



Undergraduate School

Degree Program in Economics and Business (Major in Finance)

Course of Macroeconomics

Minimum Wage and Productivity: Evidence from Sectoral Transmission Channels

Giovanna Vallanti

SUPERVISOR

Filippo Campanella
(298181)

CANDIDATE

Academic Year 2025/2026

Abstract

This thesis examines whether minimum wage intensity is associated with labour productivity and whether this relationship varies across sectoral production structures. The analysis starts from the theoretical ambiguity of minimum wage effects. While the competitive labour market model interprets a binding minimum wage mainly as a cost shock, monopsony, search-and-matching models and efficiency wage theory suggest that wage floors may have more complex effects depending on labour-market frictions, firm wage-setting power and productivity-related adjustment margins. Building on this framework, the empirical analysis uses an OECD country-sector-year panel from 1995 to 2021. Labour productivity is measured as real value added per hour worked, while minimum wage exposure is captured by the logarithm of the ratio between the statutory minimum wage and the average wage. The main specifications include country, sector and year fixed effects, with standard errors clustered at the country level. The results do not provide evidence of a robust direct association between minimum wage intensity and labour productivity. However, they show that sectoral characteristics matter. The high-skilled labour share is positively associated with productivity in the simpler specifications, and capital intensity becomes strongly significant when interacted with minimum wage intensity. The most stable heterogeneous pattern is the negative interaction between minimum wage intensity and capital intensity, which remains significant also when the minimum wage is lagged by one period. A separate regression with R&D intensity as dependent variable does not show a significant direct relationship between minimum wages and innovation intensity. Overall, the findings suggest that minimum wage effects should not be interpreted as uniform across sectors. Rather, the relationship between minimum wages and productivity appears conditional on the productive structure of sectors, especially their degree of capital intensity.

Table of Contents

<i>Abstract</i>	<i>1</i>
<i>1. Introduction</i>	<i>1</i>
1.1 Productivity as the missing margin.....	1
1.2 Sectoral transmission channels.....	2
1.3 Contribution and structure.....	3
<i>2. Literature review.....</i>	<i>4</i>
2.1 Competitive labour market.....	4
2.2 Monopsony model.....	8
2.3 Search and matching models.....	12
2.4 Efficiency wage theory	15
2.5 Theoretical implications for the empirical analysis.....	18
<i>3. Empirical analysis</i>	<i>21</i>
3.1 Data and variable construction	21
3.2 Econometric methodology	22
3.3 Baseline specification.....	23
3.4 Interaction specifications	24
3.5 Extended specification and R&D outcome regression	26
3.6 Limitations and robustness considerations	28
<i>4. Discussion of results.....</i>	<i>31</i>
4.1 The absence of a robust direct minimum wage effect	31
4.2 Human capital and the limits of the skill-composition channel	33
4.3 Capital intensity as the main source of heterogeneity	34

4.4 Innovation and the separate R&D regression.....	35
4.5 Comparison with the theoretical models	36
4.6 Overall interpretation	37
5. Conclusion	39
5.1 Main findings	39
5.2 Interpretation in light of the theoretical models.....	41
5.3 Limitations and directions for future research.....	42
5.4 Final remarks	43
<i>Bibliography</i>	45
<i>Appendix</i>	47

1. Introduction

The minimum wage remains one of the most debated institutions in labour economics because it combines redistributive aims with possible efficiency effects. It protects workers at the bottom of the wage distribution, but it may also affect firms' costs, hiring decisions and production choices. Recent literature shows that the debate has moved beyond the traditional question of employment losses and increasingly considers broader margins of adjustment, including prices, profits, worker reallocation, capital, skills and productivity.¹

This broader perspective is necessary because the empirical evidence does not support a single mechanical prediction. Studies for the United States find substantial wage gains for low-paid workers with limited changes in aggregate low-wage employment, while evidence from Germany suggests that a national minimum wage increased wages and contributed to worker reallocation toward more productive establishments². These findings suggest that firms do not respond through employment alone, but may also reorganize production or adjust input composition.

1.1 Productivity as the missing margin

Productivity is therefore a crucial margin of analysis. If firms absorb higher wage costs through better organization, lower turnover, stronger effort, improved worker selection, training, technological upgrading or the exit of less productive producers, the effect of a minimum wage cannot be understood by looking only at employment. Evidence from the

¹ Dube Arindrajit & Lindner Attila S. (2024). Minimum Wages in the 21st Century. *National Bureau of Economics Research (NBER) Working Paper*.

² Cengiz Doruk, Dube Arindrajit et al. (2019). The effect of minimum wages on low-wage jobs: evidence from the United States using a bunching. *National Bureau of Economic Research*.

United Kingdom links the National Minimum Wage to higher firm productivity, with mechanisms consistent with organizational change, efficiency wages and training responses.³

This focus also connects the empirical analysis to the theoretical models discussed in Chapter 2. In a competitive labour market, a binding minimum wage mainly acts as a cost shock. By contrast, monopsony, search-based and efficiency-wage models allow for more nuanced outcomes, because wage floors may correct employer wage-setting power, improve effort or reduce turnover. The key implication is that the impact of minimum wages depends on labour-market structure and on the margins through which firms adjust.

1.2 Sectoral transmission channels

Against this background, the central question of this thesis is whether minimum wage intensity is associated with labour productivity and whether this relationship differs across sectoral characteristics. The empirical analysis uses a country-sector-year panel for OECD countries and models labour productivity as output per hour worked. The baseline specification is then extended through interaction terms between minimum wage intensity and three structural variables: high-skilled labour share, R&D intensity and capital intensity.

These interactions are important because the productivity response to minimum wages is unlikely to be homogeneous. Skill-intensive sectors may adjust differently from low-skill sectors; innovation-intensive sectors may have greater scope for technological or organizational responses; and capital-intensive sectors may rely less on labour as the main adjustment margin. For this reason, the analysis moves from estimating an average association to examining whether the minimum-wage-productivity relationship varies with production structure.

³ Riley Rebecca & Bondibene Chiara Rosazza (2015). Raising the standard: minimum wages and firm productivity. *National Institute of Economic and National Research (NIESR)*.

1.3 Contribution and structure

The contribution of this thesis is twofold. First, it shifts attention from the standard employment-centered debate to labour productivity, treating productivity as a central adjustment margin. Second, it tests whether the relationship between minimum wage intensity and productivity is moderated by workforce composition, innovation and capital intensity. This sectoral perspective is useful because the same wage policy may be absorbed differently across production structures and institutional settings.

The thesis is organized as follows. Chapter 2 reviews the main theoretical models of minimum-wage effects: competitive labour markets, monopsony, search and matching models, and efficiency-wage theory. Chapter 3 presents the empirical analysis, including the dataset, methodology and regression results. Chapter 4 connects the empirical findings to the theoretical mechanisms, while Chapter 5 concludes by summarizing the main findings, limitations and implications of the analysis.

2. Literature review

Minimum wage policies are among the most debated forms of labour market regulation because they directly affect the price of low-paid labour. From a theoretical perspective, their impact is not univocal. Different models lead to different predictions depending on how the labour market is assumed to work, how firms set wages, and how workers and firms interact. The competitive model provides the traditional benchmark. In this framework, a binding minimum wage is expected to reduce employment because it raises the cost of labour above the market-clearing level. However, this prediction depends on strong assumptions, especially perfect competition, full information, and the absence of frictions. Once these assumptions are relaxed, the effects of minimum wages become more complex. Models based on monopsony power, search frictions, and efficiency wages suggest that a wage floor may, under specific conditions, increase wages without necessarily producing large employment losses. In some cases, it may even improve employment or productivity by changing firms' incentives and workers' behaviour.

For this reason, the purpose of this chapter is not simply to present alternative theories of the minimum wage, but to clarify the mechanisms through which minimum wage increases may affect labour market outcomes. Particular attention is given to employment, wages, worker welfare, and the channels through which firms may adjust. This theoretical discussion is useful for the empirical part of the thesis, where the focus is not only on whether minimum wages affect labour market outcomes, but also on whether these effects differ across sectors and interact with productivity-related characteristics.

2.1 Competitive labour market

The competitive labour market model is the standard starting point for analysing minimum wage effects. Its relevance lies in the fact that it provides a clear benchmark prediction: when the wage is set above the equilibrium level, employment falls. However, this result follows from a specific set of assumptions that describe an idealised labour market rather than the full complexity of real-world labour relations.

In this framework, wages are determined by the interaction between labour supply and labour demand. Firms demand labour up to the point at which the wage equals the marginal product of labour, while workers supply labour according to their willingness to work at different wage levels. The equilibrium wage is therefore the wage at which the quantity of labour demanded by firms equals the quantity of labour supplied by workers.

The model relies on three main assumptions.

First, there are many employers and many workers, so no individual firm or worker has the power to influence the wage. All agents are therefore wage takers. Firms cannot set wages below the market rate because workers could move to other employers, and workers cannot demand wages above the market rate because firms could hire other workers.

Second, the market is assumed to be frictionless. Workers and firms have full information, job search is costless, and matches between vacancies and workers occur immediately. Under these assumptions, unemployment cannot persist in equilibrium, because any excess supply or demand for labour is eliminated through wage adjustment.

