

The Exporter Hump

*Trade, Firm Premia and the Wage Distribution
across Spanish Regions, 2006–2022*

Jorge Menéndez-Pidal

Project Frontier, Madrid

ABSTRACT. With heterogeneous firms, trade raises wage inequality where few firms export and lowers it where most do. I show that the variance of log wages in a local labour market is hump-shaped in the employment share of exporters for any productivity distribution, and that the shape of the hump carries information about the tail of the productivity distribution that can be read from wage data alone: with a Pareto tail, wage dispersion among exporters does not vary with the exporter share, and the peak lies between $1/e$ and $1/2$. Foreign demand raises inequality up to and beyond that peak. Using five waves of Spain's Wage Structure Survey (2006–2022), which record whether each worker's employer sells mainly abroad, I find the hump across region–sector cells, with a peak at 0.35 (95% interval 0.29–0.43). Within-exporter wage dispersion does not vary with the exporter share, as a Pareto tail implies; the location of the peak points the same way but is affected by the survey's understatement of exporting. Most Spanish employment lies on the rising side of the hump. Three instruments—EU demand for each sector's products, the regional product mix, and exposure to the construction bust—have weak first stages, so the causal effect of exporting is not identified. Model-based counterfactuals bound it: a five-point rise in the exporter share raises the variance of log wages by about 2 points ($\times 1000$) according to within-cell estimates and an exact accounting, and by 13 according to the cross-section, against an observed fall of 34 points over 2006–2022. The 2008 bust raised measured inequality through the loss of mid-wage construction jobs but compressed the wage structure more. The 2019 minimum-wage increase compressed the lower tail in proportion to its regional bite.

Preliminary draft. Section 3 derives the model; Sections 4–9 take it to the microdata of Spain's Wage Structure Survey, 2006–2022, and to EU and Spanish customs data. This paper replaces the author's 2021 master's thesis on the same question. Data and code are available from the author. Comments welcome; please do not cite without permission.

Keywords: trade and inequality, exporter wage premium, firm heterogeneity, composition effects, minimum wage, Spain. **JEL:** F16, F14, J31, J38, R23.

1 Introduction

Does trade raise wage inequality? For a long time the answer was sought in relative factor prices: trade raises the return to the abundant factor. Modern trade theory moved the question inside industries. Exporters are larger and more productive and pay more (Bernard and Jensen, 1995; Bernard et al., 2007), so when trade expands exporting it also expands the wage gap between the firms that export and those that do not (Helpman et al., 2010). That mechanism has a feature that factor-price stories lack: it cannot raise inequality forever. When every firm exports, the gap between exporters and non-exporters disappears from the wage distribution. Trade should raise inequality where few firms export and lower it where most do.

This paper makes that statement precise and takes it to Spanish microdata. Its object is the variance of log wages in a local labour market—a region–sector cell—and its question is how three margins move it: which firms export, how much more exporters pay, and which jobs exist at all.

The model. The model has one equation. Wages combine a worker component, a firm premium that rises with productivity, and an additional premium for exporting; firms export when their productivity exceeds a threshold. I show that the variance of log wages is hump-shaped in the employment share of exporters *for any productivity distribution* (Proposition 1). The hump has two sources: the exporter premium itself and, more importantly, the covariance between productivity and exporter status. The location of its peak is informative about the tail of the productivity distribution: thin tails put it at one half, Pareto productivity between $1/e$ and one half (Proposition 2); with a Pareto tail above a different body the location also depends on the body. A rise in foreign demand raises inequality on the rising side of the hump and at its peak, because the exporter premium always widens; any turning point lies beyond the peak (Proposition 3). The model adds the two other forces that moved Spanish wages in this period: the composition of employment, whose effect on measured inequality depends on where in the distribution the destroyed jobs sat (Proposition 4), and a wage floor (Proposition 5).

