, 1–47. 65
- De Chaisemartin, Clément and Xavier d’Haultfoeuille (2018) “Fuzzy differences-in-differences,” *The Review of Economic Studies*, 85 (2), 999–1028. 29
- (2020) “Two-way fixed effects estimators with heterogeneous treatment effects,” *American economic review*, 110 (9), 2964–2996. 15, 16
- DellaVigna, Stefano and Matthew Gentzkow (2019) “Uniform pricing in us retail chains,” *The Quarterly Journal of Economics*, 134 (4), 2011–2084. 58
- Demir, Banu, Ana Cecilia Fieler, Daniel Yi Xu, and Kelly Kaili Yang (2024) “O-Ring Production Networks,” *Journal of Political Economy*, 132 (1), 200–247. 53

- Derenoncourt, Ellora, François Gérard, Lorenzo Lagos, and Claire Montialoux (2021) “Racial inequality, minimum wage spillovers, and the informal sector.” [38](#), [68](#)
- Derenoncourt, Ellora and Claire Montialoux (2021) “Minimum wages and racial inequality,” *The Quarterly Journal of Economics*, 136 (1), 169–228. [5](#), [24](#), [25](#), [27](#), [35](#), [37](#), [38](#)
- Dettling, Lisa J and Joanne W Hsu (2020) “Minimum Wages and Consumer Credit: Effects on Access and Borrowing,” *The Review of Financial Studies*, 34 (5), 2549–2579. [65](#)
- Dickens, Richard, Stephen Machin, and Alan Manning (1998) “Estimating the effect of minimum wages on employment from the distribution of wages: A critical view,” *Labour Economics*, 5 (2), 109 – 134. [25](#)
- (1999) “The effects of minimum wages on employment: Theory and evidence from Britain,” *Journal of labor economics*, 17 (1), 1–22. [30](#), [32](#), [58](#)
- DiNardo, John, Nicole M Fortin, and Thomas Lemieux (1996) “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 64 (5), 1001–1044. [61](#), [62](#)
- Dolado, Juan, Francis Kramarz, Stephen Machin, Alan Manning, David Margolis, and Coen Teulings (1996) “Minimum wages: the European experience,” *Economic Policy*, 23 (3), 317–357. [5](#)
- Dolton, Peter, Chiara Rosazza Bondibene, and Jonathan Wadsworth (2012) “Employment, Inequality and the UK National Minimum Wage over the Medium-Term*,” *Oxford Bulletin of Economics and Statistics*, 74 (1), 78–106. [30](#)
- Doucouliafos, Hristos and Tom D Stanley (2009) “Publication selection bias in minimum-wage research? A meta-regression analysis,” *British Journal of Industrial Relations*, 47 (2), 406–428. [9](#), [33](#)
- Downey, Mitch (2021) “Partial automation and the technology-enabled deskilling of routine jobs,” *Labour Economics*, 69, 101973. [47](#), [100](#)
- Draca, Mirko, Stephen Machin, and John Van Reenen (2011) “Minimum Wages and Firm Profitability,” *American Economic Journal: Applied Economics*, 3 (1), 129–51. [31](#), [48](#), [49](#), [54](#), [100](#), [102](#)
- Drucker, Lev, Katya Mazirov, and David Neumark (2021) “Who pays for and who benefits from minimum wage increases? Evidence from Israeli tax data on business owners and workers,” *Journal of Public Economics*, 199, 104423. [54](#), [102](#)
- Dube, Arindrajit (2019a) “Impacts of minimum wages: review of the international evidence,” *Independent Report. UK Government Publication*, 268–304. [6](#), [30](#), [33](#), [38](#)
- (2019b) “Minimum wages and the distribution of family incomes,” *American Economic Journal: Applied Economics*, 11 (4), 268–304. [64](#)

- Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M Taylor (2023) “A Local Projections Approach to Difference-in-Differences,” Working Paper 31184, National Bureau of Economic Research. [16](#), [21](#), [22](#), [114](#)
- Dube, Arindrajit, Laura Giuliano, and Jonathan Leonard (2019) “Fairness and Frictions: The Impact of Unequal Raises on Quit Behavior,” *American Economic Review*, 109 (2), 620–63. [51](#), [61](#), [102](#)
- Dube, Arindrajit, T. William Lester, and Michael Reich (2010) “Minimum wage effects across state borders: estimates using contiguous counties,” *The Review of Economics and Statistics*, 92 (4), 945–964. [15](#), [22](#), [38](#)
- Dube, Arindrajit, T William Lester, and Michael Reich (2016) “Minimum wage shocks, employment flows, and labor market frictions,” *Journal of Labor Economics*, 34 (3), 663–704. [23](#), [58](#), [102](#)
- Dube, Arindrajit and Attila Lindner (2021) “City limits: What do local-area minimum wages do?” *Journal of Economic Perspectives*, 35 (1), 27–50. [13](#), [21](#), [25](#)
- Dube, Arindrajit, Alan Manning, and Suresh Naidu (2020) “Monopsony and Employer Mis-optimization Explain Why Wages Bunch at Round Numbers.” [59](#), [62](#)
- Dube, Arindrajit, Suresh Naidu, and Adam D Reich (2022) “Power and dignity in the low-wage labor market: Theory and evidence from wal-mart workers,” Technical report, National Bureau of Economic Research. [45](#)
- Dube, Arindrajit, Suresh Naidu, and Michael Reich (2007) “The Economic Effects of a Citywide Minimum Wage,” *ILR Review*, 60 (4), 522–543. [41](#), [101](#), [102](#)
- Dube, Arindrajit, Michael Reich, Akash Bhatt, and Denis Sosinskiy (2024) “Restaurant Employment, Minimum Wages, and Border Discontinuities,” Working Paper 32902, National Bureau of Economic Research, [10.3386/w32902](#). [22](#)
- Dube, Arindrajit and Ben Zipperer (2024) “Own-Wage Elasticity: Quantifying the Impact of Minimum Wages on Employment,” Working Paper 32925, National Bureau of Economic Research, [10.3386/w32925](#). [33](#), [34](#), [35](#), [94](#), [95](#), [110](#)
- Dube, Eric Freeman, Arindrajit and Michael Reich (2010) “Employee Replacement Costs,” *Institute for Research on Labor and Employment, UC Berkeley Working Paper* (201–10). [55](#)
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge (2022) “Reallocation effects of the minimum wage,” *The Quarterly Journal of Economics*, 137 (1), 267–328. [31](#), [32](#), [48](#), [49](#), [56](#), [100](#), [102](#)
- Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg (2009) “Revisiting the German wage structure,” *The Quarterly journal of economics*, 124 (2), 843–881. [61](#)
- Emanuel, Natalia and Emma Harrington (2020) “The payoffs of higher pay: elasticities of productivity and labor supply with respect to wages.” [58](#)
- Engbom, Niklas and Christian Moser (2022) “Earnings inequality and the minimum wage: Evidence from Brazil,” *American Economic Review*, 112 (12), 3803–3847. [58](#), [62](#), [68](#)

- Fan, Haichao, Yichuan Hu, and Lixin Tang (2021) “Labor costs and the adoption of robots in China,” *Journal of Economic Behavior & Organization*, 186, 608–631. [47](#), [100](#)
- Ferman, Bruno and Cristine Pinto (2019) “Inference in differences-in-differences with few treated groups and heteroskedasticity,” *Review of Economics and Statistics*, 101 (3), 452–467. [16](#)
- Fishback, Price V and Andrew J Seltzer (2021) “The rise of American minimum wages, 1912–1968,” *Journal of Economic Perspectives*, 35 (1), 73–96. [4](#)
- Flinn, Christopher J (2006) “Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates,” *Econometrica*, 74 (4), 1013–1062. [58](#), [62](#)
- (2010) *The Minimum Wage and Labor Market Outcomes*: MIT Press. [58](#), [61](#)
- Fortin, Nicole M., Thomas Lemieux, and Neil Lloyd (2021) “Labor Market Institutions and the Distribution of Wages: The Role of Spillover Effects,” *Journal of Labor Economics*, 39 (S2), S369–S412. [62](#)
- Freeman, Richard (1989) *Labor Markets in Action*: Cambridge: Harvard University Press. [7](#)
- Freeman, Richard B (1990) “Employment and Earnings of Disadvantaged Young Men in a Labor Shortage Economy,” Working Paper 3444, National Bureau of Economic Research. [18](#)
- Friedman, Jerome H (2001) “Greedy function approximation: a gradient boosting machine,” *Annals of statistics*, 1189–1232. [126](#)
- Friedman, Milton (1953) *The Methodology of Positive Economics. In Essays on Positive Economics.*: Chicago: University of Chicago Press. [8](#)
- Gardner, John, Neil Thakral, Linh T. Tô, and Yap Luther (2024) “Two-stage differences in differences.” [16](#)
- Geng, Heng (Griffin), Yi Huang, Chen Lin, and Sibio Liu (2022) “Minimum Wage and Corporate Investment: Evidence from Manufacturing Firms in China,” *Journal of Financial and Quantitative Analysis*, 57 (1), 94–126. [47](#), [100](#)
- Gibbons, Robert and John Roberts (2015) “Organizational economics,” *Emerging Trends in the Social and Behavioral Sciences. American Cancer Society*, 1–15. [59](#)
- Gittings, R Kaj and Ian M Schmutte (2016) “Getting handcuffs on an octopus: Minimum wages, employment, and turnover,” *ILR Review*, 69 (5), 1133–1170. [102](#)
- Giuliano, Laura (2013) “Minimum Wage Effects on Employment, Substitution, and the Teenage Labor Supply: Evidence from Personnel Data,” *Journal of Labor Economics*, 31 (1), 155–194. [36](#)
- Giulietti, Corrado (2014) “Is the minimum wage a pull factor for immigrants?” *ILR Review*, 67 (3.suppl), 649–674. [50](#), [100](#)
- Giupponi, Giulia, Robert Joyce, Attila Lindner, Tom Waters, Thomas Wernham, and Xiaowei Xu (2024) “The employment and distributional impacts of nationwide minimum wage changes,” *Journal of Labor Economics*, 42 (S1), S293–S333. [30](#), [32](#), [36](#), [39](#), [50](#), [63](#), [64](#), [100](#)

- Giupponi, Giulia and Stephen Machin (2022a) “Labour market inequality,” *IFS Deaton Review*, 15. [61](#)
- (2022b) “Company Wage Policy in a Low-Wage Labor Market,” CEP Discussion Paper CEPDP1869, Centre for Economic Performance, LSE. [48](#)
- Godoe, Anna and Michael Reich (2021) “Are Minimum Wage Effects Greater in Low-Wage Areas?” *Industrial Relations: A Journal of Economy and Society*, 60 (1), 36–83. [24](#), [38](#), [41](#), [64](#)
- Godøy, Anna and Ken Jacobs (2021) “The downstream benefits of higher incomes and wages,” *Federal Reserve Bank of Boston Community Development Discussion Paper* (21-1). [66](#)
- Godøy, Anna, Michael Reich, Jesse Wursten, and Sylvia Allegretto (2024) “Parental labor supply: Evidence from minimum wage changes,” *Journal of Human Resources*, 59 (2), 416–442. [38](#), [51](#), [100](#)
- Goodman-Bacon, Andrew (2021) “Difference-in-differences with variation in treatment timing,” *Journal of econometrics*, 225 (2), 254–277. [15](#), [16](#)
- Gopalan, Radhakrishnan, Barton H Hamilton, Ankit Kalda, and David Sovich (2021) “State minimum wages, employment, and wage spillovers: Evidence from administrative payroll data,” *Journal of Labor Economics*, 39 (3), 673–707. [37](#), [41](#), [51](#), [61](#), [63](#), [102](#)
- Goraus-Tańska, Karolina and Piotr Lewandowski (2019) “Minimum wage violation in central and eastern Europe,” *International Labour Review*, 158 (2), 297–336. [44](#), [99](#)
- Grimshaw, Damian and Marcela Miozzo (2003) *Minimum wages and pay equity in Latin America*, 12: International Labour Office. [5](#)
- Grossberg, Adam J and Paul Sicilian (1999) “Minimum wages, on-the-job training, and wage growth,” *Southern Economic Journal*, 65 (3), 539–556. [46](#), [99](#)
- Grossman, Jonathan (1978) “Fair Labor Standards Act of 1938: Maximum struggle for a minimum wage,” *Monthly Labor Review*, 101 (6), 22–30. [4](#)
- Gülal, Filiz and Adam Ayaita (2020) “The impact of minimum wages on well-being: evidence from a quasi-experiment in Germany,” *Journal of happiness studies*, 21, 2669–2692. [46](#)
- Haanwinckel, Daniel (2023) “Supply, demand, institutions, and firms: A theory of labor market sorting and the wage distribution,” Technical report, National Bureau of Economic Research. [58](#), [68](#)
- Hamermesh, Daniel S. (1995) “Labour Demand and the Source of Adjustment Costs,” *The Economic Journal*, 105 (430), 620–634. [7](#), [8](#)
- Hammond, Matthew B (1915) “Wages Boards in Australia: IV. Social and Economic Results of Wages Boards,” *The Quarterly Journal of Economics*, 563–630. [4](#)
- Hampton, Matt and Evan Totty (2023) “Minimum wages, retirement timing, and labor supply,” *Journal of Public Economics*, 224, 104924. [26](#), [36](#), [41](#), [51](#), [100](#)
- Hara, Hiromi (2017) “Minimum wage effects on firm-provided and worker-initiated training,” *Labour Economics*, 47, 149–162, EALE conference issue 2016. [46](#), [99](#)

- Harasztosi, Péter and Attila Lindner (2019) “Who pays for the minimum wage?” *American Economic Review*, 109 (8), 2693–2727. [31](#), [33](#), [37](#), [38](#), [39](#), [45](#), [47](#), [48](#), [49](#), [52](#), [53](#), [54](#), [56](#), [65](#), [99](#), [100](#), [101](#), [102](#)
- Hau, Harald, Yi Huang, and Gewei Wang (2020) “Firm response to competitive shocks: Evidence from China’s minimum wage policy,” *The Review of Economic Studies*, 87 (6), 2639–2671. [47](#), [55](#), [100](#), [102](#)
- Hazell, Jonathon, Christina Patterson, Heather Sarsons, and Bledi Taska (2022) “National wage setting,” *University of Chicago, Becker Friedman Institute for Economics Working Paper* (2022-150). [58](#)
- Hicks, John (1932) *The Theory of Wages*: London: Macmillan. [6](#)
- Hirsch, Barry T, Bruce E Kaufman, and Tetyana Zelenska (2015) “Minimum wage channels of adjustment,” *Industrial Relations: A Journal of Economy and Society*, 54 (2), 199–239. [51](#), [56](#), [57](#), [102](#)
- Holzer, Harry J., Lawrence F. Katz, and Alan B. Krueger (1991) “Job Queues and Wages,” *The Quarterly Journal of Economics*, 106 (3), 739–768. [46](#)
- Horton, John J (Forthcoming) “Price floors and employer preferences: Evidence from a minimum wage experiment,” *American Economic Review*. [47](#), [100](#)
- Hsieh, Chang-Tai and Peter J Klenow (2009) “Misallocation and manufacturing TFP in China and India,” *The Quarterly journal of economics*, 124 (4), 1403–1448. [67](#)
- Huang, Yi, Mr Prakash Loungani, and Gewei Wang (2014) *Minimum wages and firm employment: Evidence from China*: International Monetary Fund. [68](#)
- ILO (2017) “Minimum wage policy guide.” [5](#)
- ILO, OECD, IMF, and World Bank (2012) “Boosting jobs and living standards in G20 countries.” [2](#)
- Jacobs, Ken, Michael Reich, Tynan Challenor, and Aida Farmand (2024) “Gig Passenger and Delivery Driver Pay in Five Metro Areas.” [43](#)
- Jardim, Ekaterina, Mark C Long, Robert Plotnick, Emma van Inwegen, Jacob Vigdor, and Hilary Wething (2017) “Minimum wage increases, wages, and low-wage employment: Evidence from Seattle,” NBER Working Paper No. 23532. [23](#)
- Jardim, Ekaterina, Mark C. Long, Robert Plotnick, Emma van Inwegen, Jacob Vigdor, and Hilary Wething (2022) “Minimum-Wage Increases and Low-Wage Employment: Evidence from Seattle,” *American Economic Journal: Economic Policy*, 14 (2), 263–314. [13](#), [25](#), [26](#), [41](#), [102](#)
- Jha, Priyaranjan, David Neumark, and Antonio Rodriguez-Lopez (2024) “What’s across the border? Re-evaluating the cross-border evidence on minimum wage effects,” *Journal of Political Economy Microeconomics*, Forthcoming. [21](#), [22](#), [23](#), [38](#)
- Kabátek, Jan (2021) “Happy birthday, You’re Fired! The Effects of an Age-Dependent Minimum Wage on Youth Employment Flows in the Netherlands,” *ILR Review*, 74 (4), 1008–1035. [32](#), [48](#)