Third, wages are assumed to adjust freely until the market clears. If labour supply exceeds labour demand, wages fall; if labour demand exceeds labour supply, wages rise. The equilibrium wage is therefore the price that clears the labour market and ensures that all workers willing to work at that wage can find employment.

Within this framework, a minimum wage becomes economically relevant only when it is binding, meaning that it is set above the competitive equilibrium wage. If the minimum wage is equal to or below the equilibrium wage, it does not affect firms' hiring decisions or workers' labour supply. By contrast, when the minimum wage is above the equilibrium level, it prevents the wage from falling back to the market-clearing rate and creates a distortion.

The most direct predicted effect concerns employment. Since labour has become more expensive, firms reduce the quantity of labour demanded. In the standard textbook representation, if the competitive equilibrium level of employment were 1,000 workers, the introduction of a binding minimum wage could reduce employment to 800 workers. The


intuition is that some workers who were profitable to hire at the equilibrium wage are no longer profitable at the higher legal wage. As a result, firms hire fewer workers.

At the same time, the higher wage increases the number of individuals willing to work. This generates excess labour supply: more workers want a job at the minimum wage, but firms offer fewer jobs. The gap between labour supplied and labour demanded represents classical unemployment. In other words, some workers who would have accepted employment at the lower competitive wage are unable to find a job at the higher legal wage.

The effect on wages is therefore asymmetric. Workers who remain employed benefit from a higher hourly wage, while workers who lose their jobs or fail to enter employment receive no wage at all. The policy mainly affects workers whose previous wage was below the new minimum; workers already earning above the threshold are not directly affected. This distinction is important because the average wage of employed workers may rise even if total employment falls.

The welfare implications are also mixed. Employed workers who receive the higher wage are better off, but unemployed workers are worse off because they lose access to jobs they would have accepted at the competitive wage. The competitive model therefore predicts a redistribution among low-wage workers: some gain through higher earnings, while others lose employment opportunities. From the perspective of total surplus, the outcome is inefficient because fewer mutually beneficial labour transactions take place. This inefficiency is usually represented as a deadweight loss: some workers value employment at a wage below the minimum, and some firms would be willing to hire them at that wage, but the legal wage floor prevents the transaction from occurring.

Figure 1: Deadweight Loss from Minimum Wage



Note: When the minimum wage is set above W^* , labour demand falls to L_D and labour supply rises to L_S . The resulting gap $L_S - L_D$ represents unemployment, while the shaded area captures the associated deadweight loss.

Source: Economicshelp.org

Overall, the competitive model predicts that a binding minimum wage raises earnings for workers who keep their jobs, but reduces employment and creates involuntary unemployment for others. Its central implication is therefore a trade-off between higher wages and lower employment. This prediction is theoretically clear, but it depends on assumptions that may not always hold in actual labour markets. For this reason, later models relax the competitive benchmark and show that the effects of minimum wages can differ when firms possess wage-setting power, when job search is costly, or when wages influence productivity.

Under the classical assumptions, the minimum wage is therefore interpreted as a labour market distortion. Its main effect is redistributive: it increases the pay of workers who remain employed, but it may reduce employment opportunities for those whose productivity is below the new wage floor. The model consequently predicts lower employment than in an unregulated competitive equilibrium and the emergence of efficiency costs, usually

represented as a deadweight loss. This perspective has traditionally shaped the economic debate on minimum wages, providing the benchmark expectation that binding wage floors generate negative employment effects. At the same time, its conclusions depend on the restrictive assumptions of perfect competition, frictionless adjustment, and full wage flexibility, which motivates the analysis of alternative frameworks in the following sections.

2.2 Monopsony model

A different prediction emerges when the assumption of perfect competition is relaxed. In many labour markets, firms may have some degree of wage-setting power, especially when workers face limited outside options, mobility costs, or imperfect information about alternative jobs. The monopsony model captures this situation and provides an important alternative to the competitive benchmark.

In its simplest form, monopsony refers to a labour market in which a single employer, or a small number of dominant employers, hires a large share of workers in a given area or occupation. This can occur, for example, in geographically isolated labour markets, in sectors requiring specific skills, or in situations where workers cannot easily switch employer. The key implication is that the firm does not take the wage as given. Instead, it faces an upward-sloping labour supply curve: to attract more workers, it must offer a higher wage.

This feature changes the firm's hiring decision. In a competitive market, the wage paid to workers is also the marginal cost of hiring an additional worker. In a monopsonistic market, however, the marginal cost of labour is higher than the wage. The reason is that hiring one extra worker may require the firm to raise the wage not only for the new worker, but also for existing workers. As a result, the firm chooses employment where the marginal cost of labour equals the marginal revenue product of labour, and then pays the lowest wage necessary to attract that number of workers.

The monopsony outcome is therefore inefficient even before any policy intervention. Employment and wages are both below the competitive levels:

$$L_m < L^* \text{ and } W_m < W^*$$

The firm restricts employment in order to keep wages low. Workers are paid less than their marginal contribution to output, and some workers who would be employed in a competitive market are not hired. Differently from the competitive model, the initial distortion does not come from the minimum wage, but from the firm's labour market power.

This is why the introduction of a minimum wage can have a different effect under monopsony. If the minimum wage is set above the monopsony wage, W_m , but not above the competitive wage, W^* , it can increase both wages and employment. The wage floor limits the firm's ability to pay workers below their marginal product and reduces the gap between the wage and the marginal cost of labour.

The intuition is that, once a legal minimum wage is introduced, the firm can hire additional workers at the same mandated wage, at least up to a certain point. In this range, the marginal cost of hiring an extra worker becomes flatter than before, because the firm no longer needs to increase wages gradually in order to attract each additional worker. The firm will then expand employment until the marginal revenue product of labour equals the minimum wage.

If the minimum wage is set exactly at the competitive wage, the monopsonist may increase employment up to the competitive level, L^* . In this case, the minimum wage does not create unemployment. Instead, it corrects part of the initial monopsony distortion. Wages rise, employment rises, and part of the surplus previously captured by the firm is transferred to workers. This result is the main reason why the monopsony model produces a prediction opposite to the competitive model: a binding minimum wage can increase employment rather than reduce it.


However, this positive result does not hold for any level of the minimum wage. If the wage floor is set above the competitive wage, the firm's labour demand begins to fall, because some workers become more expensive than the value of their marginal product. At that point, the monopsony model becomes closer to the competitive case: labour supply exceeds labour demand, and unemployment may emerge.

Therefore, the relationship between the minimum wage and employment in the monopsony model is not necessarily negative. It depends on the level at which the wage floor is set. A

moderate minimum wage can increase employment by reducing monopsonistic under-hiring, while an excessively high minimum wage can reduce employment by raising labour costs beyond the productive value of some workers.

Overall, the monopsony model changes the interpretation of minimum wage policy. Rather than being automatically viewed as a distortion, the minimum wage can be understood as a corrective instrument when firms possess wage-setting power. Its effect depends on whether the policy moves the labour market closer to the competitive benchmark or beyond it. This framework is particularly relevant for empirical analysis because it suggests that minimum wage effects may vary across sectors, depending on the degree of employer concentration, worker mobility, and labour market frictions.

Figure 2 Minimum wage in a monopsonistic labour market



Note: A monopsonist hires Q_2 workers at wage W_2 , below the competitive outcome (Q_1, W_1). A minimum wage set at W_1 can increase both wages and employment by correcting monopsonistic under-hiring. If set above W_1 , however, it may reduce employment.

Source: Economicshelp.org.

From a welfare perspective, the monopsony model implies that a well-calibrated minimum wage can improve the position of low-paid workers. In the initial monopsony equilibrium, workers are harmed not only because wages are below the competitive level, but also because employment is inefficiently restricted. The firm hires fewer workers than would be employed under competition in order to keep wage costs low. This creates a deadweight loss even before the introduction of any minimum wage, since some employment relationships that would be mutually beneficial in a competitive market do not take place.

A minimum wage set between the monopsony wage, W_m and the competitive wage, W^* , can reduce this distortion. Workers who were already employed receive a higher wage, closer to their marginal contribution to production, while additional workers may be hired because the firm's marginal cost of labour is no longer above the wage over the relevant range. In this case, the policy improves both equity and efficiency: equity improves because income is redistributed from monopsony rents toward low-paid workers, while efficiency improves because employment moves closer to the competitive level.

This result contrasts sharply with the competitive model. In the competitive framework, the minimum wage creates a trade-off between higher wages for employed workers and lower employment for others. Under monopsony, instead, this trade-off may be absent as long as the wage floor corrects the firm's wage-setting power without exceeding the competitive wage. The policy can therefore raise wages without necessarily generating unemployment, because it offsets an existing market failure rather than creating a new one.

However, the positive welfare effect is not automatic. If the minimum wage is set above the competitive level, the firm may again reduce employment because some workers become more costly than their marginal product. In that case, the consequences become closer to those predicted by the competitive model, with job losses among lower-productivity workers and a possible decline in total welfare.

The main contribution of the monopsony model is therefore to show that the effects of minimum wages depend on the structure of the labour market. Where employers possess wage-setting power, a moderate minimum wage may increase wages, expand employment,

and reduce the inefficiency created by underemployment. This theoretical insight is particularly relevant for empirical analysis, because it suggests that the impact of minimum wage policies may differ across sectors depending on employer concentration, worker mobility, and the strength of labour market frictions.