The data. Spain is a good place to look. Its Wage Structure Survey records, for about 200,000 workers per wave, both the worker's wage and whether the employer's main market is the domestic, EU or world market. Between 2006 and 2022 Spain went through a construction-led boom, one of the deepest busts in Europe, an export expansion during and after that bust (Almunia et al., 2021), a labour reform that cut temporary contracts, and a 22% increase in the minimum wage. I use five waves of the survey's microdata, together with EU import data and Spanish customs records by province and product.

Findings. Six results stand out.

- (i) *The hump exists.* Across region–sector cells, wage inequality is concave in the exporter share, with a peak at 0.35 (95% interval [0.29, 0.43]). The hump survives sector-by-year and region fixed effects and controls for worker characteristics, and it disappears in a placebo that classifies firms by whether they sell beyond their region rather than abroad.
- (ii) *It is driven by the productivity covariance, and the evidence points to a Pareto tail.* The exporter

premium alone generates a tenth of the curvature; the productivity–exporting covariance accounts for the rest. The interval for the peak contains $1/e$ and excludes one half, although the survey’s understatement of exporting could produce that pattern on its own. More robustly, residual wage dispersion among exporters does not vary with the exporter share, while dispersion among non-exporters does, as expected if exporters occupy a Pareto tail above a log-normal body.

- (iii) *Most of Spain is on the rising side.* 90% of private employment is in cells whose measured exporter share lies below the lower end of the confidence interval for the peak, where more exporting raises inequality. But foreign-demand shift-share designs do not identify the causal effect: EU demand for a sector’s products does not predict Spain’s sectoral exports, consistent with exports having been pushed out by the domestic slump rather than pulled by foreign demand, the regional design has a weak first stage, and a design based on exposure to the construction bust has none. The regional design’s reduced forms show the predicted pattern—the upper tail widens, the lower tail does not move—without statistical significance.
- (iv) *The bust widened inequality through composition and narrowed it through the wage structure.* The construction jobs destroyed in 2008–2010 were homogeneous and paid near the median, so their loss raised measured inequality, by 0.019 in variance once all composition changes are accounted for. Holding composition fixed, the wage structure changed the variance by -0.034 . Hourly wage inequality fell.
- (v) *The minimum wage compressed the lower tail in proportion to its bite.* Cells where 10 percentage points more workers earned less than the new 2019 floor saw their $\ln(P50/P10)$ fall by 6.6 log points more by 2022, with no change in the upper tail.
- (vi) *The effect of exporting is bounded, and small.* Moving cells along the estimated hump, a five-point rise in the exporter share—the size of Spain’s rise between 2006 and 2014—raises the variance of log wages by 2.1 points ($\times 1000$, out of about 200) with the within-cell hump and by 12.9 with the cross-sectional one. An exact accounting of 2006–2014 gives 2.6. The two time-series estimates agree, and the cross-section is five times larger, which measures how much of the cross-sectional hump reflects persistent cell characteristics. Even the upper bound is small next to the observed fall of 33.7 points over 2006–2022.

Contribution. The paper contributes a distribution-free statement of the exporter hump, building on Helpman et al. (2010), and two tests of the productivity tail that do not require observing productivity; evidence on both from survey microdata; a decomposition of Spain’s adjustment that separates trade, composition and wage floors; model-based bounds on the effect of exporting when instruments fail; and a reinterpretation of the regional regressions in which these forces are usually confounded. It replaces the author’s 2021 master’s thesis, whose question it shares and whose conclusion—trade raised Spanish wage inequality only marginally—it confirms with a different method, data and argument.

Section 2 reviews the literature. Section 3 presents the model, Section 4 the data. Sections 5 to 9 report the evidence on the hump, the three instruments, the accounting of the bust, the

model-based bounds, and the minimum wage. Section 10 discusses what remains of the thesis, the limitations and next steps, and Section 11 concludes.