- Kaitz, Hyman (1970) “Experience of the past: the national minimum,” *Youth unemployment and minimum wages*, 30–54. [6](#)
- Karabarbounis, Loukas, Jeremy Lise, and Anusha Nath (2022) “Minimum wages and labor markets in the Twin Cities,” Technical report, National Bureau of Economic Research. [13](#)
- Katz, Lawrence F and Alan B Krueger (1992) “The effect of the minimum wage on the fast-food industry,” *ILR Review*, 46 (1), 6–21. [48](#), [101](#)
- Katz, Lawrence F and Kevin M Murphy (1992) “Changes in relative wages, 1963–1987: supply and demand factors,” *The Quarterly Journal of Economics*, 107 (1), 35–78. [8](#)
- Kellogg, Maxwell, Magne Mogstad, Guillaume A Pouliot, and Alexander Torgovitsky (2021) “Combining matching and synthetic control to tradeoff biases from extrapolation and interpolation,” *Journal of the American Statistical Association*, 116 (536), 1804–1816. [17](#)
- Kennan, John (1995) “The elusive effects of minimum wages,” *Journal of Economic Literature*, 33 (4), 1950–1965. [60](#)
- Kertesi, Gábor and János Köllő (2003) “Fighting “Low Equilibria” by Doubling the Minimum Wage? Hungary’s Experiment.” [30](#)
- Khurana, Saloni and Kanika Mahajan (2020) “Evolution of wage inequality in India (1983-2017): The role of occupational task content,” *WIDER working paper*. [90](#), [119](#)
- Khurana, Saloni, Kanika Mahajan, and Kunal Sen (2023a) “Minimum Wages and Changing Wage Inequality in India,” Technical report, Institute of Labor Economics (IZA). [68](#)
- (2023b) “Minimum wages and changing wage inequality in India,” *IZA discussion papers*. [90](#), [119](#)
- Kreiner, Claus Thustrup, Daniel Reck, and Peer Ebbesen Skov (2020) “Do Lower Minimum Wages for Young Workers Raise Their Employment? Evidence from a Danish Discontinuity,” *The Review of Economics and Statistics*, 102 (2), 339–354. [32](#), [48](#)
- Krueger, Alan B (1994) “The effect of the minimum wage when it really bites: A reexamination of the evidence from Puerto Rico,” Working Paper 4757, National Bureau of Economic Research. [28](#), [50](#), [100](#)
- Ku, Hyejin (2022) “Does minimum wage increase labor productivity? Evidence from piece rate workers,” *Journal of Labor Economics*, 40 (2), 325–359. [55](#), [102](#)
- Kudlyak, Marianna, Murat Tasci, and Didem Tuzemen (2023) “Minimum Wage Increases and Vacancies, Working Paper 2022-10.” [55](#)
- Lavecchia, Adam M. (2020) “Minimum wage policy with optimal taxes and unemployment,” *Journal of Public Economics*, 190, 104228. [51](#), [100](#)
- Laws, Athene (2018) “Do minimum wages increase search effort?” Technical report, Faculty of Economics, University of Cambridge. [51](#), [100](#)

- Lee, Chang Hyung (2020) “Minimum wage policy and community college enrollment patterns,” *ILR review*, 73 (1), 178–210. [66](#)
- Lee, David S (1999) “Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage?” *The Quarterly Journal of Economics*, 114 (3), 977–1023. [62](#), [63](#)
- Lee, David and Emmanuel Saez (2012a) “Optimal minimum wage policy in competitive labor markets,” *Journal of Public Economics*, 96 (9-10), 739–749. [10](#)
- (2012b) “Optimal minimum wage policy in competitive labor markets,” *Journal of Public Economics*, 96 (9-10), 739–749. [40](#)
- Legal, Diego and Eric R. Young (2024) “The effect of minimum wages on consumer bankruptcy,” *Journal of Economics and Business*, 129, 106171, Inequality in Consumer Credit and Payments. [65](#)
- Leibenstein, Harvey (1966) “Allocative Efficiency vs. ”X-Efficiency”,” *The American Economic Review*, 56 (3), 392–415. [57](#)
- Leigh, J Paul, Wesley A Leigh, and Juan Du (2019) “Minimum wages and public health: a literature review,” *Preventive medicine*, 118, 122–134. [66](#)
- Lemos, Sara (2006) “Anticipated effects of the minimum wage on prices,” *Applied Economics*, 38 (3), 325–337. [68](#), [101](#)
- Lemos, Sara, Roberto Rigobon, and Kevin Lang (2004) “Minimum wage policy and employment effects: Evidence from brazil [with comments],” *Economia*, 5 (1), 219–266. [68](#)
- Leonard, Thomas (2000) “The Very Idea of Applying Economics: The Modern Minimum-Wage Controversy and Its Antecedents,” *History of Political Economy*, 32 (5), 117–144. [9](#)
- Lester, Richard A. (1946) “Shortcomings of Marginal Analysis for Wage-Employment Problems,” *The American Economic Review*, 36 (1), 63–82. [2](#), [7](#), [8](#), [56](#)
- Leung, Justin H. (2021) “Minimum Wage and Real Wage Inequality: Evidence from Pass-Through to Retail Prices,” *The Review of Economics and Statistics*, 103 (4), 754–769. [52](#), [101](#), [102](#)
- Levin-Waldman, Oren M (2000) “The effects of the minimum wage: A business response,” *Journal of Economic Issues*, 34 (3), 723–730. [56](#)
- Liu, Shanshan, Thomas J Hyclak, and Krishna Regmi (2016) “Impact of the minimum wage on youth labor markets,” *Labour*, 30 (1), 18–37. [102](#)
- Lordan, Grace and David Neumark (2018) “People versus machines: The impact of minimum wages on automatable jobs,” *Labour Economics*, 52, 40–53. [47](#), [100](#)
- Luca, Dara Lee and Michael Luca (2019) “Survival of the Fittest: The Impact of the Minimum Wage on Firm Exit,” Working Paper 25806, National Bureau of Economic Research. [48](#), [49](#), [100](#)
- Luna-Alpizar, Jose Luis (2019) “Worker Heterogeneity and the Asymmetric Effects of Minimum Wages,” *CERGE-EI Working Papers*. [51](#), [100](#)

- Machin, Stephen (2024) “Wage controversies: real wage stagnation, inequality and labour market institutions,” *LSE Public Policy Review*, 3 (2). [61](#)
- Machin, Stephen, Alan Manning, and Lupin Rahman (2003) “Where the minimum wage bites hard: Introduction of minimum wages to a low wage sector,” *Journal of the European Economic Association*, 1 (1), 154–180. [31](#)
- Machlup, Fritz (1946) “Marginal Analysis and Empirical Research,” *The American Economic Review*, 36 (4), 519–554. [8](#)
- MaCurdy, Thomas (2015) “How effective is the minimum wage at supporting the poor?” *Journal of Political Economy*, 123 (2), 497–545. [59](#), [65](#)
- Malan, T (1978) “Wage control and minimum wages in Africa,” *Africa Insight*, 8 (1), 3–17. [5](#)
- Manning, Alan (2003) *Monopsony in Motion: Imperfect Competition in Labor Markets*: Princeton University Press. [54](#), [58](#)
- (2021) “The elusive employment effect of the minimum wage,” *Journal of Economic Perspectives*, 35 (1), 3–26. [2](#), [14](#), [21](#), [29](#), [30](#), [36](#), [38](#)
- Mansoor, Kashif and Donal O’Neill (2021) “Minimum wage compliance and household welfare: An analysis of over 1500 minimum wages in India,” *World Development*, 147, 105653. [65](#)
- Marcus, Michelle and Pedro HC Sant’Anna (2021) “The role of parallel trends in event study settings: An application to environmental economics,” *Journal of the Association of Environmental and Resource Economists*, 8 (2), 235–275. [114](#)
- Marks, Mindy S (2011) “Minimum Wages, Employer-Provided Health Insurance, and the Non-discrimination Law,” *Industrial Relations: A Journal of Economy and Society*, 50 (2), 241–262. [45](#), [99](#)
- Marshall, Alfred (1897) “The Old Generation of Economists and the New,” *The Quarterly Journal of Economics*, 11 (2), 115–135. [6](#), [7](#)
- Mayneris, Florian, Sandra Poncet, and Tao Zhang (2018) “Improving or disappearing: Firm-level adjustments to minimum wages in China,” *Journal of Development Economics*, 135, 20–42. [48](#), [54](#), [55](#), [68](#), [100](#), [102](#)
- McGuinness, Seamus and Paul Redmond (2018) “Estimating the Effect of an Increase in the Minimum Wage on Hours Worked and Employment in Ireland,” IZA Discussion Papers 11632, Institute of Labor Economics (IZA). [41](#)
- Meer, Jonathan and Jeremy West (2016) “Effects of the minimum wage on employment dynamics,” *Journal of Human Resources*, 51 (2), 500–522. [13](#), [18](#)
- Meiselbach, Mark K and Jean M Abraham (2023) “Do minimum wage laws affect employer-sponsored insurance provision?” *Journal of health economics*, 92, 102825. [45](#), [99](#)
- Metcalf, David (1999) “The Low Pay Commission and the national minimum wage,” *The Economic Journal*, 109 (453), 46–66. [4](#)

- Meyer, Brett (2016) “Learning to love the government: Trade unions and late adoption of the minimum wage,” *World Politics*, 68 (3), 538–575. [5](#)
- Meyer, Robert H and David A Wise (1983) “The Effects of the Minimum Wage on the Employment and Earnings of Youth,” *Journal of Labor Economics*, 1 (1), 66–100. [25](#)
- Mill, John Stuart (1848) *Principles of Political Economy*: John W. Parker. [6](#)
- Ministry of Manpower in Singapore (2018) “Opinion Editorial by Minister Josephine Teo on Minimum Wage.” [5](#)
- Minton, Robert and Brian Wheaton (2023) “Minimum Wages and Internal Migration.” [50](#), [100](#)
- Monras, Joan (2019) “Minimum wages and spatial equilibrium: Theory and evidence,” *Journal of Labor Economics*, 37 (3), 853–904. [35](#), [38](#), [50](#), [100](#)
- Neumark, David (2024) “The effects of minimum wages on (almost) everything? A review of recent evidence on health and related behaviors,” *LABOUR*, 38 (1), 1–65. [66](#)
- Neumark, David and Luis Felipe Munguía Corella (2021) “Do minimum wages reduce employment in developing countries? A survey and exploration of conflicting evidence,” *World Development*, 137, 105165. [67](#)
- Neumark, David, JM Ian Salas, and William Wascher (2014) “Revisiting the minimum wage—Employment debate: Throwing out the baby with the bathwater?” *ILR Review*, 67 (3-suppl), 608–648. [13](#), [14](#), [15](#), [21](#), [23](#), [38](#)
- Neumark, David, Mark Schweitzer, and William Wascher (2004) “Minimum wage effects throughout the wage distribution,” *Journal of Human Resources*, 39 (2), 425–450. [35](#), [40](#)
- (2005) “The effects of minimum wages on the distribution of family incomes: A nonparametric analysis,” *Journal of Human Resources*, 40 (4), 867–894. [64](#)
- Neumark, David and Peter Shirley (2022) “Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the United States?” *Industrial Relations: A Journal of Economy and Society*, 61 (4), 384–417. [9](#), [33](#)
- Neumark, David and William Wascher (1992) “Employment effects of minimum and subminimum wages: panel data on state minimum wage laws,” *ILR Review*, 46 (1), 55–81. [8](#), [13](#), [14](#)
- (1994) “Employment Effects of Minimum and Subminimum Wages: Reply to Card, Katz, and Krueger,” *Industrial and Labor Relations Review*, 47 (3), 497–512. [48](#)
- (1995a) “Minimum wage effects on employment and school enrollment,” *Journal of Business & Economic Statistics*, 13 (2), 199–206. [66](#)
- (1995b) “Minimum-wage effects on school and work transitions of teenagers,” *The American Economic Review*, 85 (2), 244–249. [66](#)
- (2000) “Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania: Comment,” *American Economic Review*, 90 (5), 1362–1396. [13](#)

- (2001a) “Minimum Wages and Training Revisited,” *Journal of Labor Economics*, 19 (3), 563–595. [99](#)
- (2001b) “Using the EITC to help poor families: New evidence and a comparison with the minimum wage,” *National Tax Journal*, 54 (2), 281–317. [46](#)
- (2003) “Minimum wages and skill acquisition: Another look at schooling effects,” *Economics of Education Review*, 22 (1), 1–10. [66](#)
- (2008) *Minimum wages*, Cambridge, MA: MIT Press. [33](#)
- (2011) “Does a Higher Minimum Wage Enhance the Effectiveness of the Earned Income Tax Credit?” *ILR Review*, 64 (4), 712–746. [40](#)
- Newell, Pamela (2009) “A Historical Overview of the Fair Labor Standards Act,” *Florida Coastal Law Review*, 10 (4), 675. [4](#)
- Nordlund, Willis J (1997) *The quest for a living wage: The history of the federal minimum wage program* (48): Greenwood publishing group. [4](#)
- Okudaira, Hiroko, Miho Takizawa, and Kenta Yamanouchi (2019) “Minimum wage effects across heterogeneous markets,” *Labour Economics*, 59, 110–122. [40](#)
- Orazem, Peter E. and J. Peter Mattila (2002) “Minimum wage effects on hours, employment, and number of firms: The Iowa case,” *Journal of labor research.*, 23 (1). [49](#), [100](#)
- Orrenius, Pia M and Madeline Zavodny (2008) “The effect of minimum wages on immigrants’ employment and earnings,” *ILR Review*, 61 (4), 544–563. [13](#), [45](#), [50](#), [100](#)
- Osterman, Paul (2011) “Institutional labor economics, the new personnel economics, and internal labor markets: A reconsideration,” *ILR Review*, 64 (4), 637–653. [8](#)
- Pereira, Sonia C (2003) “The Impact of Minimum Wages on Youth Employment in Portugal,” *European Economic Review*, 47 (2), 229–244. [48](#)
- Phelan, Brian J (2019) “Hedonic-based labor supply substitution and the ripple effect of minimum wages,” *Journal of Labor Economics*, 37 (3), 905–947. [62](#)
- Piqueras, Jon (2023) “Search Effort and the Minimum Wage.” [51](#), [100](#)
- Portugal, Pedro and Ana Rute Cardoso (2006) “Disentangling the Minimum Wage Puzzle: An Analysis of Worker Accessions and Separations,” *Journal of the European Economic Association*, 4 (5), 988–1013. [48](#), [102](#)
- Rao, Nirupama and Max Risch (2024) “Who’s Afraid of the Minimum Wage? Measuring the Impacts on Independent Businesses Using Matched US Tax Returns.” [31](#), [38](#), [48](#), [54](#), [55](#), [56](#), [102](#)
- Rebitzer, James B and Lowell J Taylor (1995) “The consequences of minimum wage laws some new theoretical ideas,” *Journal of Public Economics*, 56 (2), 245–255. [58](#)