2.3 Search and matching models

Search and matching models provide a further extension of the standard competitive framework by introducing frictions into the labour market. In these models, workers and firms do not meet instantaneously, information is imperfect, and finding a suitable job or filling a vacancy requires time and effort. As a result, unemployment can exist even in equilibrium, not because wages fail to adjust in a purely mechanical sense, but because the process of matching workers with firms is costly and imperfect.

This approach is associated with the work of Diamond, Mortensen, and Pissarides⁴, and is particularly useful for analysing minimum wages because it highlights the importance of job search, vacancy creation, and bargaining power. Differently from the competitive model, wages are not determined only by the intersection of labour supply and labour demand. They may instead result from bargaining between workers and firms, or from firms' wage-posting strategies. Therefore, the effect of a minimum wage depends on how it changes both workers' incentives to search and firms' incentives to create vacancies.

The basic assumptions of search and matching models can be summarised in three points.

First, there is frictional unemployment. At any moment, some workers are unemployed and searching for jobs, while some firms have unfilled vacancies. Matches are created gradually through a matching process that depends on the number of job seekers and vacancies. This

⁴ Diamond, Mortensen and Pissarides received the 2010 Nobel Prize in Economic Sciences for their analysis of markets with search frictions. The labour-market search-and-matching framework is commonly known as the Diamond–Mortensen–Pissarides model; among its key references are Pissarides (1985) and Mortensen and Pissarides (1994).

means that unemployment and vacancies can coexist even when the labour market is in equilibrium.

Second, wages are often determined through bargaining or wage posting. Since a job match creates a surplus for both the worker and the firm, the wage depends on how this surplus is divided. If workers have weak bargaining power, firms may be able to pay wages below the level that would prevail in a fully competitive market. In this sense, search frictions can generate a form of monopsony-like power even when there are many firms. Workers may accept lower wages because looking for another job is costly, uncertain, or time-consuming.

Third, workers' search effort and firms' vacancy creation are endogenous. A higher expected wage may encourage unemployed workers to search more actively or to enter the labour force. At the same time, a higher wage may reduce firms' incentives to open vacancies, because each job becomes less profitable. Minimum wage policy therefore affects both sides of the matching process: it can increase the attractiveness of employment for workers, but it can also increase labour costs for firms.

For this reason, the employment effect of a minimum wage is theoretically ambiguous. The final outcome depends on the balance between two opposite mechanisms: the search-intensity effect and the vacancy-reduction effect.

When the initial wage is very low, workers may have weak bargaining power and firms may capture a large share of the match surplus. In this case, a moderate minimum wage can improve labour market outcomes. By increasing the payoff from employment, the minimum wage can make job search more attractive. Unemployed workers may search more intensively, send more applications, or enter the labour force because the expected return from working is higher. If firms were initially paying wages well below workers' productivity, the increase in labour costs may not strongly reduce vacancy creation. As a result, the positive effect on search effort may dominate the negative effect on vacancies. Employment may then rise, or unemployment may fall, because more workers actively search and more successful matches are created.

This mechanism is close to the monopsony logic discussed above. Search frictions give firms some wage-setting power because workers cannot immediately move to alternative employers. A minimum wage can therefore push wages closer to workers' productivity and reduce the inefficiency created by excessively low wages. In this case, the policy may increase both wages and employment.

The opposite result may occur when the initial wage is already relatively high. If workers have strong bargaining power, or if wages are already close to the competitive level, a higher minimum wage does not correct an underpayment problem. Instead, it mainly increases labour costs for firms. Firms may respond by posting fewer vacancies, since each filled job generates lower expected profit. With fewer vacancies available, unemployed workers face a lower probability of finding a job, and some may reduce their search effort. In this case, the vacancy-reduction effect dominates, leading to lower employment and higher unemployment.

The search and matching framework therefore does not produce a single prediction. It suggests that the effect of a minimum wage depends on the initial structure of the labour market. If wages are initially too low because workers have limited outside options, a moderate wage floor may improve matching and increase employment. If wages are already high enough, the same policy may reduce vacancy creation and generate disemployment effects.

This also affects the interpretation of worker welfare. A higher minimum wage increases earnings for workers who remain employed or who find jobs after the policy change. If the policy also increases successful matches, worker welfare clearly improves because both wages and employment rise. However, if the minimum wage substantially reduces vacancies, unemployed workers may experience longer job-search durations and lower probabilities of employment. In that case, the gains for employed workers must be weighed against the losses suffered by those who remain unemployed.

Some search models suggest that welfare may increase even when unemployment rises slightly, provided that the wage gains for employed workers and the improvement in match

quality are large enough to offset the cost of additional unemployment. However, this result depends on the level of the minimum wage. If the wage floor is set too high, the reduction in vacancy creation becomes dominant and worker welfare may decline, especially for low-skilled workers who are more exposed to job loss.

Overall, search and matching models occupy an intermediate position between the competitive and monopsony frameworks. They show that labour markets may contain frictions that give firms some wage-setting power, even without a single dominant employer. A minimum wage can therefore either correct an existing inefficiency or create a new distortion, depending on workers' bargaining power, search incentives, and firms' vacancy decisions. This perspective is particularly relevant for empirical analysis because it helps explain why minimum wage effects may be small, heterogeneous, or even positive in some contexts, while negative in others.

2.4 Efficiency wage theory

Efficiency wage theory introduces a different mechanism through which wages may affect labour market outcomes. While the competitive model treats wages mainly as a cost for firms, efficiency wage models emphasise that wages can also influence worker productivity. In this framework, firms may voluntarily pay wages above the market-clearing level because higher wages can improve effort, reduce turnover, attract better applicants, or increase workers' attachment to the firm.

The central idea is that the wage is not only the price of labour, but also an instrument that affects the quality and intensity of labour supplied by workers. If paying a higher wage increases productivity sufficiently, the firm may find it profitable to pay more than the minimum wage required to hire workers. Therefore, the firm chooses the wage that minimises labour cost per effective unit of work, rather than simply the lowest possible wage.

Several mechanisms can explain this result.

First, firms may face imperfect monitoring. If workers' effort cannot be perfectly observed, employees may have an incentive to shirk. By paying a wage above outside alternatives, the

firm makes the job more valuable to workers. Losing the job then becomes more costly, which encourages employees to exert greater effort. This logic is captured by the Shapiro-Stiglitz “no-shirking condition”, according to which firms may need to pay an above-market wage to induce workers to supply full effort.

Second, higher wages may reduce turnover. When workers are better paid, they are less likely to quit, and firms can save on recruitment and training costs. In this case, a higher wage may be partly offset by lower turnover-related expenses. A related interpretation is the “gift exchange” mechanism, where workers respond to fair or generous wages with greater loyalty and effort.

Third, higher wages may help firms attract more productive workers. If applicants differ in ability and firms cannot perfectly observe worker quality in advance, offering higher wages can improve the pool of candidates. In low-income contexts, higher wages may also affect productivity through better nutrition, health, or stability, although this mechanism is more relevant in specific economic environments.

A key implication of efficiency wage theory is that the labour market may not clear. Since firms deliberately pay wages above the market-clearing level, some workers who would be willing to work at a lower wage remain unemployed. Unemployment is therefore not necessarily caused by minimum wage legislation. It can arise as an equilibrium outcome of firms’ own wage-setting decisions when productivity depends on wages.

The effect of a minimum wage in this framework depends on how the legal wage floor compares with the efficiency wage already chosen by firms.

If the minimum wage is below the wage that firms already pay for efficiency reasons, it is not binding. In this case, the policy has little or no direct effect on employment or wages, because firms would continue to pay the higher efficiency wage even without legal intervention. For example, if a firm finds that a certain wage is necessary to reduce shirking or turnover, a lower statutory minimum wage does not alter its behaviour.

The more relevant case occurs when the minimum wage forces some firms to raise wages above the level they would otherwise choose. The standard prediction is cautious. If the firm had already selected the wage that maximised effort or minimised cost per effective unit of labour, a higher legal wage may increase labour costs without generating proportional productivity gains. In that case, employment may fall because each worker becomes more expensive relative to the value produced. This resembles the competitive result, although the mechanism is different: the problem is not simply that the wage is above the market-clearing level, but that it may be above the firm's productivity-maximising wage.

However, efficiency wage theory also suggests that some productivity offsets may reduce the negative employment effect of a minimum wage. If higher wages induce workers to exert more effort, reduce quits, improve morale, or encourage firms to reorganise production more efficiently, the increase in labour costs may be partly absorbed by higher productivity. In this case, firms may reduce employment less than predicted by the competitive model.

Some theoretical contributions even show that, under specific conditions, a higher minimum wage may increase employment. For example, in shirking models, a higher wage floor can increase workers' effort and improve the functioning of the low-wage sector. The intuition is that the minimum wage may move the labour market from a low-wage, low-effort equilibrium toward a higher-wage, higher-effort equilibrium. Nevertheless, this is a special case rather than the general prediction of efficiency wage theory. A common objection is that, if higher wages were clearly profitable because of productivity gains, firms would often have an incentive to raise wages voluntarily even without a legal mandate.