2 Related literature

Trade and inequality with heterogeneous firms. Exporters are larger, more productive and pay more (Bernard and Jensen, 1995; Bernard et al., 2007). In the models of Egger and Kreickemeier (2009), Helpman et al. (2010) and Amiti and Davis (2012) these facts link trade to wage inequality, and Helpman et al. (2010) show that inequality is first increasing and then decreasing in trade openness. Helpman et al. (2017) estimate such a model on Brazilian data and find that the between-firm component dominates. This paper derives the hump in a form that holds for any productivity distribution (Proposition 1), shows that the location of its peak and the behaviour of within-exporter dispersion are informative about the tail of that distribution (Proposition 2), and proves that the turning point of the effect of foreign demand always lies beyond the peak (Proposition 3). These statements are directly testable with survey microdata that record whether a worker's employer sells abroad, which is what the Spanish Wage Structure Survey provides. Verhoogen (2008) and Frías et al. (2012) document how exporting changes wage distributions within plants in Mexico; I work at the level of local labour markets.

Firm effects. The AKM literature (Abowd et al., 1999; Card et al., 2013, 2018) shows that firm wage premia account for a sizeable part of wage dispersion and that rent sharing links them to productivity. In the model the exporter premium is one such firm effect and the covariance between productivity and exporting is its sorting counterpart across firms. Without matched employer–employee panels I cannot estimate AKM effects; I adjust exporter premia for observable worker characteristics instead.

Spain's adjustment. Spanish exports grew strongly during and after the 2008–2013 slump, and Almunia et al. (2021) show that firms facing larger falls in domestic demand exported more. The collapse of construction and the reliance on temporary contracts shaped who lost jobs (Bentolila et al., 2012; Fernández-Villaverde et al., 2013), and earnings inequality in Social Security data was strongly countercyclical (Bonhomme and Hospido, 2017). Donoso et al. (2015) study Chinese import competition across Spanish provinces. I show that, for hourly wages in the private sector, the bust widened inequality through composition but compressed it through the wage structure, and that exporting added little to either.

Minimum wages. Regional variation in the bite of a national minimum wage (Card, 1992) and its effect on the lower tail of the wage distribution (Lee, 1999; Cengiz et al., 2019) provide the template for Section 9. Proposition 5 adds that the size of the compression depends on whether the floor pins the tenth percentile, and that the shift in the median identifies the employment loss from wage data alone.

Shift-share designs. The foreign-demand shocks are shift-share variables (Autor et al., 2013; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022). Because each sector receives a single shock per period, I absorb it with sector-by-period fixed effects and identify the model from

differential exposure across regions, with inference based on permuting shocks across sectors in the spirit of Adão et al. (2019).

Regional trade and inequality. Rodríguez-Pose (2012) relates trade openness to inequality *between* regions. The present paper studies inequality *within* regional labour markets. The distinction matters, and conflating the two was a weakness of the master's thesis that this paper replaces.

3 A model of the exporter hump

The model has one object, the variance of log wages in a local labour market, and asks how trade moves it. A local labour market, or *cell*, is a region–sector pair. It contains firms that differ in productivity, and the more productive of them export. Everything that follows is a statement about how three margins—which firms export, how much more exporters pay, and which jobs exist at all—map into the cell's wage distribution.

3.1 Environment

Let x denote the log productivity of a worker's employer, relative to the domestic survival cutoff, with distribution F *weighted by employment*. The wage of a worker is

$$\ln w = a + \lambda x + \varphi X, \quad X = \mathbf{1}[x \geq x_0], \quad (1)$$

where a is a worker component with variance σ_a^2 , independent of x ; $\lambda \in (0, 1)$ is the elasticity of the firm wage premium to productivity (rent sharing); X indicates that the employer exports; and $\varphi > 0$ is the additional premium paid by exporters. Firms export when their productivity exceeds a threshold x_0 , and the employment share of exporters is $\rho = 1 - F(x_0)$.