- Reich, Michael and Amy Laitenen (2003) “Raising low pay in a high income economy: The economics of a San Francisco minimum wage.” [56](#)
- Reich, Michael and Rachel West (2015) “The effects of minimum wages on food stamp enrollment and expenditures,” *Industrial Relations: A Journal of Economy and Society*, 54 (4), 668–694. [64](#)
- Renkin, Tobias, Claire Montialoux, and Michael Siegenthaler (2022) “The Pass-Through of Minimum Wages into U.S. Retail Prices: Evidence from Supermarket Scanner Data,” *The Review of Economics and Statistics*, 104 (5), 890–908. [52](#), [65](#), [101](#), [102](#)
- Riley, Rebecca and Chiara Rosazza Bondibene (2017) “Raising the standard: Minimum wages and firm productivity,” *Labour Economics*, 44, 27–50. [55](#)
- Rohlin, Shawn M. (2011) “State minimum wages and business location: Evidence from a refined border approach,” *Journal of Urban Economics*, 69 (1), 103–117. [48](#), [49](#), [100](#)
- Rositani, Annunziata (2017) “Work and Wages in the Code of Hammurabi,” *Work and Wages in the Code of Hammurabi*, 47–71. [4](#)
- Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski, and John Poe (2023) “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature,” *Journal of Econometrics*, 235 (2), 2218–2244. [22](#), [114](#)
- Rothstein, Jesse (2010) “Is the EITC as good as an NIT? Conditional cash transfers and tax incidence,” *American economic Journal: economic policy*, 2 (1), 177–208. [40](#)
- Ruffini, Krista (2022) “Worker Earnings, Service Quality, and Firm Profitability: Evidence from Nursing Homes and Minimum Wage Reforms,” *The Review of Economics and Statistics*, 1–46. [38](#), [52](#), [55](#), [102](#)
- Sabia, Joseph J. (2009) “The Effects of Minimum Wage Increases on Retail Employment and Hours: New Evidence from Monthly CPS Data,” *Journal of Labor Research*, 30, 75–97. [40](#)
- Sabia, Joseph J and Robert B Nielsen (2015) “Minimum wages, poverty, and material hardship: new evidence from the SIPP,” *Review of Economics of the Household*, 13, 95–134. [64](#)
- Schanzenbach, Diane Whitmore, Julia A Turner, and Sarah Turner (2024) “Raising state minimum wages, lowering community college enrollment,” *Review of Economics and Statistics*, 1–29. [66](#)
- Shiller, Robert J (1994) *Macro markets: creating institutions for managing society’s largest economic risks*: OUP Oxford. [46](#)
- Sidgwick, Henry (1886) “Economic Socialism,” *History of Economic Thought Articles*, 50, 620–631. [7](#)
- Simon, Kosali Ilayperuma and Robert Kaestner (2004) “Do minimum wages affect non-wage job attributes? Evidence on fringe benefits,” *ILR Review*, 58 (1), 52–70. [45](#), [99](#)
- Smith, Alexander A (2021) “The minimum wage and teen educational attainment,” *Labour Economics*, 73, 102061. [66](#)

- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter (2018) “Firming Up Inequality,” *The Quarterly Journal of Economics*, 134 (1), 1–50. [61](#)
- Soundararajan, Vidhya (2019) “Heterogeneous effects of imperfectly enforced minimum wages in low-wage labor markets,” *Journal of Development Economics*, 140, 355–374. [23](#), [40](#)
- Stansbury, Anna (2024) “Incentives to Comply with the Minimum Wage in the US and UK,” *ILR Review*, forthcoming. [44](#), [67](#)
- Starr, Gerald (1981) “Minimum wage fixing: international experience with alternative roles,” *International Labour Review*, 120, 545. [4](#), [5](#)
- Stewart, Mark B. (2002) “Estimating the Impact of the Minimum Wage Using Geographical Wage Variation,” *Oxford Bulletin of Economics and Statistics*, 64 (supplement), 583–605. [30](#)
- Stewart, Mark B. and Joanna K. Swaffield (2008) “The Other Margin: Do Minimum Wages Cause Working Hours Adjustments for Low-Wage Workers?” *Economica*, 75 (297), 148–167. [41](#)
- Stigler, George J. (1946) “The Economics of Minimum Wage Legislation,” *The American Economic Review*, 36 (3), 358–365. [2](#), [6](#), [7](#), [8](#)
- Sun, Liyang and Sarah Abraham (2021) “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225 (2), 175–199. [15](#), [17](#), [114](#), [116](#)
- Sun, Liyang and Jesse M. Shapiro (2022) “A Linear Panel Model with Heterogeneous Coefficients and Variation in Exposure,” *Journal of Economic Perspectives*, 36 (4), 193–204. [29](#)
- Teulings, Coen N (2000) “Aggregation bias in elasticities of substitution and the minimum wage paradox,” *International Economic Review*, 41 (2), 359–398. [62](#)
- Totty, Evan (2017) “The effect of minimum wages on employment: A factor model approach,” *Economic Inquiry*, 55 (4), 1712–1737. [15](#)
- Umkehrer, Matthias and Philipp Berge (2020) “Evaluating the Minimum-Wage Exemption of the Long-Term Unemployed in Germany,” *ILR Review*, 73. [32](#)
- Vaghul, Kavya and Ben Zipperer (2016) “Historical state and sub-state minimum wage data,” *Washington Center for Equitable Growth Working Paper*, <http://cdn.equitablegrowth.org/wp-content/uploads/2016/09/02153029/090716-WP-Historical-min-wage-data.pdf>. [120](#), [122](#)
- Van Den Berg, Gerard J (2003) “Multiple equilibria and minimum wages in labor markets with informational frictions and heterogeneous production technologies,” *International Economic Review*, 44 (4), 1337–1357. [58](#)
- Vergara, Damian (2023) “Minimum wages and optimal redistribution: The role of firm profits.” [38](#), [54](#), [64](#), [102](#)
- Webb, Sidney (1912) “The economic theory of a legal minimum wage,” *Journal of Political Economy*, 20 (10), 973–998. [4](#)

- Wessels, Walter (2005) “Does the Minimum Wage Drive Teenagers Out of the Labor Force?” *Journal of Labor Research*, 26 (1), 169–176. [51](#), [100](#)
- Wiltshire, Justin (2022) “Walmart Supercenters and Monopsony Power: How a Large, Low-Wage Employer Impacts Local Labor Markets.” [39](#)
- Wiltshire, Justin C, Carl McPherson, Michael Reich, and Denis Sosinskiy (2023) “Minimum wage effects and monopsony explanations.” [102](#)
- Wooldridge, J.M. (2012) *Introductory Econometrics: A Modern Approach*: South-Western Cengage Learning. [14](#)
- Wursten, Jesse and Michael Reich (2023) “Racial inequality in frictional labor markets: Evidence from minimum wages,” *Labour Economics*, 82, 102344. [25](#), [37](#), [38](#), [102](#)
- Yao, Wenyun, Yuhang Qian, Hang Yang, and Wei Xu (2023) “Does minimum wages affect executive compensation? – Evidence from China,” *Pacific-Basin Finance Journal*, 80, 102080. [51](#), [102](#)
- Zavodny, Madeline (2000) “The effect of the minimum wage on employment and hours,” *Labour Economics*, 7 (6), 729–750. [40](#)

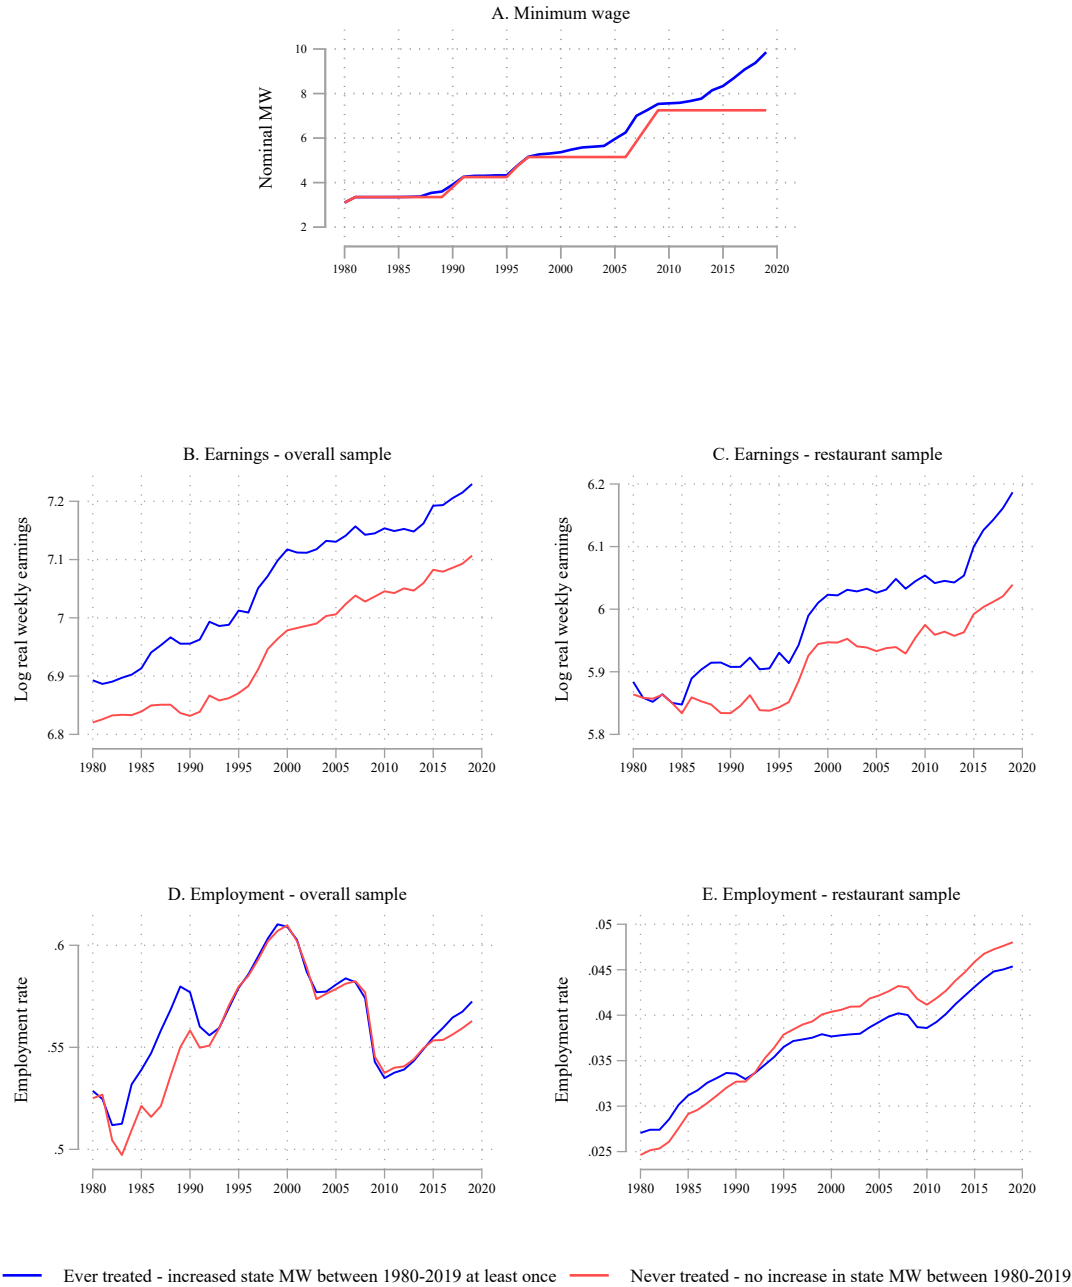
Figures

Figure 1: Evolution of minimum wage-to-average wage ratio (Kaitz index) over time in selected countries



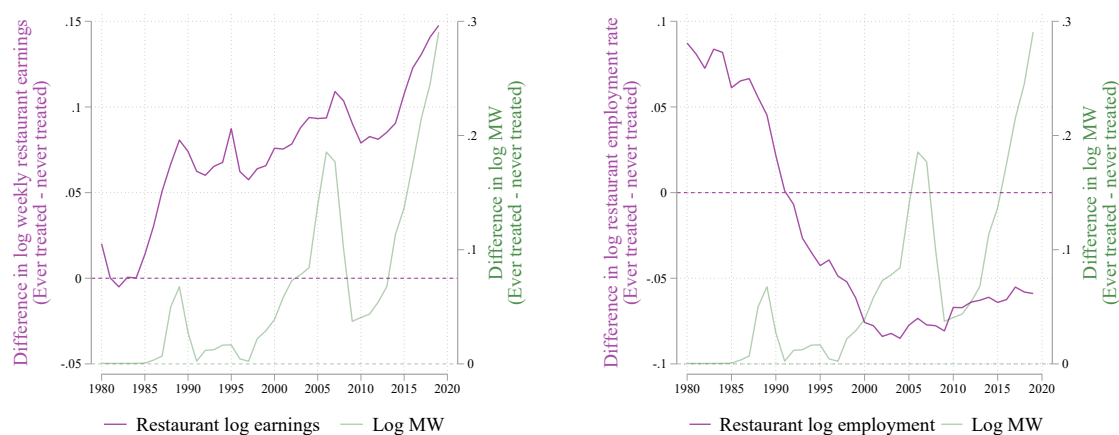
Notes: Data for France, Germany, Spain, UK and US comes from OECD - we scale the reported Kaitz index by 1.1 to account for the difference between average *full-time* wage and overall average wage. The USA-state estimate is a weighted Kaitz index taking into account state minimum wages and populations. For India, we use province-level data used in [Khurana et al. \(2023b\)](#) and [Khurana and Mahajan \(2020\)](#), provided to us by the authors, and average these for a national estimate. Data for China is from the Chinese National Bureau of Statistics and the China Statistical Yearbook. Data for Brazil comes from IPEA and the PNAD in later years. See Appendix C for detailed explanations.

Figure 2: Minimum wages, earnings and employment in ever-treated and never-treated states



Notes: These figures show the evolution of the minimum wage, earnings, and employment in ever treated and never treated states - the 35 ever treated states had at least one increase in MW above the federal level between 1980-2019, while the 15 never treated states had no such increase. Panel A shows how the nominal minimum wage changed and B and C show the progression of log average weekly earnings (in 2023 dollars) for the full sample and the restaurant sample respectively. Panel D and E show the log employment rate for these samples. All averages for the two groups are weighted by state population. For both earnings and employment in the restaurant sample, post 1990 data is as per the NAICS 3-digit definition. From 1980-1990, the data has been imputed using a scaling factor derived from the SIC/NAICS ratio in 1990. In addition, missing restaurant sector data for Alaska, Delaware, and Rhode Island in the 1990s has been imputed from sub-sector data using a scaling ratio of sector to sub-sector numbers. (details are in Appendix D).