The welfare implications are therefore mixed. Workers who remain employed benefit from higher earnings, and possibly from more stable employment relationships if turnover falls. Firms may also benefit from higher productivity or lower recruitment costs. However, if the wage floor pushes labour costs above the productivity gains generated by higher wages, firms may hire fewer workers. In that case, the policy creates the familiar trade-off between higher wages for employed workers and reduced opportunities for those who lose their jobs or cannot enter employment.

For this thesis, the main relevance of efficiency wage theory is that it introduces a direct link between minimum wages and productivity. A wage increase may not only affect employment through labour costs, but also through workers' effort, turnover, job stability, and firm organisation. This is important because the empirical analysis focuses on productivity-related outcomes and sectoral heterogeneity. Sectors with high turnover costs or stronger effort-related mechanisms may respond differently to minimum wage changes than sectors where wages have limited effects on productivity.

Overall, efficiency wage theory does not imply that minimum wages automatically increase employment. Rather, it suggests that the employment effect may be softened, neutral, or negative depending on whether productivity gains are large enough to offset higher wage costs. If the minimum wage is consistent with the wage level justified by productivity improvements, its employment effects may be limited and worker welfare may improve. If it exceeds that level, the policy may increase unemployment. As in the previous models, the theoretical prediction is therefore context-dependent and may vary across sectors.

2.5 Theoretical implications for the empirical analysis

The theoretical literature does not provide a single prediction on the effects of minimum wage policies. Rather, the expected impact depends on the assumptions made about the structure of the labour market, the behaviour of firms, and the way workers respond to wage changes. The minimum wage can therefore be interpreted in different ways: as a distortion in a competitive market, as a corrective instrument in a monopsonistic setting, or as a policy whose effects depend on search frictions, bargaining power, and productivity responses.

The competitive model offers the clearest benchmark. Under perfect competition, firms are wage takers, wages adjust freely, and the labour market clears through the interaction between labour supply and labour demand. In this framework, a binding minimum wage raises labour costs above the equilibrium level and reduces the quantity of labour demanded. The policy therefore creates a trade-off: workers who remain employed earn higher wages, while others may lose employment opportunities. The result is lower employment and a deadweight loss relative to the unregulated equilibrium.

However, this prediction becomes less straightforward once the assumption of perfect competition is relaxed. The monopsony model shows that, when firms have wage-setting power, the initial market outcome may already be inefficient. In that case, wages and employment are below their competitive levels because the employer restricts hiring in order to keep wages low. A moderate minimum wage can then improve outcomes by reducing monopsony power, increasing wages, and moving employment closer to the competitive level. The minimum wage is therefore not necessarily a distortion; under specific conditions, it can correct an existing market failure.

Search and matching models add a further layer of complexity by introducing frictions in the process through which workers and firms meet. In these models, unemployment can exist even in equilibrium because job search takes time and vacancies are costly to create. A minimum wage affects both sides of the market. On the one hand, higher wages may encourage workers to search more actively or participate in the labour market. On the other hand, they may reduce firms' incentives to post vacancies. The final employment effect depends on which mechanism dominates. If wages are initially very low and workers have weak bargaining power, a minimum wage may improve matching and reduce unemployment. If wages are already close to productivity levels, the same policy may reduce vacancy creation and increase unemployment.

Efficiency wage theory shifts the focus from employment alone to the relationship between wages and productivity. In this framework, firms may pay above-market wages because higher wages can increase effort, reduce turnover, attract better workers, or improve job stability. This perspective is particularly relevant for the present thesis because it suggests that minimum wages may affect productivity-related outcomes, not only employment. Nevertheless, the theory does not imply that higher minimum wages automatically increase employment. If the wage floor remains consistent with productivity gains, employment effects may be limited. If it exceeds the level justified by productivity improvements, firms may reduce hiring.

Taken together, these models show that the effects of minimum wages are theoretically ambiguous. A binding minimum wage can reduce employment, have little effect, or even

increase employment depending on the degree of competition, the presence of employer market power, the intensity of search frictions, and the extent to which higher wages translate into productivity gains. This ambiguity is important for the empirical analysis of the thesis. It suggests that minimum wage effects should not be expected to be uniform across sectors. Sectors differ in labour intensity, skill composition, capital intensity, turnover costs, and productivity dynamics. As a result, the same minimum wage policy may generate different adjustment mechanisms across different parts of the economy.

For this reason, the theoretical chapter provides the basis for the empirical strategy developed in the following section. The aim of the empirical analysis is not only to test whether minimum wages affect labour productivity, but also to examine whether these effects vary across sectors and depend on structural characteristics. The models reviewed in this chapter therefore support the idea that the relationship between minimum wages and productivity should be analysed as a context-dependent process rather than as a single mechanical effect.

3. Empirical analysis

3.1 Data and variable construction

The empirical analysis is based on a final country–sector–year dataset constructed for the regression analysis. The dataset contains the country identifier, the NACE Rev.2 sector code, the year, the dependent variables and the explanatory ones used in the econometric specifications. The final unit of observation is the country–sector–year, which makes it possible to combine cross-country variation, sectoral heterogeneity and time variation in labour productivity.

Labour productivity is constructed from EU KLEMS data and defined as real value added per total hours worked. Such logarithmic variable captures differences in labour productivity across country–sector units and over time while reducing the skewness of the original productivity measure.

To control for differences in labour quality, the analysis includes the share of high-skilled workers, obtained by dividing the percentage share of high-skilled workers by 100, so that it ranges on a 0–1 scale, making the econometric interpretation more standard and directly comparable to the other intensity variables.

Capital intensity is measured using the EU KLEMS variable referring to the all-assets capital stock at the country–sector–year level. Rather than using the absolute value, it has been decided to construct a capital intensity measure by dividing it by total hours worked. This transformation captures the amount of productive capital available per unit of labour input and therefore allows comparisons across sectors and countries of different sizes. In the empirical specification, this variable controls for differences in capital deepening across country–sector units and over time, since more capital-intensive sectors may display systematically higher labour productivity.

Innovation intensity is captured using the EU KLEMS variable measuring the research and development capital stock at the country–sector–year level. This variable reflects the

accumulated stock of R&D investment rather than current R&D expenditure, and can therefore be interpreted as a proxy for the existing innovative capacity of each sector. To make the variable comparable across sectors and countries, R&D intensity has been constructed by dividing it by total hours worked.

Finally, minimum wage exposure is measured using OECD data as the ratio between the statutory minimum wage and the average wage of full-time workers. The logarithm of this relative measure captures the bindingness of the minimum wage rather than its nominal level, making it more comparable across countries and over time. Since the minimum wage variable is observed at the country–year level, the same minimum wage intensity is assigned to all sectors within a given country and year.

3.2 Econometric methodology

The empirical strategy is organized around two related blocks of regressions. The first block studies the association between minimum wage intensity and labour productivity. In this part of the analysis, the dependent variable is the logarithm of labour productivity, while the main explanatory variable is the logarithm of minimum wage intensity. The use of the logarithmic transformation makes the coefficient easier to interpret in proportional terms and reduces the influence of scale differences in the minimum wage measure.

The productivity analysis is organized around four specifications. The first is the baseline model, where log labour productivity is regressed on log minimum wage intensity, the high-skilled labour share and capital intensity. The second specification adds an interaction between log minimum wage intensity and the high-skilled share, in order to test whether the minimum wage-productivity relationship differs according to workforce composition. The third specification instead interacts log minimum wage intensity with capital intensity, testing whether the relationship varies with the capital intensity of production. Finally, the extended specification includes both interaction terms simultaneously.

Instead, R&D intensity is treated as a separate dependent variable in an additional specification. This allows the analysis to examine whether minimum wage intensity is

associated with innovation intensity itself, rather than treating R&D only as a channel inside the productivity equation.

The empirical analysis relies on a single final dataset in which the minimum wage measure is merged onto the sectoral dataset by country and year. Since minimum wage intensity is observed at the country-year level, the same country-level value is assigned to all sectors within a country in a given year. The final dataset is an unbalanced country-sector panel observed annually from 1995 to 2021.

All main specifications include country fixed effects, sector fixed effects and year fixed effects separately. Country fixed effects absorb time-invariant differences across countries, such as institutional characteristics, persistent labour market structures and long-run macroeconomic conditions. Sector fixed effects control for structural differences across industries, such as production technologies, sectoral organization and average productivity levels. Year fixed effects absorb common shocks affecting all countries and sectors in a given year. Standard errors are clustered at the country level, allowing for arbitrary correlation of the error term within countries over time and across sectors.

3.3 Baseline specification

The baseline specification can be written as follows:

$$\ln(Y)_{cst} = \beta_1 \ln(MW_{ct}) + \beta_2 HighSkill_{cst} + \beta_3 K_{cst} + \gamma_c + \delta_s + \lambda_t + \epsilon_{cst}$$

Where:

- $\ln(Y)_{cst}$ denotes log labour productivity in country c , sector s and year t ;
- $\ln(MW_{ct})$ denotes log minimum wage intensity at the country–year level;
- $HighSkill_{cst}$ denotes the share of high-skilled workers;
- K_{cst} denotes capital intensity;
- The terms γ_c , δ_s and λ_t denote country, sector and year fixed effects, respectively;
- ϵ_{cst} captures the standard errors clustered at the country level to allow for arbitrary correlation of the error term within countries over time and across sectors.