Equation (1) is the reduced form of the heterogeneous-firm labour-market models of Helpman et al. (2010) and Amiti and Davis (2012). With search frictions and ex-post bargaining, firms pay a share of revenue per worker; revenue rises with productivity, and access to the export market multiplies it by a factor $\Upsilon(A^*) > 1$ that increases with foreign demand A^* . Hence $\varphi = \lambda \ln \Upsilon(A^*)$. Two features matter for what follows. First, the exporter premium is a *firm* effect, as in the AKM decompositions of Abowd et al. (1999) and Card et al. (2013), not a return to worker skill. Second, exporters are not a random draw: they are the most productive firms, so the exporter dummy and the productivity component of the premium covary. This covariance turns out to drive the evidence.

3.2 The hump

Let $\Delta(\rho) \equiv \mathbb{E}[x \mid X = 1] - \mathbb{E}[x \mid X = 0] > 0$ be the productivity gap between exporters and non-exporters when a share ρ of employment is in exporters.

Proposition 1 (The exporter hump). *For any productivity distribution F , the variance of log wages in the cell is*

$$V(\rho) = \sigma_a^2 + \lambda^2 \text{Var}(x) + \rho(1 - \rho)[\varphi^2 + 2\lambda\varphi\Delta(\rho)]. \quad (2)$$

The term in ρ is positive on $(0, 1)$ and vanishes as $\rho \rightarrow 0$ and $\rho \rightarrow 1$. Wage inequality is therefore hump-shaped in the employment share of exporters: it is lowest when no firm or every firm exports.

The first bracketed term is the familiar between-group component: exporters pay φ more. The second is the covariance between productivity and exporter status, $\text{Cov}(x, X) = \rho(1-\rho)\Delta(\rho)$, scaled by $2\lambda\varphi$. When few firms export, they are a thin and very productive tail; adding exporters raises inequality. When most firms export, the remaining non-exporters are the outliers; adding exporters lowers it. With $\lambda = 0$ the model collapses to a two-type mixture and V peaks at $\rho^* = \frac{1}{2} + (\sigma_X^2 - \sigma_N^2)/(2\varphi^2)$, where σ_X^2, σ_N^2 are the within-type variances (Appendix A).

The location of the peak is informative about the tail of the productivity distribution.

Proposition 2 (Pareto tail). *If productivity is Pareto, so that x is exponential with parameter $\tilde{k}$ in the employment-weighted distribution, then $\text{Cov}(x, X) = -\rho \ln \rho / \tilde{k}$ and*

$$V(\rho) = \sigma_a^2 + \frac{\lambda^2}{\tilde{k}^2} + \varphi^2 \rho(1-\rho) - \frac{2\lambda\varphi}{\tilde{k}} \rho \ln \rho. \quad (3)$$

V is strictly concave, $V''(\rho) = -2\varphi^2 - 2\lambda\varphi/(\tilde{k}\rho)$, and its maximum ρ^* lies in $(1/e, 1/2)$. If instead x is normal (log-normal productivity), the covariance term is symmetric in ρ and V peaks at $\rho^* = 1/2$.

The covariance term $-\rho \ln \rho$ is steep where exporters are scarce, $V'' \rightarrow -\infty$ as $\rho \rightarrow 0$, which is where most Spanish cells are. Thin-tailed distributions place the peak at one half; a Pareto tail pulls it towards $1/e \approx 0.37$. The peak is thus a test of the tail that does not require observing productivity.

Remark (Body and tail). Firm size and productivity distributions are well described by a log-normal body with a Pareto upper tail (Axtell, 2001; Head et al., 2014; Nigai, 2017). If exporters lie in the Pareto tail, the memorylessness of the exponential implies that the dispersion of x among exporters does not depend on ρ , whereas the dispersion among non-exporters, who span the body, does. Equation (2) still holds, and (3) holds up to a term linear in ρ . The remark yields a second, independent test: within-exporter dispersion should be flat in ρ , within-non-exporter dispersion should not.