Figure 3: Difference in minimum wage, restaurant employment and restaurant earnings between ever-treated and never-treated states




A. Log weekly restaurant earnings

B. Log restaurant employment rate

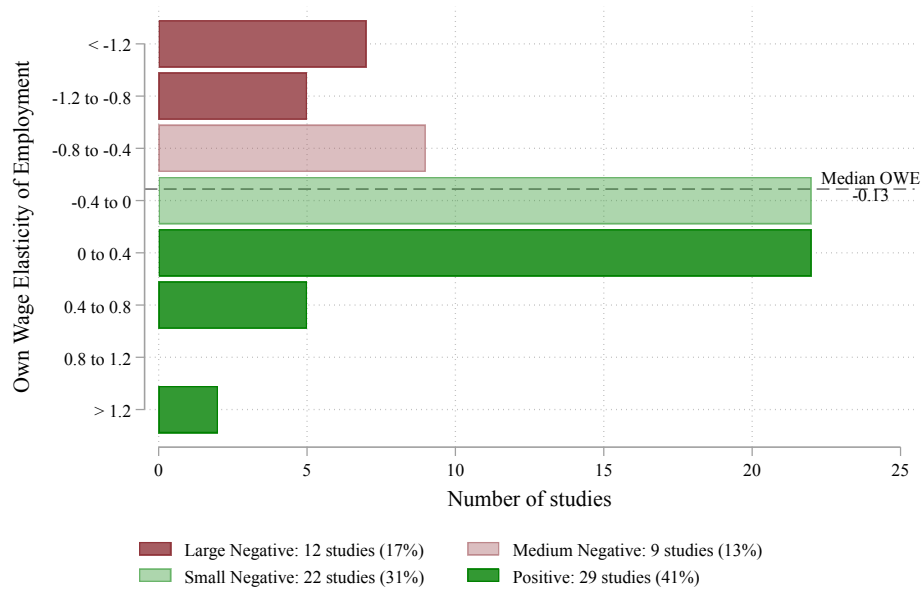
Notes: This figure plots the population-weighted difference in restaurant log earnings/log employment between 35 ever treated and 15 never treated states on one axis, and the difference in log MW between these states on the other. Log average weekly earnings (in 2023 dollars) are shown in Panel A, and log employment rate in Panel B. Ever treated states had at least one state MW increase, over and above any federal increases from 1980-2019. Never treated states had no such increase.

Figure 4: Impact of minimum wages on the frequency distribution of wages



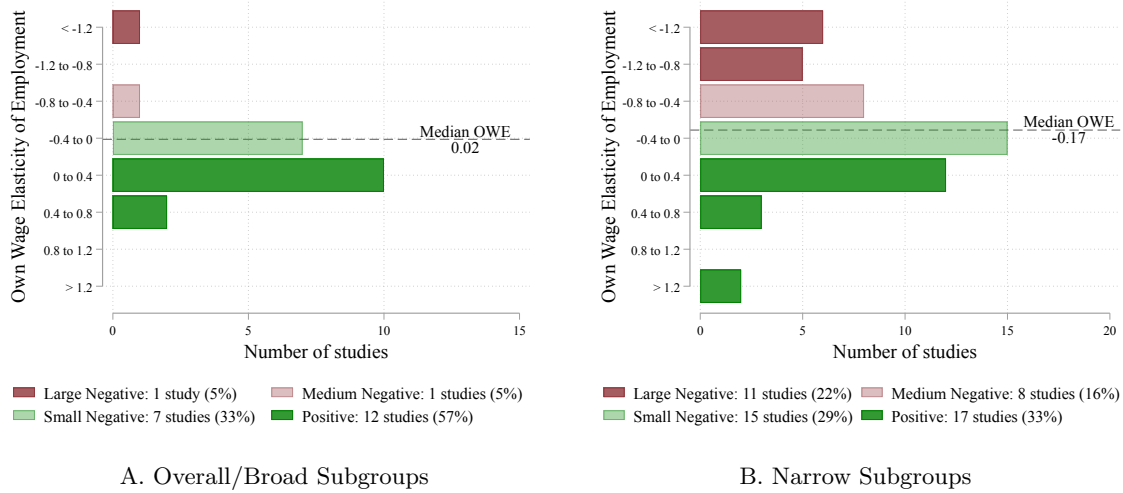
Source: Based on estimates underlying Figure 2 in [Cengiz et al. \(2019\)](#). Notes: The figure shows bin-by-bin employment changes resulting from 138 state-level minimum wage changes between 1979 and 2016. For each dollar bin (relative to the new minimum wage), the bars show the estimated average employment changes in that bin during the five-year post-treatment relative to the total employment in the state one year before the treatment. The dashed line plots the running sum of employment changes up to the respective wage bin.

Figure 5: Distribution of published studies, by own-wage elasticity estimate range




Notes: The figure shows the frequency distribution of own-wage elasticities from the [Dube and Zipperer \(2024\)](#) database. The bars are additionally colored differently for four broad groups: Positive ($OWE \geq 0$); Small Negative ($-0.4 \leq OWE < 0$); Medium Negative ($-0.8 \leq OWE < -0.4$); and Large Negative ($OWE < -0.8$).

Figure 6: Distribution of published studies, overall/broad and narrow subgroups



Notes: The figure shows the frequency distribution of own-wage elasticities from [Dube and Zipperer \(2024\)](#), separately for overall/broad and narrow subgroups. Common examples of narrow subgroups are teens or restaurant workers. Positive ($OWE \geq 0$); Small Negative ($-0.4 \leq OWE < 0$); Medium Negative ($-0.8 \leq OWE < -0.4$); and Large Negative ($OWE < -0.8$).

Figure 7: Event-by-event employment and earnings estimates for the restaurant sector



Notes: These figures plot the event-by-event earnings and employment estimates for the restaurant industry. Thus, each point estimate is from an event-study regression (Equation 5) with a particular combined event and its corresponding clean controls. The dependent variable is log average weekly earnings (2023 dollars) for Panel A, and log employment rate for Panel B. All regressions use state population weights. The 95% confidence intervals use Ferman-Pinto standard errors. We identify 60 events, where successive minimum wage increases are combined into a single event (see Appendix E). The events are sorted by total increase in log minimum wage.

Tables

Table 1: Effects of increased MW - Event study and TWFE estimates

	(1) Restaurant sample	(2) Overall sample	(3) High prob sample	(4) Teen sample	(5) High recall sample	(6) Low prob sample
A. 1980-2019 event study (Total events: 60): $\Delta \log MW = 0.261$ (s.e. = 0.016)						
Log wages	0.031** (0.006)	0.014** (0.004)	0.034** (0.006)	0.059** (0.012)	0.019* (0.008)	0.008 (0.005)
Log employment	-0.001 (0.005)	0.005 (0.004)	-0.002 (0.009)	0.007 (0.016)	0.006 (0.007)	0.003 (0.002)
OWE	-0.026 (0.151)	0.331 (0.217)	-0.067 (0.277)	0.116 (0.267)	0.335 (0.303)	- -
B. 1998-2019 event study (Total events: 46): $\Delta \log MW = 0.288$ (s.e. = 0.015)						
Log wages	0.035** (0.005)	0.014* (0.007)	0.034** (0.006)	0.062** (0.015)	0.018* (0.008)	0.010* (0.005)
Log employment	0.005 (0.005)	0.007 (0.004)	0.003 (0.011)	0.010 (0.016)	0.009 (0.007)	0.004 (0.002)
OWE	0.146 (0.140)	0.494** (0.177)	0.083 (0.337)	0.157 (0.260)	0.512 (0.369)	- -
C. 1980-2019 TWFE: implied $\Delta \log MW = 1$						
Log wages	0.281** (0.033)	0.158** (0.042)	0.183** (0.021)	0.405** (0.037)	0.039* (0.019)	0.042 (0.027)
Log employment	-0.146 (0.101)	-0.020 (0.037)	-0.210** (0.048)	-0.227** (0.059)	0.040 (0.036)	0.040** (0.015)
OWE	-0.519 (0.348)	-0.128 (0.230)	-1.147** (0.286)	-0.561** (0.163)	1.015 (1.146)	- -
D. 1998-2019 TWFE: implied $\Delta \log MW = 1$						
Log wages	0.211** (0.023)	0.029 (0.026)	0.183** (0.035)	0.397** (0.055)	0.078* (0.034)	-0.021 (0.018)
Log employment	0.106 (0.088)	0.066 (0.040)	-0.116* (0.045)	-0.039 (0.062)	0.021 (0.026)	0.024* (0.010)
OWE	0.502 (0.406)	- -	-0.636* (0.242)	-0.097 (0.154)	0.263 (0.350)	- -
Data source	QCEW	QCEW	CPS	CPS	CPS	CPS

Each cell represents a coefficient from a separate regression. The columns are sub-samples, while the rows are outcomes. The first 2 rows in panels A and B report event study estimates from equation 5. The dependent variables are log real average wages and log employment rate. Events are starting years of prominent MW increases as explained in Appendix E. The estimated $\Delta \log MW$ reported in Panels A and B is the coefficient from the same regression with log MW as dependent variable. The first two rows in panels C and D are TWFE estimates from regressions of the outcome on log minimum wage (as per equation 3). As the independent variable itself is log MW for this regression, the implied effect on log MW is 1. The third row in each panel shows own wage elasticity which is the ratio of the employment estimate to the earnings estimate. For wages, the restaurant and overall samples use log real average weekly earnings, while all other samples use log real average hourly wage. Standard errors in parentheses are clustered by state, and all regressions use state population weights.

* $p < 0.05$, ** $p < 0.01$

Table 2: Margins of adjustment: Non-compliance and fringe benefits

Adjustment	Evidence	Summary
Non-compliance	Significant increase in wages (both in administrative and survey data) Modest (Ashenfelter and Smith, 1979 ; Bernhardt et al., 2013 ; Caliendo et al., 2019) to low non-compliance (Goraus-Tańska and Lewandowski, 2019 ; Clemens and Strain, 2022)	Not clear; likely to be small
Cutting fringe benefits	Amenities: No change in fringe benefits and working conditions (Simon and Kaestner, 2004); No change in non-cash benefits (Harasztosi and Lindner, 2019) Health insurance: Reduction in health insurance coverage (Meiselbach and Abraham, 2023) offsetting 9-16% of wage increase for low wage workers (Clemens et al., 2018); reduction or no effect depending on the presence of non-discriminatory laws (Marks, 2011) On-the-job training: Mixed evidence: reduction in training (Neumark and Wascher, 2001a ; Hara, 2017); No change in training (Grossberg and Sicilian, 1999); increase in training (Acemoglu and Pischke, 2003 ; Arulampalam et al., 2004); No change in training, some decrease in intensity (Bellmann et al., 2017)	At most a small offset of the wage increase

Notes: This table reviews the literature on the margins of adjustment associated with non-compliance and fringe benefits.

Table 3: Margins of Adjustment: Employment and Refinements

Adjustment	Evidence	Summary
Reducing employment	See Section 3.3 on employment effects	Limited Employment change
Substitution (capital-labor)	<p>Firm-level on capital: evidence for some capital-labor substitution (Harasztosi and Lindner (2019) in Hungary; Hau et al. (2020); Geng et al. (2022) in China)</p> <p>Firm-level on robots: evidence for increased robot adoption in some periods, but not in others in China (Fan et al., 2021)</p> <p>Automation (inferred from occupation change): A raise in automation, disagreement on employment consequences (Lordan and Neumark, 2018; Aaronson and Phelan, 2019); Decrease in automation(Downey, 2021)</p>	Some indication for substitution away from certain types of workers
Substitution (labor-labor)	<p>Lack of heterogeneous impact by demographic groups (see Section 3.3.3)</p> <p>Substitution to more skilled or productive workers (Clemens et al., 2021; Horton, Forthcoming)</p>	No substitution b/w broad demographic groups, some substitution from lower to higher productive ones
Exit & entry	<p>Exit rate: Increase in exit rate (Draca et al., 2011; Mayneris et al., 2018; Aaronson et al., 2018; Harasztosi and Lindner, 2019; Luca and Luca, 2019; Dustmann et al., 2022; Chava et al., 2023); No increase in exit (Rohlin, 2011)</p> <p>Entry rate: Increase in entry rate (Aaronson et al., 2018); No effect on entry rate (Harasztosi and Lindner, 2019); Decrease on entry rate (Rohlin, 2011; Draca et al., 2011; Luca and Luca, 2019)</p> <p>Number of firms: A reduction in number of firms (Orazem and Mattila, 2002; Dustmann et al., 2022)</p>	Increase in exit rate, unclear effect on firm entry. Number of firms fall
Migration	<p>Low-skilled population size. Evidence for outmigration (Castillo-Freeman and Freeman, 1992) though results are fragile (Krueger, 1994). Decrease in share of low-skilled population (Monras, 2019) in U.S.; Increase in low-skilled population (Minton and Wheaton (2023) in U.S.; Giupponi et al. (2024) in UK, Ahlfeldt et al. (2018) in Germany).</p> <p>Immigrant population. Mixed evidence: Reduction in immigration (Orrenius and Zavodny, 2008; Cadena, 2014). Increase in immigration (Boffy-Ramirez, 2013; Giulietti, 2014)</p>	In most cases no significant change in the local population, unclear evidence on immigrants
Participation	<p>Extensive margin: Increase in overall participation (Laws, 2018; Agan and Makowsky, 2021), for prime-age unskilled (Luna-Alpizar, 2019); close to retirement (Borgschulte and Cho, 2020; Hampton and Totty, 2023) for parents with young children (Godøy et al., 2024); Reduction in participation overall (Wessels, 2005), for young (Lavecchia, 2020); No change in participation (Cengiz et al., 2022)</p> <p>Unemployed’s job search: Modest negative effect (Laws, 2018); No effect (Adams et al., 2022); Positive effect (Piqueras, 2023)</p>	Mixed evidence. Probably a limited effect on the participation, and some increase in job search

Notes: This table reviews the literature on the margins of adjustment associated with employment and various refinements of the employment estimates including capital-labor and labor-labor substitution, firm dynamics, migration and participation.

Table 4: The Effect of Minimum Wages on Prices

Study	Industry	Effect on Prices	Pass-through
Lemos (2006)	All consumers	Increase	Not reported (likely to be close to full-pass through)
Cooper et al. (2020)	All consumers	Increase	Not reported (likely to be close to full-pass through)
Aaronson (2001)	Restaurant	Increase	Pass-through depends on specification
Dube et al. (2007)	Restaurant	Increase for Fast-food restaurant	More than full pass-through
Aaronson et al. (2008)	Restaurant	Increase for Fast-food restaurant	More than full pass-through
Allegretto and Reich (2018)	Restaurant	Increase	Nearly full pass-through
Katz and Krueger (1992)	Fast-food restaurant	No effect (imprecisely estimated)	-
Card and Krueger (1995a)	Fast-food restaurant	Mixed (imprecisely estimated)	-
Basker and Khan (2016)	Fast-food restaurant	Increase	Sizeable pass-through
Ashenfelter and Jurajda (2022)	Fast-food restaurant	Increase	Full pass-through
Leung (2021)	Retail	Increase in grocery stores	More than full pass-through
Renkin et al. (2022)	Retail	Increase of grocery prices	Full pass-through
Agarwal et al. (2024)	Hotel	No effect on price	-
Harasztosi and Lindner (2019)	Manufacturing	Increase	Full pass-through

Notes: The table provides an overview of the impact of minimum wages on prices. It also shows the implied pass-through rate: the size of the change in the unit price of the output relative to the change in unit labor costs. In [Lemos \(2006\)](#) and [Cooper et al. \(2020\)](#), the overall increase is small for prices as a whole, nonetheless it likely represents close to a fully pass through given the small share of minimum wage workers in the economy.