The baseline results are reported in the first column of Table A1. The coefficient on log minimum wage is positive, but small and statistically insignificant. This suggests that, once persistent country-level differences, sector-level characteristics and common yearly shocks are controlled for, there is no robust evidence of a direct average association between minimum wage intensity and log labour productivity.

The control variables provide more systematic information. The coefficient on the high-skilled share is positive and statistically significant. This indicates that country-sector observations with a larger share of high-skilled workers tend to display higher labour productivity, conditional on country, sector and year fixed effects. The result is consistent with the idea that labour quality and human capital are important correlates of productivity differences across sectors and countries.

Capital intensity is positive in the baseline model, but it is not statistically significant in this simpler specification. This does not imply that capital is irrelevant for productivity. Rather, it suggests that the independent association between capital intensity and labour productivity is less clearly identified when the model does not allow the minimum wage-productivity relationship to vary with the degree of capital intensity.

3.4 Interaction specifications

The interaction specifications extend the baseline model by allowing the relationship between log minimum wage intensity and labour productivity to vary across country-sector characteristics. The objective is not to estimate a single homogeneous effect of the minimum wage, but to examine whether the association differs according to workforce composition and capital intensity. Each specification includes country, sector and year fixed effects separately, with standard errors clustered at the country level.

The first interaction specification examines whether the association between minimum wage intensity and labour productivity depends on the skill composition of the workforce. The model introduces an interaction between log minimum wage intensity and the high-skilled share:

$$\ln(Y)_{cst} = \beta_1 \ln(MW_{ct}) + \beta_2 HighSkill_{cst} + \beta_3 (\ln(MW_{ct}) \times HighSkill_{cst}) + \beta_4 K_{cst} + \gamma_c + \delta_s + \lambda_t + \epsilon_{cst}$$

As Table A1 suggests, the results do not provide evidence of a statistically significant moderating role of the high-skilled share. The interaction between log minimum wage intensity and high-skilled labour share is close to zero and statistically insignificant. Moreover, the high-skilled share itself loses statistical significance in this specification, which is likely related to the inclusion of the interaction term and the resulting increase in standard errors.

Economically, this means that workforce composition remains important for productivity in the baseline and other specifications, but the data do not show that high-skilled sectors experience a systematically different productivity association with minimum wage intensity. In other words, sectors with a higher share of skilled workers tend to be more productive, but they do not appear to react differently to minimum wage intensity in a statistically robust way.

The second interaction specification tests whether the minimum wage-productivity relationship varies with capital intensity. The model is specified as follows:

$$\ln(Y)_{cst} = \beta_1 \ln(MW_{ct}) + \beta_2 HighSkill_{cst} + \beta_3 K_{cst} + \beta_4 (\ln(MW_{ct}) \times K_{cst}) + \gamma_c + \delta_s + \lambda_t + \epsilon_{cst}$$

In this case, the results are more informative. Capital intensity becomes positive and statistically significant, while the interaction between log minimum wage intensity and capital intensity is negative and statistically significant. This suggests that capital-intensive sectors are associated with higher labour productivity, but that the marginal association between minimum wage intensity and productivity is weaker when capital intensity is higher.

A possible interpretation is that, in more capital-intensive sectors, productivity is already strongly linked to installed capital, technology and production equipment. Therefore, labour may represent a less flexible margin of adjustment compared with more labor-intensive sectors. Under this interpretation, minimum wage increases may have less scope to generate

productivity improvements through effort, organization or workforce adjustment in sectors where production depends more heavily on capital. However, the magnitude of the interaction coefficient is small, so the result should be interpreted cautiously as evidence of statistically detectable heterogeneity rather than as evidence of a large economic effect.

Taken together, the interaction specifications suggest that the direct coefficient on log minimum wage intensity is not the central result. It remains statistically insignificant across models. The main source of heterogeneity comes from the interaction with capital intensity, while the interaction with high-skilled labour share does not provide statistically robust evidence of moderation.

3.5 Extended specification and R&D outcome regression

The extended productivity specification includes both interaction terms simultaneously. This allows the moderating role of workforce composition and capital intensity to be assessed jointly, rather than considering each interaction separately. The model can be written as follows:

$$\ln(Y)_{cst} = \beta_1 \ln(MW_{ct}) + \beta_2 HighSkill_{cst} + \beta_3 K_{cst} + \beta_4 (\ln(MW_{ct}) \times HighSkill_{cst}) \\ + \beta_5 (\ln(MW_{ct}) \times K_{cst}) + \gamma_c + \delta_s + \lambda_t + \epsilon_{cst}$$

As shown in the extended column of Table 1, the coefficient on log minimum wage remains positive but statistically insignificant. This confirms that the analysis does not identify a robust direct average association between minimum wage intensity and labour productivity once fixed effects, controls and interaction terms are included.

The high-skilled share is no longer statistically significant in the extended specification, although its coefficient remains positive. This should be interpreted with caution, because the inclusion of the interaction between log minimum wage and high-skilled share substantially increases the standard error. By contrast, capital intensity remains positive and statistically significant, confirming that more capital-intensive country-sector observations tend to display higher labour productivity when the model allows for heterogeneous minimum wage effects.

Among the interaction terms, only the interaction between log minimum wage intensity and capital intensity is statistically significant. Its coefficient is negative and significant at the 1 percent level. This result is consistent with the previous interaction specification and suggests that the association between minimum wage intensity and labour productivity becomes less favorable in more capital-intensive sectors. The interaction with the high-skilled share remains statistically insignificant, suggesting that there is no robust evidence that the minimum wage-productivity relationship differs according to skill composition once the two interaction channels are considered jointly.

Overall, the extended productivity model indicates that the direct coefficient of log minimum wage intensity is not the main empirical finding. Instead, the evidence points to sectoral structure, especially capital intensity, as the relevant source of heterogeneity. High-skilled labour is positively associated with productivity in simpler specifications, while capital intensity becomes strongly significant once its interaction with log minimum wage is included.

The second part of the empirical analysis treats R&D intensity as a dependent variable. This additional specification examines whether minimum wage intensity is associated with innovation intensity itself. The extended R&D specification can be written as follows:

$$\begin{aligned} \text{R\&D}_{cst} = & \beta_1 \ln(MW_{ct}) + \beta_2 \text{HighSkill}_{cst} + \beta_3 K_{cst} + \beta_4 (\ln(MW_{ct}) \times \text{HighSkill}_{cst}) \\ & + \beta_5 (\ln(MW_{ct}) \times K_{cst}) + \gamma_c + \delta_s + \lambda_t + \epsilon_{cst} \end{aligned}$$

The results of this regression are reported in Table A2. The coefficient on log minimum wage is negative but statistically insignificant, indicating that there is no robust evidence of a direct association between minimum wage intensity and R&D intensity. The high-skilled share is also not statistically significant in this specification.

Capital intensity, instead, is positive and statistically significant. This suggests that more capital-intensive sectors also tend to display higher R&D intensity, which is consistent with the idea that physical and intangible investment may be complementary in technologically advanced sectors. However, the interaction between log minimum wage intensity and capital intensity is negative and statistically significant. This mirrors the pattern found in the

productivity regressions: in more capital-intensive sectors, the association between minimum wage intensity and the outcome variable becomes weaker or more negative.

Therefore, the separate R&D regression does not support the idea that minimum wage intensity directly increases innovation intensity. Rather, it suggests that R&D intensity is more strongly associated with capital intensity, while the role of minimum wage intensity remains indirect and heterogeneous. As with the productivity regressions, this result should be interpreted cautiously and not as definitive causal evidence.

3.6 Limitations and robustness considerations

While the empirical analysis provides useful insights into the relationship between minimum wages, productivity and innovation intensity, several limitations should be acknowledged.

A first important limitation concerns the potential endogeneity of minimum wage policies. Minimum wages are not randomly assigned across countries and over time, but are often determined in response to economic conditions, labour market performance and political factors. As a result, the estimated coefficients may capture not only the effect of minimum wage intensity, but also reverse causality or omitted variable bias. For instance, countries experiencing stronger productivity growth or broader structural changes may also be more likely to adjust minimum wage policies. Although the inclusion of country, sector and year fixed effects helps mitigate this issue by controlling for time-invariant country characteristics, time-invariant sector characteristics and common yearly shocks, it does not fully eliminate concerns related to time-varying unobserved factors.

A second limitation relates to the measurement of minimum wage intensity. The analysis uses a relative measure, defined as the ratio between the statutory minimum wage and the average wage, and then applies a logarithmic transformation. This improves comparability across countries and makes the interpretation more proportional, but it may not fully capture the effective bindingness of minimum wage policies. In particular, the measure may be influenced by changes in the wage distribution, it does not account for compliance or

enforcement differences, and it ignores institutional features such as collective bargaining systems.

In addition, although the analysis exploits sector-level variation, the use of relatively aggregated sectors based on the NACE Rev. 2 classification may mask important within-sector heterogeneity. Firms within the same sector may differ substantially in terms of productivity, technology, capital intensity, innovation activity and workforce composition. Consequently, the estimated coefficients should be interpreted as average sector-level relationships, rather than firm-level responses.