3.3 Foreign demand

A rise in foreign demand $\ln A^*$ moves two margins. On the extensive margin it lowers the export threshold: $d \ln \rho / d \ln A^* = \zeta > 0$. On the intensive margin it raises the exporter premium: $d\varphi / d \ln A^* = \Lambda \equiv \lambda s_x > 0$, where s_x is the export share of exporters' sales.

Proposition 3 (Foreign demand and inequality). *For any productivity distribution,*

$$\frac{dV}{d \ln A^*} = \zeta \rho G(\rho) + \Lambda \partial_\varphi V, \quad \partial_\varphi V = 2\rho(1-\rho)[\varphi + \lambda\Delta(\rho)] > 0, \quad (4)$$

where $G = \partial V / \partial \rho$ is the slope of the hump.

- (a) *The intensive margin always raises inequality.*
- (b) *At the peak of the hump $G = 0$, so foreign demand raises inequality there. Any turning point $\tilde{\rho}$, above which foreign demand compresses wages, lies to the right of the peak: $\tilde{\rho} > \rho^*$.*
- (c) *With a Pareto tail, $dV / d \ln A^* = \rho H(\rho)$ with $H(\rho) = \zeta[\varphi^2(1-2\rho) + c\varphi(-\ln \rho - 1)] + \Lambda[2\varphi(1-\rho) - c \ln \rho]$, $c \equiv 2\lambda/\tilde{k}$. H is strictly decreasing, $H' = -(\Lambda + \varphi\zeta)(c + 2\varphi\rho)/\rho$, with $H(0^+) = +\infty$ and $H(1) < 0$: the turning point is unique and $\tilde{\rho} \in (\rho^*, 1)$.*

- (d) With a Pareto tail the effect $E(\rho) = \rho H(\rho)$ is itself a concave hump: $E(0) = 0$, $E'(0^+) = +\infty$, $E'' = -(\Lambda + \varphi\zeta)(c/\rho + 4\varphi) < 0$.

Part (b) is the result the evidence needs. It does not depend on the productivity distribution: wherever the hump is still rising, a foreign demand shock raises wage inequality, and it can only compress it in cells already well past the peak. Part (d) yields the estimating equation. Writing the shock as Z , the model predicts

$$\Delta\rho_c = \zeta \rho_{c0} Z_c, \quad \Delta V_c = \underbrace{\zeta \rho_{c0} G_{c0}}_{\text{extensive}} Z_c + \underbrace{\Lambda \partial_\varphi V_{c0}}_{\text{intensive}} Z_c, \quad (5)$$

so a cell's response to a common shock scales with its initial exporter share. Section 6 estimates both equations.

Import competition. An import shock that lowers the average wage of non-exporters raises φ and, by (4), inequality. If it also destroys low-wage jobs in exposed industries, the composition margin of the next subsection operates with the opposite sign.

3.4 Composition

Downturns change who is employed. Let a share d of jobs be destroyed, and let the destroyed jobs have mean log wage m_D and variance V_D , against m and V for all jobs before the shock.

Proposition 4 (Composition). *The variance of log wages among surviving jobs changes by*

$$V_R - V = \frac{d}{1-d} \left[V - V_D - \frac{(m - m_D)^2}{1-d} \right]. \quad (6)$$

Measured inequality rises if and only if the destroyed jobs were homogeneous and close to the middle of the distribution, $V_D + (m - m_D)^2/(1-d) < V$. Destroying low-wage jobs compresses the observed distribution.