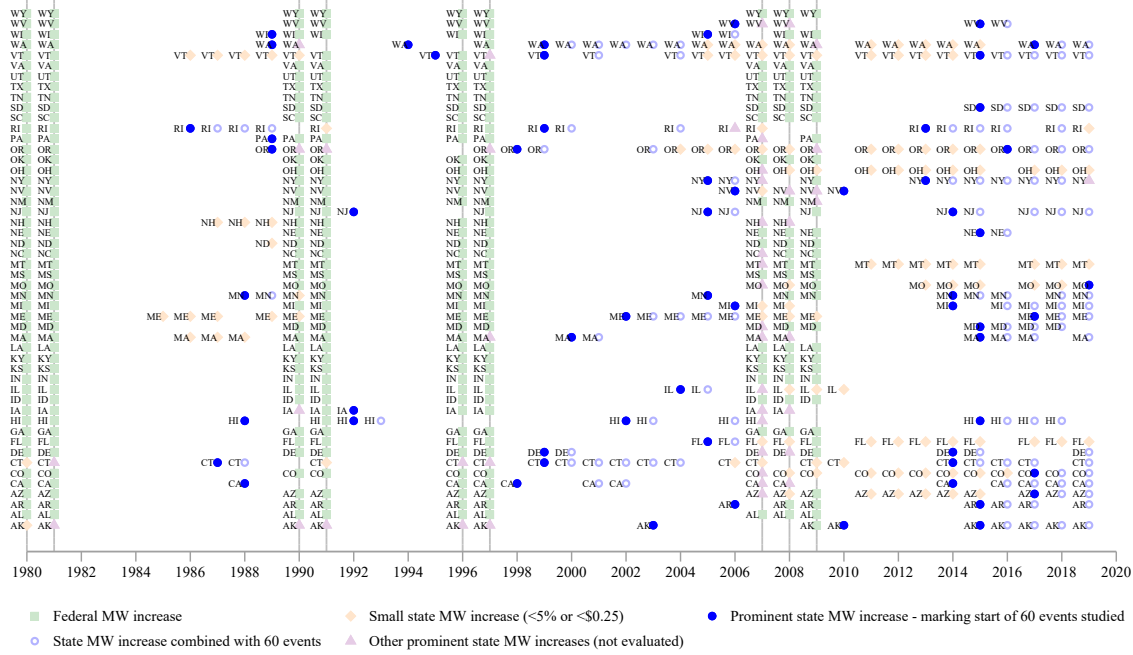
Table 5: Margins of adjustment: Incidence

Adjustment	Evidence	Summary
Prices	Significant price pass-through. See Table 4 for details	Almost full pass-through
Revenue or Quantities	Leung (2021) and Alonso (2022) find an <i>increase</i> in quantities following the minimum wage changes, while Renkin et al. (2022) find no effect. Agarwal et al. (2024) find reduction in hotel occupancy. Cooper et al. (2020) find increased quantity for food away consumption, no change in quantity for other goods. Harasztosi and Lindner (2019) find negative revenue effect in the tradable sector, and positive effect in the non-tradable sector.	No reduction in local service, significant reduction in tradable
Rents or Suppliers	No evidence on rents. Harasztosi and Lindner (2019) finds short term increase in spending on intermediate goods and no effect in the medium term	Unclear
High-skilled workers' wage	Within-firm: Small effect (Hirsch et al., 2015); Modest positive wage spillovers (Gopalan et al., 2021 ; Dube et al., 2019); Positive spillovers on executive compensations in China (Yao et al., 2023). See Section 5 for details.	Modest positive wage increase for slightly higher skilled workers, no effect on others
Profits	Stock-market value: No effect (Card and Krueger, 1994a); Large reduction (Bell and Machin, 2018) Profitability (Profit/Revenue): Reduction in the U.K. (Draca et al., 2011); No change in China (Mayneris et al., 2018) Profit: Increased profit among surviving pass-through businesses (Rao and Risch, 2024), Small effect on profit for restaurants in Georgia and Alabama (Hirsch et al., 2015); Small positive or negative effect on profits depending on the specification for a large U.S. retailer (Coviello et al., 2022); Modest fall in Hungary (Harasztosi and Lindner, 2019); Significant fall in Israel (Drucker et al., 2021); Significant fall in the U.S. (Vergara, 2023)	Mixed but likely a modest fall
Turnover	No change in turnover (Hirsch et al., 2015); Mixed evidence (Gopalan et al., 2021); Significant reduction in turnover (or increase in tenure) (Portugal and Cardoso, 2006 ; Dube et al., 2007 ; Brochu and Green, 2013 ; Dube et al., 2016 ; Gittings and Schmutte, 2016 ; Liu et al., 2016 ; Jardim et al., 2022 ; Coviello et al., 2022 ; Brochu et al., 2023 ; Wursten and Reich, 2023 ; Wiltshire et al., 2023)	Lower worker turnover
Productivity	Quantity increase: Ku (2022) in agriculture; Coviello et al. (2022) in retail; Hau et al. (2020) for exporting firms Quality increase: Ruffini (2022) in nursing home Indirect evidence on cost saving from turnover reduction: see above Change in efficiency through market allocation: Dustmann et al. (2022)	Increase in productivity

Notes: This table reviews the literature on the margins of adjustment associated with the incidence of the policy. The increase in labor costs can be borne by consumers (higher output prices, no decrease in quantity), suppliers, higher wage workers, lower profits, and higher productivity (reduced turnover, increased operational efficiency).

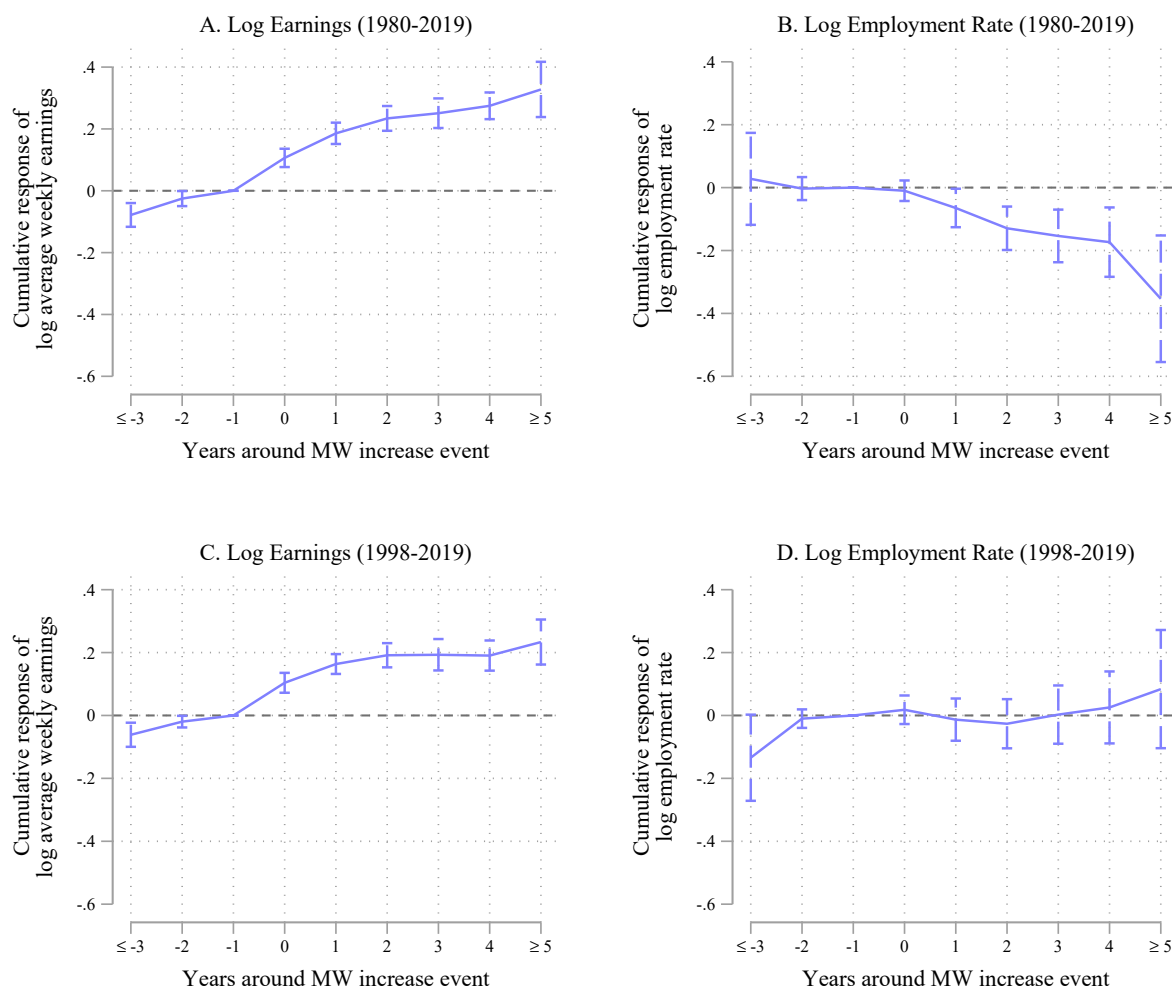
A Additional results

Figure A1: All minimum wage increases from 1980-2019



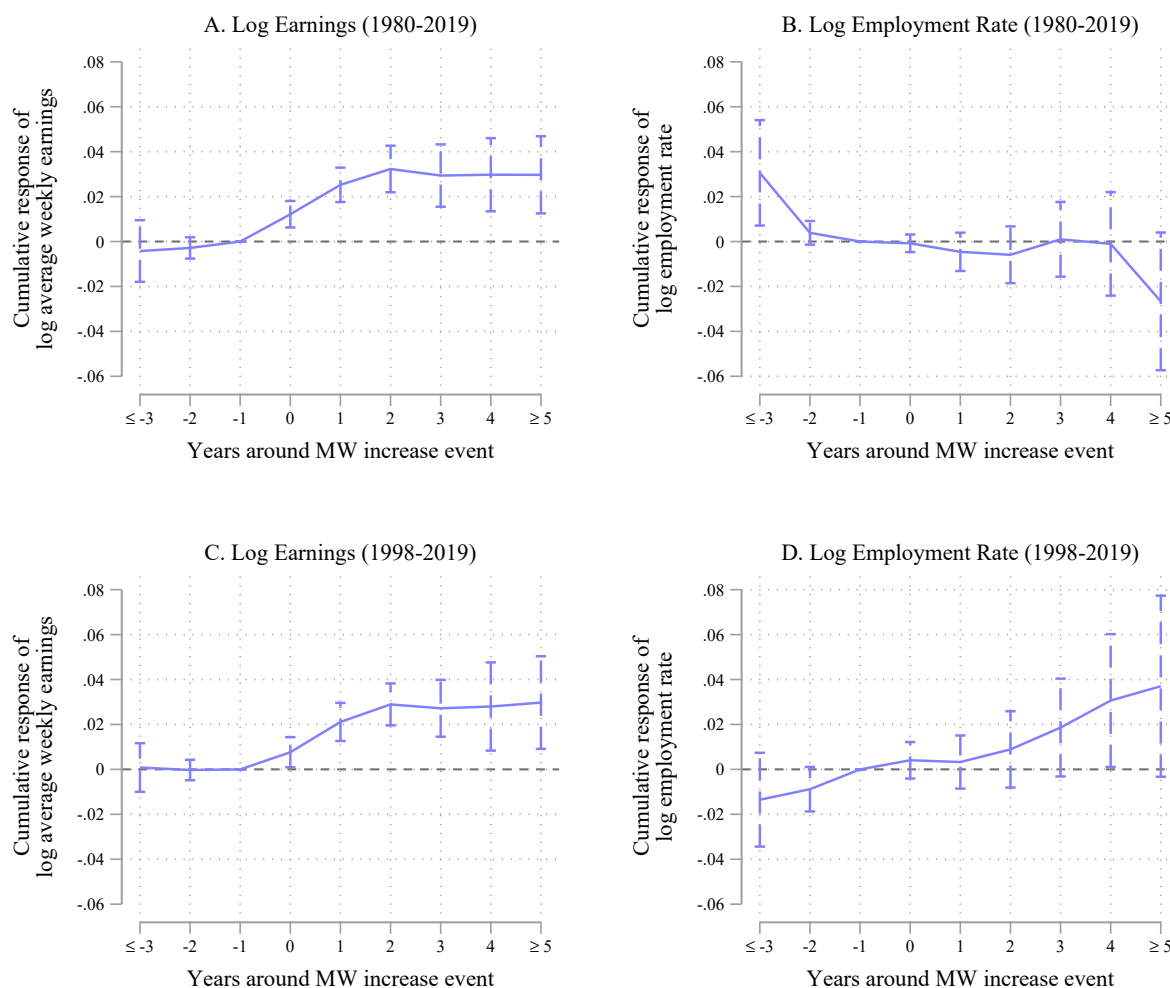
Notes: This figure plots all instances of the minimum wage increasing in any state. The blue solid circles represent the start of 60 events that we study in our main analysis. The blue hollow circles are state MW increases that are evaluated as part of the post-periods of these 60 events. There can also be state MW increases that are neither the start of an event, nor are evaluated as part of another event. These are either very small increases (some of them due to indexation) - represented by orange diamonds; or they are prominent increases, represented by purple triangles. Most of the latter are in years when the federal minimum wage also increased - the only two exceptions are RI 2006 and NY 2019, which are both both close enough to a prominent increase to not be counted as the start of a separate event, but more than 6 periods away from the last event-start year. The green squares show instances of increases due to a federal minimum wage increase. Any state that does not have a green square in a federal MW increase year either had its own increase in that year (small or prominent) which made its binding MW higher than the federal, or had a MW higher than the federal in the previous year as well, which made the federal increase redundant. More details are in Appendix Section E.

Figure A2: Impact of a log-point increase in MW on log employment rate and log average earnings in the restaurant sector over time for TWFE-logMW specification with distributed lags



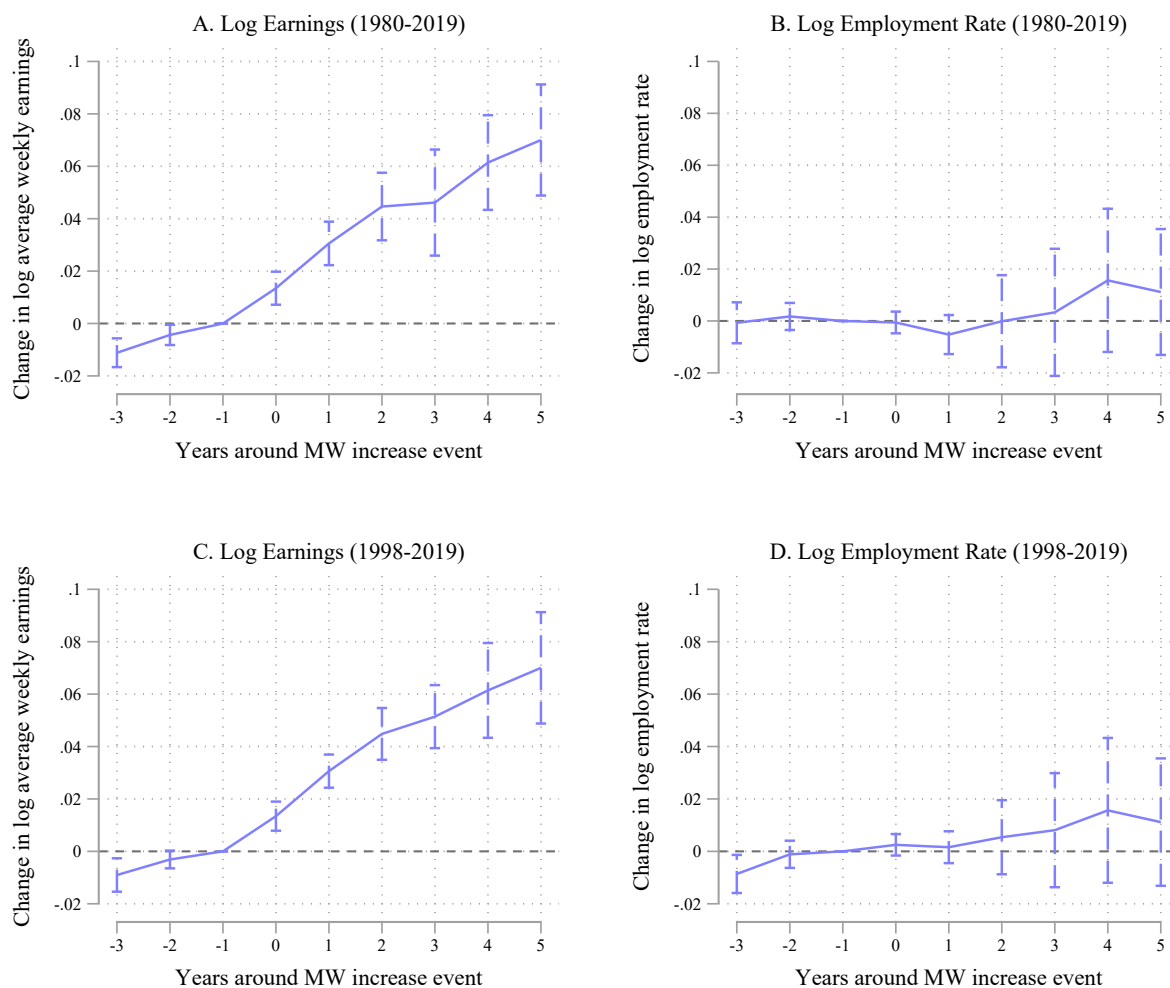
Notes: This figure plots the cumulative response of log employment rate and log average weekly earnings (2023 dollars) for restaurant workers to a log-point increase in the minimum wage (i.e., elasticities) using a distributed lag model with 2 leads and 5 lags of log minimum wage. The cumulative responses are normalized relative to event date (-1). 95% confidence intervals are based on standard errors clustered by state, and all regressions use state population weights. Panels A and C show the impact on earnings in the full and post-1998 time periods. Panels B and D show impacts on employment.