The interaction terms included in the empirical model allow the analysis to examine whether the relationship between minimum wage intensity and labour productivity differs depending on workforce composition and capital intensity. However, these variables should not be interpreted as fully identifying causal channels. Rather, they provide evidence on whether the association between minimum wages and productivity varies with structural characteristics of sectors. Firms may also respond to minimum wage changes through other margins not directly captured in the model, such as employment adjustments, price changes, changes in working hours, contract types, managerial practices or reallocation across firms.

Another limitation concerns the timing of adjustment. The main empirical specification focuses on contemporaneous relationships between log minimum wage intensity and labour productivity. However, some adjustment mechanisms, especially those related to investment, innovation and workforce restructuring, may require time to materialize. Firms may need several periods to adjust production processes, reorganize tasks, or change their capital and skill composition.

For this reason, the same productivity specifications are also estimated using the log minimum wage variable lagged by one period. The results are reported in Table A3 and can be compared with the contemporaneous specifications in Table A1. Overall, the lagged estimates are broadly consistent with the main results. The direct coefficient on lagged minimum wage remains positive but statistically insignificant across specifications,

suggesting that there is no robust evidence of an average direct association between minimum wages and labour productivity, either contemporaneously or with a one-year lag.

The robustness check also confirms the main pattern of heterogeneity. The interaction between lagged minimum wage and high-skilled share remains statistically insignificant, while the interaction between lagged minimum wage and capital intensity remains negative and statistically significant. Capital intensity is again positive and significant in the specifications that include the capital interaction. This suggests that the main result identified in the contemporaneous model is not driven only by the timing of the minimum wage variable.

Finally, the results are based on a sample of OECD countries and may not generalize to economies with different institutional frameworks, labour market structures or levels of development. In particular, the effects of minimum wages may differ substantially in countries with weaker enforcement mechanisms, different collective bargaining systems or larger informal sectors.

Overall, despite these limitations, the empirical strategy provides a consistent framework to analyze the sectoral relationship between minimum wage intensity, labour productivity and innovation intensity. The evidence does not support a robust direct effect of minimum wages on productivity. Instead, the results suggest that productivity differences are more closely related to structural sectoral characteristics, especially skill composition and capital intensity, and that the most stable form of heterogeneity is linked to capital intensity.

4. Discussion of results

The purpose of this chapter is to connect the empirical findings presented in Chapter 3 with the theoretical mechanisms discussed in Chapter 2. The literature review showed that minimum wage effects cannot be reduced to a single prediction. In the competitive model, a binding minimum wage mainly operates as a cost shock; in monopsony and search-based models, it may correct wage-setting power or labour-market frictions; in efficiency wage theory, it may affect productivity through effort, turnover and organizational responses. The empirical results must therefore be interpreted not only in terms of whether the coefficient on minimum wage intensity is positive or negative, but also in terms of whether the evidence supports a homogeneous effect or a more heterogeneous, sector-dependent mechanism.

The main finding of the empirical analysis is that log minimum wage intensity does not display a statistically significant direct association with labour productivity. This result holds both in the contemporaneous specifications and in the robustness check using the lagged minimum wage variable. At the same time, the regressions reveal systematic patterns in the role of sectoral characteristics. The high-skilled share is positively associated with labour productivity in the simpler productivity specifications, while capital intensity becomes strongly significant once the interaction between log minimum wage and capital intensity is included. The only interaction term that remains statistically significant across the main and lagged productivity specifications is the interaction between minimum wage intensity and capital intensity, which is negative. Finally, the separate R&D regression does not show a direct significant association between minimum wage intensity and innovation intensity, although capital intensity is positively related to R&D intensity.

4.1 The absence of a robust direct minimum wage effect

The first point to discuss is the absence of a robust direct relationship between minimum wage intensity and labour productivity. In the baseline productivity regression, the coefficient on log minimum wage is positive but statistically insignificant. The same pattern remains in the interaction models and in the extended specification. The lagged specifications

lead to a similar conclusion: even when the minimum wage is measured with a one-year delay, the direct coefficient remains statistically insignificant.

This result creates a first point of distance from a simple competitive interpretation. In the competitive labour market model, a binding minimum wage is expected to raise labour costs and reduce labour demand. If this mechanism translated mechanically into productivity outcomes, one might expect a clear negative association between minimum wage intensity and productivity, especially in sectors where firms are less able to absorb higher wage costs. The empirical results do not show such a strong negative average relationship. The minimum wage coefficient is not negative and significant; rather, it is positive but statistically indistinguishable from zero. Therefore, the data do not provide evidence that higher minimum wage intensity is associated with a systematic productivity loss at the average sectoral level.

At the same time, the findings also do not provide strong support for a broad efficiency-wage interpretation. Efficiency wage theory suggests that higher wages may raise productivity by increasing effort, reducing turnover, improving morale or encouraging firms to reorganize production. If this mechanism were dominant and homogeneous across sectors, one might expect a positive and statistically significant coefficient on minimum wage intensity. This is not what the regressions show. The absence of a significant direct coefficient therefore suggests that minimum wages do not operate as a simple productivity-enhancing shock across the whole sample.

The result is instead more consistent with the general message of the non-competitive models reviewed in Chapter 2: minimum wage effects are likely to be context-dependent. Monopsony, search and matching models, and efficiency wage theory all imply that outcomes depend on the structure of the labour market, the degree of frictions, the strength of worker bargaining power, and the margins along which firms adjust. The empirical evidence does not identify a single average effect, but it does point toward heterogeneity across sectoral production structures.

4.2 Human capital and the limits of the skill-composition channel

The second important result concerns the high-skilled share. In the baseline specification and in the model including the interaction with capital intensity, the coefficient on the high-skilled share is positive and statistically significant. This indicates that country-sector observations with a larger proportion of high-skilled workers tend to display higher labour productivity. This result is not surprising from an economic perspective, since human capital is closely related to workers' ability to perform complex tasks, adopt new technologies and contribute to more efficient production processes.

This finding is broadly consistent with the theoretical discussion, but only in a limited sense. Efficiency wage theory and search-based models both recognize that worker quality, effort and match quality matter for labour-market outcomes. A more skilled workforce may allow firms to adjust more effectively to cost shocks, adopt better organizational practices or use technologies more productively. Therefore, the positive coefficient on high-skilled share confirms that labour quality is an important productivity-related characteristic.

However, the interaction between log minimum wage intensity and high-skilled share is not statistically significant. This is a crucial distinction. The results suggest that high-skilled sectors are more productive on average, but they do not show that the relationship between minimum wage intensity and productivity is systematically different in sectors with a larger share of skilled workers. In other words, human capital matters for productivity, but the empirical analysis does not identify it as a robust moderating channel for minimum wage effects.

This creates a point of partial distance from an interpretation based on skill-biased adjustment. One could expect firms facing higher wage floors to respond by upgrading their workforce or relying more heavily on skilled labour. However, the interaction results do not provide strong evidence for this mechanism. The high-skilled share remains relevant as a structural determinant of productivity, but not as a clear channel through which minimum wage intensity changes productivity outcomes.

4.3 Capital intensity as the main source of heterogeneity

The most stable form of heterogeneity in the empirical analysis concerns capital intensity. In the specifications that include the interaction between log minimum wage intensity and capital intensity, capital intensity is positive and statistically significant. This means that more capital-intensive country-sector observations tend to be associated with higher labour productivity. This is consistent with standard productivity theory: when more capital is available per unit of labour input, workers can produce more output per hour.

The more interesting result is the negative and statistically significant interaction between log minimum wage intensity and capital intensity. This result appears both in the contemporaneous specifications and in the robustness check using lagged minimum wage. Therefore, it is the empirical pattern that remains most stable across alternative timing assumptions. The interpretation, however, must be cautious. The coefficient is statistically significant, but its magnitude is small. It should therefore not be read as evidence of a large negative economic impact. Rather, it suggests that the association between minimum wage intensity and labour productivity is less favorable, or weaker, in sectors where production is more capital-intensive.

This result can be linked to the theoretical models in several ways. In the competitive model, a minimum wage is primarily a cost shock. If firms cannot easily reorganize production or increase productivity, the wage floor may reduce profitability or employment. In highly capital-intensive sectors, however, production is already strongly shaped by installed capital, machinery, technology and fixed production processes. Labour may be a less flexible margin of adjustment than in labor-intensive sectors. Therefore, higher minimum wage intensity may have less scope to produce productivity improvements through effort, turnover reduction or work reorganization.

The result also speaks to efficiency wage theory. The productivity-enhancing mechanisms described in that theory operate mainly through workers: effort, discipline, loyalty, reduced quits and better worker selection. These mechanisms are likely to matter most when labour is central to the production process. In sectors where output depends heavily on capital

equipment and technological infrastructure, the link between wages and productivity-enhancing worker responses may be weaker. The negative interaction with capital intensity is therefore consistent with the idea that minimum wage-related productivity gains, if present, are less likely to appear in sectors where capital rather than labour is the dominant production input.

At the same time, the result should not be interpreted as a full confirmation of the competitive model. The competitive model would suggest a more mechanical negative effect of higher wage floors, while the regressions do not show a significant negative direct coefficient on minimum wage intensity. The evidence is better described as heterogeneous: the average effect is not statistically clear, but the association becomes less favorable as capital intensity rises. This supports the broader theoretical view that minimum wage effects depend on sectoral conditions rather than operating uniformly across the economy.