Figure A3: Impact of a binary increase in MW on log employment rate and log average earnings in the restaurant sector over time for TWFE-binary specification with distributed lags



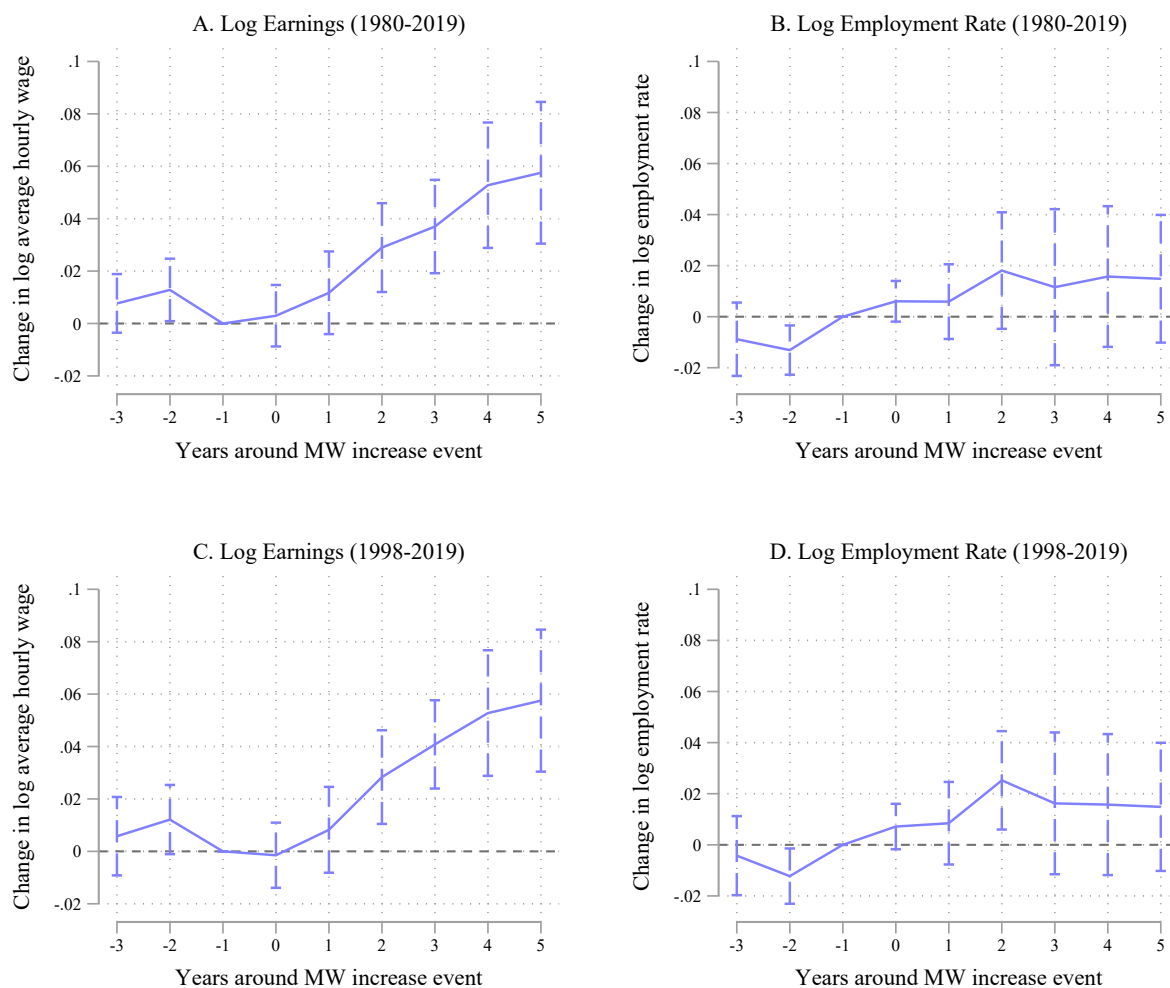
Notes: This figure plots the cumulative response of log employment rate and log average weekly earnings (2023 dollars) for restaurant workers to a binary increase in the minimum wage (i.e., semi-elasticities) using a distributed lag model with 2 leads and 5 lags of log minimum wage. The cumulative responses are normalized relative to event date (-1). 95% confidence intervals are based on standard errors clustered by state, and all regressions use state population weights. Panels A and C show the impact on earnings in the full and post-1998 time periods. Panels B and D show impacts on employment.

Figure A4: Impact of an increase in MW on log employment rate and log average earnings in the restaurant sector over time for event-study specification



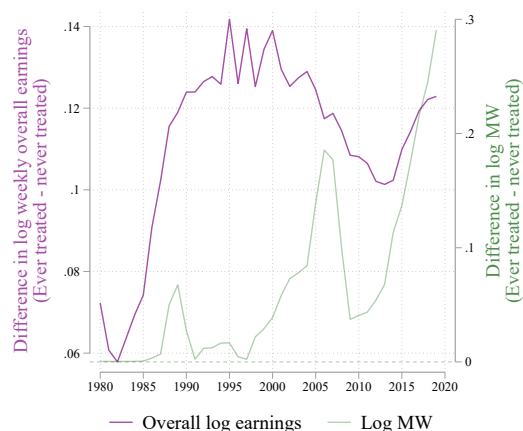
Notes: This figure shows the impact of a minimum wage increase on restaurant workers' log average weekly earnings in 2023 dollars (Panel A and Panel C) and log employment rate (Panel B and Panel D) in an event-study design. The top two panels are for the full period (1980-2019), while the bottom two only use data post-1998. The plotted points are coefficients from separate regressions using Equation 7. Standard errors are clustered by state, and all regressions use state population weights.

Figure A5: Impact of an increase in MW on log employment rate and log average earnings in the high recall sample over time for event-study specification

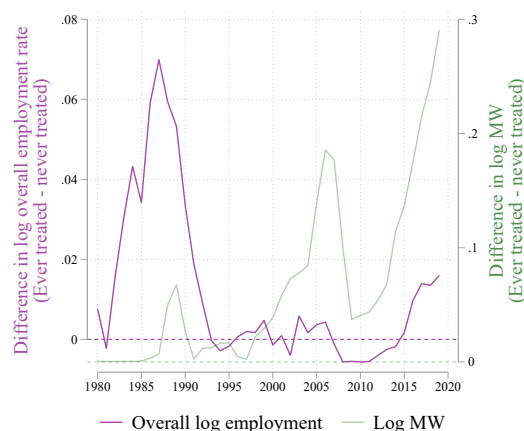


Notes: This figure shows the impact of a minimum wage increase on the high recall sample's log average hourly wage in 2023 dollars (Panel A and Panel C) and log employment rate (Panel B and Panel D) in an event-study design. The top two panels are for the full period (1980-2019), while the bottom two only use data post-1998. The plotted points are coefficients from separate regressions using Equation 7. Standard errors are clustered by state, and all regressions use state population weights.

Figure A6: Difference in minimum wage, overall employment rate and overall log average earnings between ever-treated and never-treated states




A. Log weekly overall earnings



B. Log overall employment rate

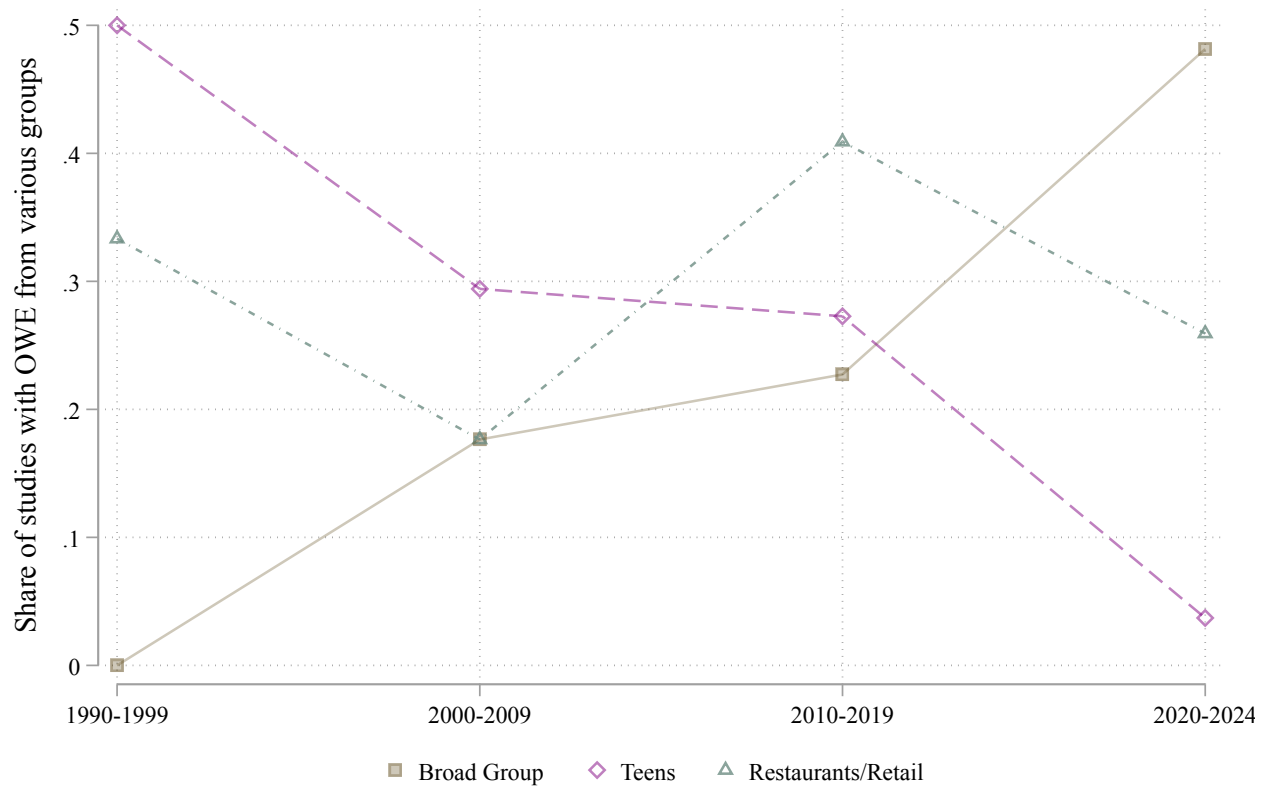
Notes: This figure plots the population-weighted difference in overall log earnings/log employment between 35 ever treated and 15 never treated states on one axis, and the difference in log MW between these states on the other. Log average weekly earnings (in 2023 dollars) are shown in Panel A, and log employment rate in Panel B. Ever treated states had at least one state MW increase, over and above any federal increases from 1980-2019. Never treated states had no such increase.

Figure A7: Event-by-event employment and earnings estimates for high-recall sample



Notes: These figures plot the event-by-event earnings and employment estimates for the high recall sample. Thus, each point estimate is from an event-study regression (Equation 5) with a particular combined event and its corresponding clean controls. The dependent variable is log average hourly wage (2023 dollars) for Panel A, and log employment rate for Panel B. All regressions use state population weights. The 95% confidence intervals use Ferman-Pinto standard errors. We identify 60 events, where successive minimum wage increases are combined into a single event (see Appendix E). The events are sorted by total increase in log minimum wage.

Figure A8: Evolution of MW studies over time: the share of groups being studied over time



Notes: This figure plots the share of MW studies in the [Dube and Zipperer \(2024\)](#) database studying various groups over time. Starting from the 1990s, the share of studies on teens has steadily declined, while a much larger number of studies now focus on broader groups. Studies on restaurant/retail sectors have declined somewhat, but not as much as teens.

Table A1: Details for the 60 combined minimum wage events, by year of the initial increase

	Number of events	Post period length(years)	No. of clean control states	Mean increase in log MW
1986	1	4	36	.238
1987	1	3	36	.232
1988	3	2	36	.172
1989	4	1	36	.116
1992	3	4	35	.158
1994	1	2	41	.142
1995	1	1	45	.057
1998	2	6	39	.249
1999	5	6	38	.269
2000	1	6	33	.251
2002	2	5	29	.261
2003	1	4	29	.235
2004	1	3	30	.233
2005	5	2	31	.245
2006	4	1	32	.2
2010	2	6	17	.078
2013	2	6	22	.368
2014	6	6	24	.302
2015	9	5	24	.266
2016	1	4	24	.196
2017	4	3	24	.305
2019	1	1	24	.091
Overall	60	3.9	30	.231

This table provides information by cohort on the 60 events used in the event-study analysis. The first column reports the number of events in that cohort/year. The second column gives the length of the post-period for that cohort in years. Post-periods are a maximum of 6 years unless interrupted by a federal MW increase year or end-of-sample. One exception is that Alaska 2010 has a five year post-period to prevent overlap with a 2016 event in the same state. The other event in the 2010 cohort has a six-year post-period (the maximum possible). The third column reports number of states that serve as clean controls for the cohort - a clean control needs to have no state MW increase in the three years before the cohort year, and also no state MW increase in the relevant post-period. Finally, the last column reports the average of the log MW increase (log MW in the last year of the post-period - log MW in year $t-1$, where t is the event year) across all events in a cohort. The last row gives the total number of events and means *by event* (not by cohort) of post-period length, clean control states, and log MW increase.

Table A2: Effects of increased MW on restaurant employment - Quasi event study estimates

	(1) Short pre-period (2010-2013)	(2) Long pre-period (1980-1988; 1990-2000; 2009-2013)
Log MW	0.145*** (0.028)	0.174*** (0.030)
Log wages	0.042*** (0.010)	0.070*** (0.014)
Log employment	0.005 (0.012)	-0.053** (0.024)
OWE	0.125 (0.266)	-0.761** (0.359)

Each cell represents a coefficient from a separate regression. The first three rows in the first column report estimates from Equation 8 for the restaurant sample, i.e, this is like an event study regression where the event start date is 2014 for 35 ever-treated states, and the sample consists of observations from only 2010-2019. The pre-period is thus 2010-2013 and the post-period is 2014-2019. The second column extends the pre-period to also include all years before 2010 in which the difference in log MW between ever-treated and never-treated states is no greater than the difference in 2010-2013. This way, we end up including 1980-1988; 1990-2000; and 2009. Both columns show effects on log MW, log real average weekly earnings, and log employment rate; and the last row reports own-wage elasticity which is the ratio of the log employment and log earnings estimates. Standard errors in parentheses are clustered by state, and all regressions use state population weights.

* $p < 0.05$, ** $p < 0.01$

Table A3: Heterogeneity analysis and robustness for impacts of increased MW

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	All events	Long events	Long post events	Top tercile events	Equally weighted ATT	Control for low prob emp	TWFE binary
A. Restaurant sample (Data source: QCEW)							
Log wages	0.031** (0.006)	0.050** (0.007)	0.045** (0.005)	0.040** (0.007)	0.032** (0.006)	0.031** (0.006)	0.032** (0.008)
Log employment	-0.001 (0.005)	0.008 (0.009)	0.004 (0.007)	0.009 (0.006)	0.001 (0.005)	-0.001 (0.005)	-0.046* (0.019)
OWE	-0.026 (0.151)	0.166 (0.176)	0.084 (0.161)	0.233 (0.145)	0.016 (0.162)	-0.019 (0.149)	-1.432* (0.659)
B. High recall sample (Data source: CPS)							
Log wages	0.019* (0.008)	0.033** (0.008)	0.029** (0.007)	0.020 (0.012)	0.019* (0.009)	0.018* (0.009)	-0.004 (0.005)
Log employment	0.006 (0.007)	0.020 (0.010)	0.018* (0.008)	0.008 (0.011)	0.007 (0.007)	0.003 (0.006)	0.003 (0.006)
OWE	0.335 (0.303)	0.605 (0.347)	0.603* (0.269)	0.413 (0.428)	0.346 (0.401)	0.167 (0.350)	-0.615 (1.626)

This table reports various robustness checks for the original event-study estimates reported in Column 1 run as per Equation 5. Columns 2, 3, and 4 run the same regression, but with a subset of the original 60 events. Column 2 only keeps events that have at least a three year clean pre-period and a three year clean post-period. Column 3 has a weaker constraint, and keeps all events that have at least a three year post-period. Having a three year post-period means not having any federal increases until at least two years after the event, and having a three year pre-period means having no federal increases in the three years preceding the event. Column 4 restricts attention to the top tercile of events in terms of increase in the log minimum wage caused by the event. Columns 5 and 6 report an equally weighted ATT, using regression adjustment. The only difference is that while Column 5 only controls for year fixed effects, Column 6 additionally controls for low probability sample employment rate. Finally, Column (7) reports estimates from a TWFE regression with the cumulative number of MW increase events in a state as the independent variable (this is a variation on Equation 3). All columns have effects on log real average wages, effects on log employment rate, and own-wage elasticity which is the ratio of the latter to the former, for both restaurant and high recall samples. For wages, the restaurant sample uses log real weekly average earnings, while the high recall sample uses log real hourly average wage. Standard errors in parentheses are clustered by state, and all regressions use state population weights.

* $p < 0.05$, ** $p < 0.01$

B Bias from heterogeneous pre-existing trends: A simulation study

Recent literature on difference-in-differences (DiD) has highlighted how, under staggered treatment adoption, heterogeneous treatment effects combined with negative weighting can lead to spurious estimated dynamics when using Two-Way Fixed Effects (TWFE) models with distributed lags (see [Sun and Abraham \(2021\)](#)). However, another key source of bias is the interaction between the negative weighting problem in TWFE models and pre-treatment trends. This interaction can lead to highly misleading estimates of dynamic responses that (1) obscure actual pre-existing trends and (2) falsely suggest treatment effects. These issues are further exacerbated by a related concern: unlike DiD models with defined event windows, the TWFE model requires a stronger assumption that potential outcomes follow parallel trends over the entire sample period ([Marcus and Sant’Anna, 2021](#); [Roth et al., 2023](#)).

In this Appendix, we use a simple Monte Carlo simulation with two treatment cohorts and staggered adoption to illustrate these issues. We demonstrate that heterogeneous violations of the parallel trends assumption can cause TWFE distributed lag models to generate highly misleading inferences, even when the true treatment effects are constant (and zero). Importantly, in our example, the average violation of parallel trends is small and occurs well before the treatment events, outside the defined event window. Consequently, while the TWFE model’s stronger parallel trends assumption results in biased estimates, modern difference-in-differences event study estimators, such as those proposed by [Callaway and Sant’Anna \(2021\)](#) and [Dube et al. \(2023\)](#) provide unbiased estimates.

Simulation design

The basic setup for the data generating process (DGP) reflects a staggered adoption process:

- Time runs from $t = 1$ to $t = 15$.
- There are 50 states, divided into three groups: 30 never-treated states, 10 early-adopter states (treated at time $t = 9$), and 10 late-adopter states (treated at time $t = 11$).
- The true treatment effect β is equal to zero.
- Errors ν_{st} at the s -by- t level are normally distributed with a standard deviation of 0.05.

The key feature of the DGP here is that early-adopter states have a **negative latent trend** until time $t = 6$. In contrast, late-adopter states have a **positive latent trend** per year until $t = 6$. There are no latent trends after $t = 6$ for either group. All trends are relative to the never-treated group. The opposite signs in the latent trends means that the overall violation of parallel trends is small. However, as we will see, in combination with the staggered adoption this can lead to a large bias in TWFE estimates.

These assumption can be expressed formally using a potential outcomes model. For state s at time t , there are state fixed effects γ_s , time fixed effects τ_t , and treatment D_{st} , which takes the value 1 if state s is treated at time t , and 0 otherwise. The treatment effect is denoted by β . The potential outcomes Y_{st}^0 are described as follows:

$$Y_{st}^0 = \gamma_s + \tau_t + \text{Latent trend}_{st} + \nu_{st}$$

The heterogeneous latent trends (i.e., violations of parallel trends) for the early and late-treated groups are specified as:

$$\text{Latent trend}_{st} = \begin{cases} -0.35 \times t & \text{if } s \in \text{early-adopter and } t < 6 \\ -1.75 & \text{if } s \in \text{early-adopter and } t \geq 6 \\ 0.45 \times t & \text{if } s \in \text{late-adopter and } t < 6 \\ 2.25 & \text{if } s \in \text{late-adopter and } t \geq 6 \\ 0 & \text{if } s \in \text{never-treated} \end{cases}$$

The overall outcome under this data generating process can be expressed as:

$$Y_{st} = \gamma_s + \tau_t + \beta D_{st} + \text{Latent trend}_{st} + \nu_{st}$$

Time effects and state effects are assumed to be distributed normally with mean zero, and standard deviation of 0.05.

Figure B1 plots the mean potential outcomes in dashed lines and outcomes under treatment in solid lines for these three groups of states. We can see the heterogeneous violations of parallel trends before $t = 5$, but afterward the potential outcomes follow parallel trends. In addition, the effect of the treatment is constant (in this case, zero), and so the potential and actual mean outcomes coincide in the units treated following treatment.

This setup illustrates a violation of the assumption of parallel trends before time $t = 5$, since the early and late adopting groups exhibit opposite-signed latent trends. However, after time $t = 5$, the assumption of parallel trends holds, as the change in expected potential outcomes is the same in all 3 groups.

Event study versus TWFE distributed lag estimates

We begin by estimating the dynamic treatment effects using well-suited difference-in-differences (DiD) event study designs, specifically the [Callaway and Sant'Anna \(2021\)](#) and LP-DiD estimators. Recall that LP-DiD, when applied without re-weighting—as in this implementation—is equivalent to the stacked regression estimator proposed by [Cengiz et al. \(2019\)](#). Recall that both treatment events occur after $t = 6$, at periods 9 and 11. This means that if we focus on an event window covering 3

years before and 3 years after treatment, DiD methods such as Callaway-Sant’anna or LP-DiD would estimate a null effect, as the true causal effect is zero, and there is no bias due to violations of parallel trends within the event windows.

Figure B2 presents event study estimates from 250 simulations, focusing on an event window from periods -3 to 3. The solid lines represent the mean estimates from the simulations, while the shaded areas show the $2 \times SD$ interval around the mean, where SD is the standard deviation of the estimates across simulations. Both the LP-DiD and Callaway-Sant’Anna estimators accurately identify the absence of pre-treatment trends, consistent with the true data-generating process. Additionally, both estimators correctly detect zero post-treatment effects, as expected, since the treatment has no impact beyond the event period.

Next, we apply a TWFE distributed lag model to the same scenario. This model includes 2 leads of treatment, the contemporaneous treatment, and 3 lags of treatment. We normalize all coefficients relative to event time -1, meaning that estimates for event time -3 or earlier are combined (or “binned”) into a single coefficient, as are estimates for event time 3 or later. We assume no changes in treatment status before time 1 or after time 14, allowing us to define leads and lags for all 750 observations. Cumulative responses are then calculated by summing the effects over event time.

The mean impulse responses from the TWFE distributed lag model also show no obvious pre-treatment trends, including in the binned -3 event time; this is not itself a problem, as the opposite signed early-period violations tend to mostly cancel each other out. However, the model suggests a small negative “short-run” effect, followed by a large negative “long-run” effect, despite the true effect being zero throughout. This spurious result arises because the negative weighting problem inherent in TWFE models interacts with a violation of parallel trends before time $t = 5$, leading to a highly misleadingly shaped impulse response function.

The most concerning issue is that while the source of the bias lies in a parallel trends violation substantially prior to the event window, the erroneous inference manifests as a lagged treatment effect. Moreover, we would not be able to detect this problem by inspecting the leading coefficients. This is the mirror image of the problem identified by [Sun and Abraham \(2021\)](#), where heterogeneous treatment effects produce misleading pre-trends, even when the parallel trends assumption is satisfied in reality.

Overall, this simulation demonstrates how heterogeneous pre-trends can result in biased inferences when using TWFE distributed lag models in a highly non-transparent fashion. With staggered treatment timing, even if the parallel trends assumption holds within the event window and the treatment effect is homogeneous, the combination of negative weighting and violations of parallel trends outside the event window produces spurious impulse response estimates. This is a consequence of the stronger form of parallel trends assumption needed for consistency under the TWFE model: that it holds not only within the event window but across the whole sample period.

Fortunately, modern difference-in-differences estimators, such as Callaway-Sant’anna and LP-DiD, which utilize clean control groups and define a clear event window, are better at guarding against these biases and provide more reliable results.

Figure B1: Mean potential outcome, and outcome under treatment for the 3 groups

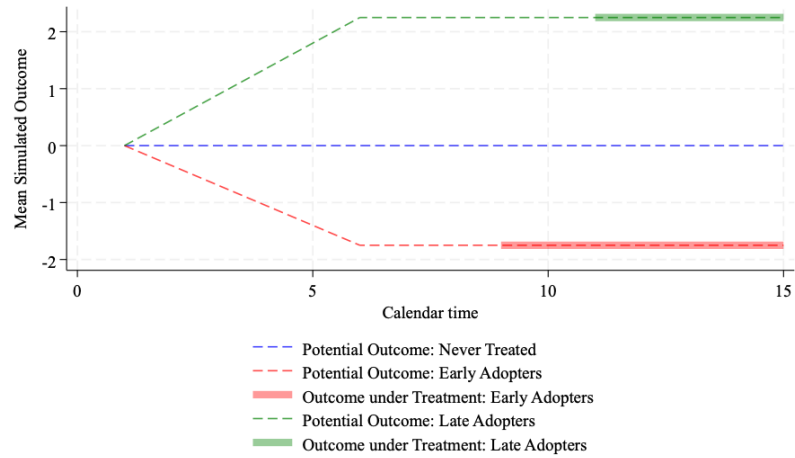



Figure B2: TWFE-distributed lag versus Event Study estimates using Callaway-Sant'Anna, and Local Projections DiD (equivalently Stacked DiD)



Notes: Results from N=100 simulations. The shaded regions represent 2xSD intervals based on the estimated treatment effects.

C Data sources for cross-country Kaitz indices

Brazil

Minimum wage data can be accessed from the *Instituto de Pesquisa Econômica Aplicada* or IPEA. Average wage for 2002-2016 also comes from IPEA, from the survey *Pesquisa Mensal de Emprego* ([PME](#)). PME was a monthly employment survey of 6 metropolitan areas (Recife, Salvador, Belo Horizonte, Rio de Janeiro, São Paulo and Porto Alegre) that was discontinued in 2016 and was replaced with the PNAD (below) with covers the whole country

For 2012-2023, we access average wage data from the household survey *Pesquisa Nacional por Amostra de Domicílios Contínua* ([PNAD](#)).

To harmonize these two sources, we created a wage variable that is equal to the average wage from PNAD for 2012+. Then, for 2002-2011, we used the metropolitan average wage from PME multiplied by the average ratio of the overall to metropolitan average wage for the overlapping years (2012-2016).

China

Average wage and employment numbers for China come from the Chinese National Bureau of Statistics (NBS) [interactive data explorer](#). Some earlier data is not available there, but was gathered manually from the China Statistical Yearbook - average wage data for 1995-1999 was sourced from the 2010 Statistical Yearbook.

NBS only reports average wage for workers employed in urban areas, so we use urban average wages. We use the following indicators from NBS: Average Wage of Employed Persons in Urban Units(yuan), Average Wage of Employed Persons in Urban Private Units(yuan), Urban Employed Persons(10000 persons), and Urban Employed Persons, Private Enterprises(10000 persons). Note that “Urban units” does not include private, so the indicator Average Wage of Employed Persons in Urban Units(yuan) is the average wages for workers in the non-private sector.

We calculate the proportion employed in each year and use that to take a weighted average of the average wage in the private and non-private sector to get the national annual average wage.

The China NBS stops reporting the number of urban employed persons in private enterprises in 2019. So, after 2019, we hold the proportion of urban employment in private enterprises constant (i.e. we use the 2019 value for 2020+).

Finally, NBS only reports private (urban) average wages starting in 2009. We calculate the private average wage prior to 2009 ($t < 2009$), by multiplying the average wage in year t by the ratio between the average private wage to average non-private wage in 2009.

Monthly nominal minimum wage data is from the [ILO](#). We multiply the monthly minimum wage by 12 to get the annual minimum.

India

We use regional data used in [Khurana et al. \(2023b\)](#) and [Khurana and Mahajan \(2020\)](#) provided to us by the authors - the data is on average daily wages and regional population for the years 1983-84, 1993-94, 1999-00, 2004-05, 2007-08, 2009-10, 2011-12, 2017-18 and 2018-19. Average wage data is for age ranges 15-59. We also use daily minimum wage data for unskilled agricultural workers (the lowest minimum wage) for the years 1999-00, 2004-05, 2007-08, 2009-10, 2011-12, 2017-18 and 2018-19 provided to us by Khurana and Mahajan. We use minimum wage data from Menon and Rodgers for the years 1983 and 1993 available [here](#).

We create a Kaitz for India by first calculating region-specific Kaitz and then using a population-weighted average of these for each year to get the national Kaitz.

The workers' wage data is based on surveys conducted from July to June, and minimum wage data (provided by Mahajan and Khurana) is from January to December. For surveys conducted from July to December in year t , the minimum wages for t are used, and for surveys from January to June in year $t + 1$, the minimum wages for $t+1$ are used. Therefore, the t - $t+1$ wages data correspond to the average of the t and $t+1$ minimum wage data, that is, the 1999-2000 worker's average wages data is linked to the average of the 1999 and 2000 administrative minimum wage data. Wages for the years 1983-84, 1993-94 are paired with minimum wage data from Rodgers and Menon for the years 1983 and 1993, respectively. In the figure, we plot Kaitz based on the first year in which the survey was conducted, so, for example, the Kaitz associated with wages and population numbers for the 1993-94 household survey is plotted as the Kaitz in 1993.

Minimum wage data from Rodgers and Menon is based on industry and occupation groups. Because there is not always an "unskilled" agriculture minimum wage for each region in their data, we take the lowest minimum wage in the agriculture sector of each region as the minimum wage for the Kaitz index.

The data for the years 1999-00, 2004-05, 2007-08, 2009-10, 2011-12, 2017-18 and 2018-19 pertains to 18 major states of India according to the Census 2001. For the newly created states from the old states, administrative minimum wages are allocated district-wise from the year these states began reporting. For example, Telangana, which was carved out of Andhra Pradesh in 2014, uses the minimum wages of Telangana for the districts in Andhra Pradesh that were transferred to Telangana since the time Telangana started reporting the minimum wages.