4.4 Innovation and the separate R&D regression

A further issue concerns innovation. In the final empirical specification, R&D intensity is not included as a control inside the productivity regression, but is instead treated as a separate dependent variable. This choice allows the analysis to ask whether minimum wage intensity is associated with innovation intensity itself. The results do not provide evidence of a direct minimum wage effect on R&D intensity. The coefficient on log minimum wage is negative but statistically insignificant, while the high-skilled share is also not statistically significant.

This finding creates a point of distance from a strong induced-innovation interpretation. If higher minimum wages systematically encouraged firms or sectors to invest more in innovation, one might expect a positive and significant association between minimum wage intensity and R&D intensity. The empirical results do not show this. Therefore, the analysis does not support the idea that minimum wage intensity directly stimulates R&D investment at the country-sector level.

However, the R&D regression also reveals that capital intensity is positively and significantly associated with R&D intensity. This suggests that physical capital and innovation-related

capital may be complementary. Sectors with greater capital intensity may also be sectors with more advanced technologies, more complex production processes and greater accumulated innovative capacity. In this sense, innovation appears more closely connected to the broader production structure of sectors than to minimum wage intensity itself.

The negative and significant interaction between log minimum wage intensity and capital intensity in the R&D regression mirrors the pattern found in the productivity regressions. Again, this does not imply a large causal effect, but it suggests that the relationship between minimum wage intensity and sectoral outcomes becomes less favorable in more capital-intensive environments. This reinforces the interpretation that capital intensity is the main dimension along which the empirical relationship varies.

4.5 Comparison with the theoretical models

The empirical results do not fit perfectly into any single theoretical model reviewed in Chapter 2. This is not necessarily a weakness. On the contrary, it reflects the central lesson of the theoretical chapter: minimum wage effects are ambiguous and depend on assumptions about market structure, frictions and firm behavior.

The results are not strongly aligned with the simple competitive model. The competitive framework predicts that a binding minimum wage acts mainly as a distortion by increasing labour costs above the equilibrium level. If that model fully described the empirical setting, one would expect clearer evidence of negative effects. Instead, the direct coefficient on log minimum wage is not statistically significant. This does not disprove the competitive mechanism, but it suggests that the relationship between minimum wages and productivity cannot be reduced to a simple cost-shock story.

The results are also not a direct confirmation of the monopsony model. The monopsony framework predicts that a moderate minimum wage may improve outcomes by correcting employer wage-setting power. However, the empirical analysis does not directly measure employer concentration, worker mobility or monopsony power. Moreover, it does not find a positive average effect of minimum wage intensity on productivity. Therefore, the results

cannot be interpreted as direct evidence that minimum wages are correcting monopsonistic underpayment. Nevertheless, the absence of a strong negative average effect is compatible with the idea that labour markets may not behave like perfectly competitive markets.

Search and matching models provide a useful intermediate interpretation. These models predict that minimum wage effects depend on the balance between different forces: higher wages may improve search incentives and match quality, but they may also reduce vacancy creation. Although the empirical analysis does not estimate search frictions directly, the lack of a uniform effect and the presence of heterogeneity are consistent with the logic of these models. Minimum wage intensity does not have one stable average association with productivity; its relationship with outcomes changes according to sectoral characteristics.

Efficiency wage theory is particularly relevant because it explicitly links wages to productivity. The empirical results, however, only partially align with this perspective. On the one hand, the focus on productivity and the positive role of high-skilled labour are consistent with the idea that workers' characteristics and productive capacity matter. On the other hand, the absence of a significant positive direct coefficient on minimum wage intensity weakens a strong efficiency-wage interpretation. The results do not show that higher minimum wage intensity generally raises productivity. Instead, any productivity-related effect appears conditional on the sectoral structure, especially capital intensity.

4.6 Overall interpretation

Overall, the empirical findings suggest that the relationship between minimum wage intensity and productivity is better understood as conditional rather than uniform. The direct coefficient on log minimum wage is not statistically significant, either contemporaneously or with a one-year lag. This means that the analysis does not find strong evidence that minimum wage intensity directly raises or lowers labour productivity on average.

At the same time, the results do not imply that minimum wage intensity is irrelevant. Rather, they suggest that its relationship with productivity depends on the productive environment in which firms operate. High-skilled labour is positively associated with productivity in the

simpler specifications, confirming the importance of labour quality. Capital intensity is also strongly connected to productivity when interaction terms are included. Most importantly, the negative interaction between minimum wage intensity and capital intensity indicates that the association between minimum wages and productivity becomes less favorable in more capital-intensive sectors.

The main point of contact between the empirical findings and Chapter 2 is therefore the idea of heterogeneity. The theoretical models show that minimum wage effects depend on market structure, frictions and firm adjustment margins. The empirical analysis supports this general view, because it does not identify a single average effect but instead finds that sectoral structure matters. The main point of distance is that the evidence does not clearly support one specific mechanism. It does not confirm the competitive model's strong negative prediction, nor does it clearly confirm a positive monopsony or efficiency-wage effect. The results are more modest but also more realistic: minimum wage intensity appears to have no robust direct association with productivity, while sectoral production characteristics, especially capital intensity, shape the way this relationship appears in the data.

This interpretation is important for the broader argument of the thesis. Minimum wage policies should not be evaluated only through average effects or through a single theoretical model. Their consequences may depend on whether sectors are labor-intensive or capital-intensive, on the skill composition of the workforce, and on the scope firms have to reorganize production. In this sense, the empirical results reinforce the need to analyze minimum wage policies as context-dependent interventions whose effects vary across the productive structure of the economy.

5. Conclusion

This thesis has analyzed the relationship between minimum wage intensity and labour productivity by combining a theoretical discussion of labour-market models with an empirical analysis based on OECD country-sector-year data. The starting point of the thesis was the observation that the minimum wage debate has progressively moved beyond the traditional focus on employment effects. While employment remains a central outcome, firms may also adjust to higher wage floors through productivity, capital deepening, workforce composition, innovation, prices, reallocation or organizational change. For this reason, the thesis focused on labour productivity as a key margin through which minimum wage policies may be transmitted across sectors.

The theoretical chapter showed that there is no single prediction about the effects of minimum wages. In the competitive labour market model, a binding minimum wage is primarily interpreted as a distortion that raises labour costs above the market-clearing level. By contrast, the monopsony model suggests that a moderate minimum wage may improve outcomes when firms possess wage-setting power. Search and matching models add further ambiguity, because higher wages may improve workers' search incentives and match quality, but may also reduce firms' incentives to create vacancies. Finally, efficiency wage theory introduces a direct link between wages and productivity, since higher wages may affect effort, turnover, worker selection and firm organization. Taken together, these models imply that minimum wage effects are likely to be context-dependent rather than uniform.

5.1 Main findings

The empirical analysis examined this question using a country-sector-year panel for OECD countries over the period 1995-2021. The main dependent variable was log labour productivity, while the key explanatory variable was the logarithm of minimum wage intensity. The regressions controlled for high-skilled labour share and capital intensity, and included country, sector and year fixed effects. The analysis then introduced interaction terms between minimum wage intensity and sectoral characteristics in order to assess whether the

relationship between minimum wages and productivity differs across productive environments.

The first main result is that minimum wage intensity does not show a statistically significant direct association with labour productivity. In both the baseline specification and the extended specifications, the coefficient on log minimum wage is positive but statistically insignificant. This result also holds when the minimum wage variable is lagged by one period. Therefore, the empirical evidence does not support the existence of a robust average direct effect of minimum wage intensity on labour productivity in the sample considered.

The second result is that productivity differences appear more closely related to structural sectoral characteristics than to minimum wage intensity itself. The share of high-skilled workers is positively and significantly associated with labour productivity in the simpler productivity specifications. This confirms the importance of human capital and labour quality for sectoral productivity. However, the interaction between minimum wage intensity and high-skilled share is not statistically significant, meaning that the analysis does not identify workforce skill composition as a robust channel through which minimum wage intensity changes the productivity relationship.

The third and most stable result concerns capital intensity. Capital intensity becomes positive and statistically significant when the interaction between minimum wage intensity and capital intensity is introduced. At the same time, the interaction between log minimum wage and capital intensity is negative and statistically significant. This pattern remains visible in the lagged robustness check. The result should be interpreted cautiously, because the magnitude of the interaction coefficient is small. Nevertheless, it suggests that the association between minimum wage intensity and labour productivity is less favorable in more capital-intensive sectors. In other words, the productive structure of sectors appears to matter for how minimum wage intensity is related to productivity outcomes.

Finally, the separate R&D regression does not provide evidence that minimum wage intensity directly increases innovation intensity. The coefficient on log minimum wage is not statistically significant when R&D intensity is used as the dependent variable. Instead, R&D

intensity is more strongly associated with capital intensity. This suggests that innovation intensity is more closely linked to the broader productive and technological structure of sectors than to minimum wage intensity itself.

5.2 Interpretation in light of the theoretical models

The empirical findings do not fully confirm any single theoretical model discussed in Chapter 2. They are not strongly consistent with a simple competitive interpretation, because the regressions do not show a significant negative direct association between minimum wage intensity and productivity. If the minimum wage operated only as a cost shock, one might expect clearer evidence of productivity losses or stronger negative effects. The absence of such a result suggests that the competitive model alone is too restrictive to describe the empirical relationship observed in the data.