For the years 1983-84, 1993-94, we merge wage and population data from Khurana and Mahajan with minimum wage data from Menon and Rodgers. The result of this is that in 1993-94, there are 16 states represented, but in 1983-84, we only have data to construct Kaitz for 9 states. In order to correct for this, we calculate Kaitz indexes in 1999 based on the subset of overlapping states between 1999 and 1983 (or 1993). Then we then multiply the Kaitz indexes in 1983 (or 1993) by the ratio of

the 1999 Kaitz using all states and the 1999 Kaitz using only overlapping states. For 1983, all states are in both datasets (9 states). In 1993, one state, Orissa, does not overlap and, thus, is dropped, so there are 15 overlapping observations.

OECD countries

We use OECD data from France, Germany, Spain, U.K. and the US, on the Kaitz index, constructed as the ratio of the minimum wage to mean wages of full-time workers. We assume that the average full-time wage is around 10% higher than the overall average wage. So we scale the OECD Kaitz index by 1.1 to construct an estimated Kaitz for all workers (not just full-time workers).

United States

For the U.S., we supplement the OECD data by additionally constructing a state-minimum-wage based Kaitz measure. Specifically, we construct two measures as follows.

1. Kaitz with federal minimum wage: we do the same as for other OECD countries, and use the Kaitz index reported by OECD and multiply times 1.1. The OECD reports the Kaitz index starting in 1974. Pre-1973, we use the average hourly wage of production and non-supervisory workers in order to calculate Kaitz from 1964-1972.
2. Kaitz with state minimum wage: we use the Kaitz index from the OECD to construct a Kaitz index using state-level minimum wages starting in 1980. We use state-level minimum wage data over time from [Vaghul and Zipperer \(2016\)](#). We then construct an annual, weighted mean of the state minimum wage, where the weights are the counts of wage earners in each state-year cell. These counties are calculated using the CPS data (NBER extracts for 1980-1981, and IPUMS extracts for 1981-2022). We then multiply the OECD (federal) Kaitz by the average state MW/federal MW ratio to obtain the Kaitz using state minimums.

D Constructing historical QCEW restaurant data

Our results on the restaurant and overall sample come from Quarterly Census of Employment and Wages (QCEW) data. In this Appendix, we describe the process of cleaning and constructing this data set.

Harmonizing restaurant data from different classification systems (NAICS and SIC)

We downloaded the annual version of QCEW data for 1980-2019—however, not all of it is available under the same classification. Data from 1990 onwards is available under the North American Industry Classification System (NAICS), while data from 1980-2000 is available under the Standard Industrial Classification (SIC).

We use two groups of outcome variables: overall earnings and employment by state-year, and earnings and employment in the restaurant sector by state-year. There are separate SIC and NAICS based versions for both groups. For the overall outcomes, we use SIC-based variables before 1990, and NAICS-based variables from 1990 onward. As these are overall numbers, they should not differ across the two classification systems.⁴⁷

However, there are differences in restaurant earnings and employment we get from both classification systems in overlap years (1990-2000). For NAICS, we take data for the three-digit code 722 (“Food services and drinking places”), while for SIC, we use code SIC_0G581 (“Eating and drinking places”). To create a harmonized series for restaurant earnings and employment for 1980-2019, we use the following procedure. We first calculate the ratio of NAICS employment (or earnings) to SIC employment (or earnings) in 1990—the first year when we have data for both classification systems. Next, we multiply the SIC-based outcomes (employment or earnings) during 1980-1989 by this ratio to obtain our new *aligned* series for these years. For 1990 and later years, we simply use the NAICS data.

Figure D1 provides an illustration of this exercise. The top panel plots restaurant employment rate in “ever-treated” and “never-treated” states from both the NAICS and SIC datasets. We can see that, while close, the levels in the two datasets are not exactly equal in the overlapping years of 1990-2000. The bottom panel plots the aligned series where the post-1990 lines remain the same, but the pre-1990 lines have been aligned as described above to prevent any discontinuities.

Dealing with missing data for Alaska, Delaware, and Rhode Island in the 1990s

Three states—Alaska, Delaware, and Rhode Island—have missing or incomplete data for the restaurant sector (NAICS code 722) in the 1990s. For these states, we impute data for the 3-digit restaurant

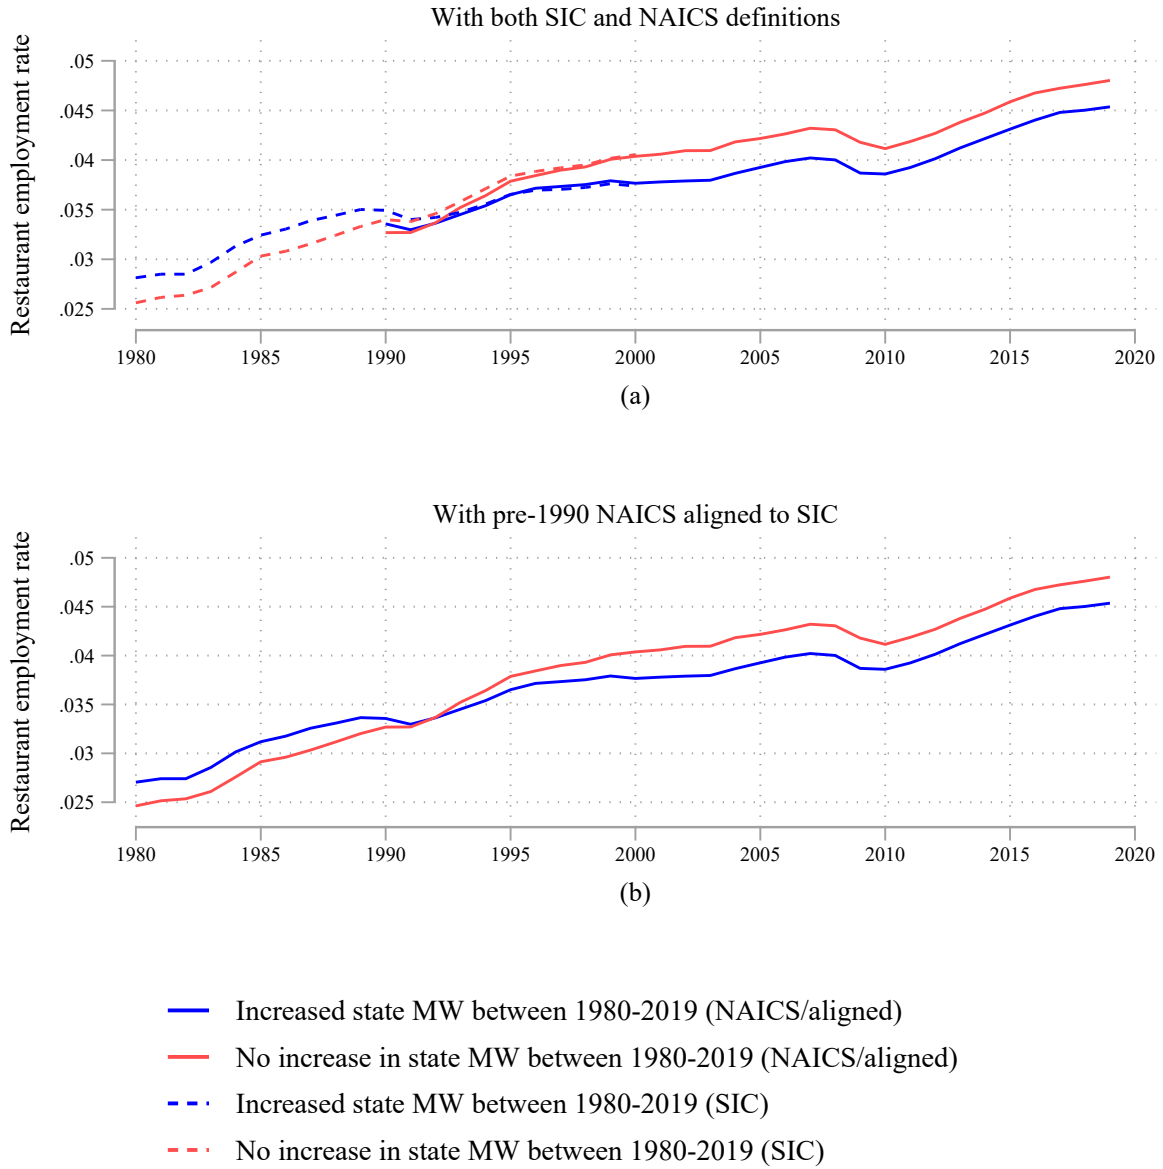
⁴⁷In practice, there are some very minor differences, but the mean difference between employment level and total wages in NAICS and SIC data in 1990 (the first year for which both overlap) is less than 0.1%

sector using available data from 4-digit sub-sectors. In the NAICS classification system, there are five 4-digit sub-sectors within the broadly defined restaurant sector (the description for NAICS code 722 is “Food services and drinking places”). These 5 sub-sectors are 7221 (Full-service restaurants), 7222 (Limited-service eating places), 7223 (Special food services), 7224 (Drinking places), and 7225 (Restaurants and other drinking places). Out of these, data on 7221 and 7222 is consistently available through the 1990s for all three of Alaska, Delaware, and Rhode Island. The other three sub-sectors are not available in one or more of the three. Thus, we use only 7221 and 7222 for our imputation. Clearly, just adding these two sub-sectors is not enough as they leave out some components of 722, which would create non-comparability of levels around 2000. We solve this using a similar method as the NAICS-SIC alignment as outlined above. For all three states, we take data for 2001 (which is the first year with full data), and take the ratio of employment under the 3-digit code 722 to the sum of employment under the 4-digit codes 7221 and 7222. We then multiply the 7221 + 7222 employment numbers from 1990-2000 by this factor. This gives us an estimate for employment for the full 3-digit sector in the 1990s for these states. We do the same for earnings (wage bill). Using this estimated NAICS data for 1990, we can apply the NAICS-SIC scaling described above to these three states as well, thus giving us a full harmonized series from 1980-2019.

Other data

We use the Current Population Survey (CPS) data to calculate earnings and employment for various groups (explained more in Appendix F). We also use CPS data to calculate the 16+ population in each state and year. Finally, we use state-level minimum wage data from [Vaghul and Zipperer \(2016\)](#) (the details are provided in Appendix E).

Figure D1: Example of aligning NAICS definition with SIC definition



Notes: This figure gives an example of the NAICS/SIC alignment process. We have data as per the SIC definition from 1980-2000, and as per the NAICS definition from 1990 onwards. The top panel plots restaurant employment in ever treated and never treated states as per both definitions, subject to data availability. The bottom panel imputes data for 1980 as per the NAICS definition by comparing the NAICS/SIC ratio in 1990 (the first year where we have overlapping data). The text in this section explains this process in more detail.

E Construction of 60 state-level minimum wage events

Identifying events

The event-study analysis reported in Table 1 uses 60 prominent state minimum wage increases, which are shown in Figure A1. We use the following process to identify these events:

1. We first classify as “initial events” all instances of a *state* minimum wage increase of at least \$0.25 and 5%.
2. Of these initial events, we retain those that occurred in years without a federal minimum wage increase as “admissible events.” This is because an event starting in a federal minimum wage year cannot have a “clean” post-period (see below for more details on clean controls). As a result, we exclude initial events in 1980, 1981, 1990, 1991, 1996, 1997, 2007, 2008, and 2009.
3. Among the admissible events, we classify an event as a “provisional combined event” (or the start of a potentially multi-year minimum wage increase) if no admissible event occurred in the preceding three years. This ensures that multi-year minimum wage increases are counted as a single event.
4. However, in some cases, this process may misidentify the true start year of a multi-year minimum wage increase. Specifically, if a multi-phase increase begins with a small increment, we might incorrectly conclude that the event started later than it actually did. To correct for this, we implement the following algorithm:
 - For all provisional combined events, we check if there was a minimum wage increase (large or small) in the preceding year, as long as the previous year was not a federal minimum wage increase year.
 - If there was such an increase, we then check whether the state had indexation at that time. If it did not, the increase is considered a legislated (albeit small) minimum wage increase, and we adjust the event start date to the previous year. (We use data on indexation from [Brummund and Strain \(2020\)](#).)
 - If the state did have indexation, we check if the small MW increase was a legislated increase. If it was indeed a legislated increase, we again move the event start date to the previous year.
 - We repeat this process until no combined event start years have a legislated non-federal-year minimum wage increase in the preceding year.
5. The preceding procedure identifies 64 events. However, since the maximum post-period is six years (as discussed below), there are five cases where an event falls within the post-period of an earlier event. These instances are: Alaska 2010 and 2015; Delaware 2014 and 2019; New Jersey 2014 and 2019; Oregon 1998 and 2003; and Rhode Island 1999 and 2004. We must decide whether to treat the latter year as part of the post-period for the earlier event, or to let the latter year be a standalone event. We choose the first option (removing the latter event) if the

minimum wage increase in the post-period of the latter event is no larger than the increase in the post-period of the former event. If the increase for the latter event is larger, we keep both events but reduce the post-period of the former event, so the latter event no longer overlaps with the former’s post-period. In practice, we choose the first option for Delaware, New Jersey, Oregon, and Rhode Island—i.e., the latter events are dropped and considered part of the earlier event. For Alaska, we keep both the 2010 and 2015 events; as a result, we restrict the post-period for Alaska 2010 to five years instead of six. This results in a final list of 60 combined events, which are plotted by state and time in the Figure [A1](#).

Post-periods and clean controls

To implement Equation [5](#), we need to calculate the differences between average outcomes over (up to) six years following the event and the outcome in the year prior to the event. In some cases, the post-period is shortened either by a federal minimum wage increase or the end of the sample period in 2019.

For example, 1998 events have a full six-year post-period (1998–2003). In contrast, 2005 events only have a two-year post-period (2005–2006, since 2007 includes a federal minimum wage increase), while 2016 events have a four-year post-period (2016–2019). Therefore, we group post-periods by event-year cohorts—every event in the same year has the same post-period.⁴⁸ Table [A1](#) summarizes the distribution of post-periods for each event-year cohort.

Post-periods also enter into the construction of clean controls. In a given year, a state is considered a clean control if it has *no* state MW increase (large or small) in the three preceding years, and *no* state MW increase in the post-period (as defined by that year’s event-year cohort). For instance, a clean control for an event in 1998 must have had no state minimum wage increase during the 1995–2003 period, whereas a clean control for an event in 2005 only needs to have no increases during the 2002–2006 period.⁴⁹

⁴⁸The one exception is Alaska in 2010, which has a five-year post-period, while the other 2010 event (Nevada) has a six-year post-period. This ensures that the 2016 Alaska event is not included in the evaluation of the 2010 Alaska event.

⁴⁹For Alaska 2010, while we reduce the post-period to five years instead of six, clean controls must still be “clean” for six years, like other events in the 2010 cohort.

F Construction of probability groups using demographic predictors

Here we provide a brief summary of the approach. The details are provided in [Cengiz et al. \(2022\)](#), which we follow to construct the probability groups.

We construct a model that predicts the likelihood of an individual being a minimum wage worker based on demographic variables, using data from the Current Population Survey over the years 1979-2019. For the minimum wage information, we select state-level events for which the given state did not experience significant minimum wage events in the past 20 quarters and which will experience a prominent change in the subsequent 12 quarters. In terms of demographic variables, the ones used in the prediction model are: age, education, sex, rural residency, marital status, race, Hispanic, and veteran status.

To find the best predictors we use only data from the 20 quarters preceding the minimum wage changes. The data is divided into two mutually exclusive samples: a training sample and a test sample. We apply a gradient boosted trees algorithm ([Friedman, 2001](#)) in the training sample, choosing the parameters using cross-validation and assessing model performance in the test sample. The trained model is then employed to estimate predicted probabilities for the comprehensive dataset. This procedure results in an individual-level dataset containing predicted probabilities of being a minimum wage worker, and encompassing all states and time periods over 1979-2019.

We then use that dataset to create three groups of individuals according to their predicted probabilities, defined as: (i) high probability – sample of top 10% of predicted probability distribution; (ii) high-recall – broader sample comprising 75% of all minimum wage workers which includes the high probability group; (iii) low probability – sample of individuals that are not contained in the high-recall group.