At the same time, the findings do not provide strong evidence for a broad efficiency-wage mechanism. If higher minimum wages systematically increased productivity through effort, reduced turnover or organizational improvements, the coefficient on minimum wage intensity would be expected to be positive and significant. Instead, the direct coefficient is not statistically different from zero. This weakens the interpretation that minimum wage intensity acts as a general productivity-enhancing force across all sectors.

The results are more consistent with the broader message of the monopsony, search-and-matching and efficiency-wage frameworks: minimum wage effects depend on the context in which firms operate. The empirical analysis does not identify a uniform average relationship, but it does show that sectoral characteristics, especially capital intensity, shape the estimated association. This supports the idea that minimum wage policies should be analyzed through heterogeneous adjustment mechanisms rather than through a single aggregate effect.

The negative interaction between minimum wage intensity and capital intensity can be interpreted in this perspective. In highly capital-intensive sectors, productivity is already strongly linked to installed capital, machinery, technology and fixed production processes. Labour may therefore be a less flexible margin of adjustment. As a result, wage increases

may have less scope to generate productivity improvements through worker effort, turnover reduction or reorganization of tasks. This does not mean that minimum wages have a large negative effect in capital-intensive sectors, but rather that the productivity association becomes weaker when capital intensity is higher.

5.3 Limitations and directions for future research

The results should be interpreted with caution. First, the analysis is observational and therefore does not identify a fully causal effect of minimum wage policies. Minimum wages are not randomly assigned across countries and over time, and they may respond to economic conditions, productivity trends or political factors. Although the inclusion of country, sector and year fixed effects controls for several sources of persistent heterogeneity and common shocks, time-varying unobserved factors may still affect the estimates.

Second, the minimum wage variable is measured at the country-year level, while the outcome variables vary at the country-sector-year level. This makes it possible to study heterogeneous sectoral associations, but it also means that the analysis cannot observe sector-specific statutory minimum wage exposure directly. The relative minimum wage measure improves comparability across countries, but it does not fully capture enforcement, collective bargaining coverage, compliance or the actual share of workers affected by the minimum wage in each sector.

Third, the use of sector-level data may hide important firm-level heterogeneity. Firms within the same sector may differ substantially in productivity, market power, technology, workforce composition and exposure to the minimum wage. Some firms may respond through reorganization or innovation, while others may adjust employment, prices or profits. These opposing responses may offset each other in aggregate sector-level data.

Future research could extend the analysis in several directions. A first improvement would be to use firm-level or worker-level data, which would allow the analysis to measure exposure to minimum wage changes more directly and to identify adjustment margins more precisely. A second direction would be to include measures of labour-market concentration

or worker mobility, in order to test more directly the predictions of monopsony and search-based models. A third extension would be to examine longer lags, since investment, innovation and organizational change may require more than one year to materialize. Finally, future work could explore whether the effects differ across labor-intensive and capital-intensive sectors using more disaggregated industry classifications.

5.4 Final remarks

The main conclusion of the thesis is that minimum wage intensity is not associated with a robust direct change in labour productivity across OECD country-sector observations. However, this does not imply that minimum wage policies are irrelevant. Rather, the evidence suggests that their relationship with productivity is mediated by the productive structure of sectors. Human capital is positively related to productivity, capital intensity is an important structural determinant, and the interaction between minimum wage intensity and capital intensity represents the most stable form of heterogeneity identified in the empirical analysis.

This conclusion reinforces the broader argument developed throughout the thesis. Minimum wage policies should not be evaluated only through average effects or through a single theoretical lens. Their consequences depend on the characteristics of labour markets and production structures. The competitive model, monopsony model, search-and-matching framework and efficiency wage theory each capture part of this complexity, but none of them alone fully explains the empirical findings. The evidence points instead to a more nuanced interpretation: minimum wages may interact with sectoral conditions, and their productivity implications depend on the margins of adjustment available to firms.

For policy purposes, this means that the impact of minimum wage increases may differ across sectors. In some environments, firms may have room to absorb higher wage costs through improved organization, better worker retention or more efficient production. In others, especially where production is already highly capital-intensive, the scope for labour-related productivity gains may be more limited. Therefore, minimum wage policy should be

understood as part of a broader economic context, in which institutions, skills, capital and technology jointly determine the way firms respond to changes in labour costs.

Bibliography

Azar, J., Huet-Vaughn, E., et al. (2024). Minimum Wage Employment Effects and Labour Market Concentration. *The Review of Economic Studies*.

Blanchard, O., Amighini, A., Giavazzi, F. (2021). Macroeconomics, a european perspective, Fourth Edition. *Pearson Education Limited*

Boeri, T., Ours, J. van. (2021). The Economics of Imperfect Labor Markets, Third Edition. *Princeton University Press*.

Calligaris, S., Calvino, F., et al. (2023). Is there a trade-off between productivity and employment? A cross-country micro-to-macro study. *OECD science, technology, and industry policy papers*.

Cengiz, D., Dube, A., et al. (2019). The effect of minimum wages on low-wage jobs: evidence from the United States using a bunching. *National Bureau of Economic Research*.

Dube, A., Lindner Attila S. (2024). Minimum Wages in the 21st Century. *National Bureau of Economics Research (NBER) Working Paper*.

Economics Help. (n.d.). *Minimum wage*. Retrieved May 2026.

EU KLEMS & INTANProd. (n.d.). EU KLEMS & INTANProd database. Luiss Lab of European Economics. Retrieved May 2026.

Mortensen, D. T., Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61, 397–415.

OECD. (n.d.). Minimum wage relative to average wages. OECD Data Explorer. Retrieved May 2026.

OECD (2025). Employment protection and minimum wages. *OECD*

Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *American Economic Review*, 75, 676–690.

Rebitzer, J. B., Taylor, L. J. (1995). The consequences of minimum wage laws Some new theoretical ideas. *Journal of Public Economics*, Elsevier, vol. 56(2), pages 245-255, February.

Riley, R., Bondibene, C. R. (2015). Raising the standard: minimum wages and firm productivity. *National Institute of Economic and National Research (NIESR)*.

Rocheteau, G., Murat, T. (2007). The Minimum Wage and the Labor Market. *Federal Reserve Bank of Cleveland*, Economic Commentary 5/1/2007.

Ryan, B. (2015). The Living Wage: beware the trade-off deniers. *Institute of Economic Affairs (IEA)*.

Ryan, B. (2017). Misconceptions of the Efficiency Wage Hypothesis. *Cato Institute*.

Saylor Academy (2012). Microeconomics – Theory Through Application. *Saylor Academy*

Appendix

Table A1. Main results: contemporaneous log minimum wage

Dependent variable: Log labour productivity

	Baseline	MW-High Skill	MW-Capital	Extended
Log minimum wage	0.0377 (0.1049)	0.0376 (0.2180)	0.0527 (0.1049)	0.0576 (0.2221)
High-skilled share	7.3904*** (2.3712)	7.3832 (9.9627)	7.4117*** (2.3725)	7.7412 (10.0767)
Capital intensity	0.0002 (0.0008)	0.0002 (0.0008)	0.0212*** (0.0060)	0.0212*** (0.0061)
Log MW $\times$ High-skilled share		0.0019 (2.4111)		-0.0850 (2.4525)
Log MW $\times$ Capital intensity			-0.0056*** (0.0016)	-0.0056*** (0.0016)
Country FE	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	5815	5815	5815	5815
R-squared	0.966	0.966	0.966	0.966

Notes: Standard errors clustered at the country level are reported in parentheses.

Country, sector and year fixed effects are included separately.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A2. R&D Intensity regression: extended specification

Dependent variable: R&D Intensity

	R&D intensity
Log minimum wage	-0.0034 (0.0614)
High-skilled share	-1.8385 (4.0897)
Capital intensity	0.0101*** (0.0027)
Log MW $\times$ High-skilled share	-0.0969 (1.0180)
Log MW $\times$ Capital intensity	-0.0027*** (0.0007)
Country FE	Yes
Sector FE	Yes
Year FE	Yes
Observations	4933
R-squared	0.340

Notes: Standard errors clustered at the country level are reported in parentheses.,
Country, sector and year fixed effects are included separately.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A3. Lagged MW specifications

Dependent variable: Log labour productivity

	Baseline	MW-High Skill	MW-Capital	Extended
Lagged minimum wage	0.0881 (0.1076)	0.0959 (0.2202)	0.0978 (0.1082)	0.1083 (0.2233)
High-skilled share	7.3258*** (2.4103)	7.8295 (9.6489)	7.3553*** (2.4083)	8.0386 (9.7243)
Capital intensity	0.0001 (0.0008)	0.0001 (0.0008)	0.0191*** (0.0061)	0.0192*** (0.0061)
Lagged MW × High-skilled share		-0.1301 (2.3312)		-0.1765 (2.3581)
Lagged MW × Capital intensity			-0.0050*** (0.0016)	-0.0050*** (0.0016)
Country FE	Yes	Yes	Yes	Yes
Sector FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	5440	5440	5440	5440
R-squared	0.966	0.966	0.966	0.966

Notes: Standard errors clustered at the country level are reported in parentheses.

The log minimum wage variable is lagged by one period.

Country, sector and year fixed effects are included separately.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$