The proposition is an accounting statement; the economics lies in which jobs are destroyed. Two features of Spanish labour markets determine that. With downward wage rigidity (Schmitt-Grohé and Uribe, 2016), large negative shocks are absorbed through employment rather than wages: if wages can fall by at most ω , a shock of size x to labour demand rations employment by $x - (\varepsilon + \nu)\omega$, where ε and ν are the demand and supply elasticities, whereas positive shocks of the same size move wages. With a dual labour market (Bentolila et al., 2012), the jobs destroyed are temporary ones, since permanent contracts carry firing costs: the temporary separation rate is $\min\{1, d/\tau\}$ for a temporary share τ . In Spain's 2008 bust the destroyed jobs were temporary jobs in construction, which paid near the median and were compressed by collective agreements. Proposition 4 then predicts that the bust *widened* measured inequality through composition alone (Appendix A states these remarks formally).

3.5 Wage floors

Let latent log wages have distribution F and let a minimum wage m bind. A worker below the floor loses the job with probability e and is otherwise paid m .

- Proposition 5** (Minimum wage). (a) *The observed tenth percentile equals the floor if and only if $F(m) \geq 0.1/(1 - 0.9e)$.*
- (b) *$\ln(P50/P10)$ is non-increasing in the bite $F(m)$.*
- (c) *Disemployment widens the observed $\ln(P50/P10)$ when the floor pins the tenth percentile, and narrows it otherwise.*
- (d) *The observed median is $F^{-1}(\frac{1}{2} + \frac{1}{2}eF(m))$, so its shift identifies the employment loss $eF(m)$ from wage data alone.*
- (e) *If the floor binds only among non-exporters, it lowers σ_N^2 and φ ; with $\sigma_X^2 \geq \sigma_N^2$ after the increase, the two-type peak ρ^* moves right.*

Part (b) gives the test in Section 9: compression of the lower tail should be proportional to the share of workers affected, with no effect on the upper tail. Part (c) warns that the size of the compression is not a sufficient statistic for the employment effect.

3.6 An exact accounting

Between two dates, the change in the two-type variance decomposes exactly (Shapley):

$$\Delta V = \underbrace{\bar{\rho} \Delta \sigma_X^2 + (1 - \bar{\rho}) \Delta \sigma_N^2}_{\text{within-type dispersion}} + \underbrace{\Delta \rho (\bar{\sigma}_X^2 - \bar{\sigma}_N^2)}_{\text{within-type composition}} + \underbrace{\Delta[\rho(1 - \rho)] \bar{\varphi}^2}_{\text{exporter extensive}} + \underbrace{\Delta \varphi^2 \bar{\rho}(1 - \bar{\rho})}_{\text{exporter intensive}}, \quad (7)$$

where a bar denotes the average of the two dates. Section 7 applies it to Spain. The binary accounting attributes the covariance term of Proposition 1 to within-type dispersion, so it is a lower bound on the role of trade.

Verification. Every proposition is verified symbolically and numerically in the replication code (`code/01_model_checks.py`, 50 checks): closed forms with `sympy`, comparative statics with finite differences, and Propositions 1–3 by Monte Carlo on simulated workforces under Pareto, log-normal, gamma, uniform and body-tail productivity.

4 Data

Wage microdata. The Spanish Wage Structure Survey (*Encuesta de Estructura Salarial*, EES) is a matched employer–employee survey run by the National Statistics Institute every four years. I use the 2006, 2010, 2014, 2018 and 2022 waves, which are publicly available with about 210,000–240,000 workers each. For each worker the EES records October earnings and hours, contract type, education, occupation, age, tenure and sex; for the employer it records sector, size, ownership, the bargaining regime and, crucially, the firm’s *main market*: local or regional, national, the European Union, or the rest of the world. Region is coded at the NUTS-1 level (seven regions).

The wage is log gross hourly earnings in October, constructed exactly as the INE does: base pay, complements, overtime and extraordinary October payments, divided by contractual hours plus overtime. I keep private-sector employees paid for the full month, trim the top and bottom 0.1% of hourly wages in each wave, and use the survey weights throughout. A

worker's employer is an *exporter* if its main market is the EU or the world. This understates exporting, because many firms whose main market is national also export; Section 5 discusses the implications. Sectors are harmonised across the CNAE-93 classification of 2006 and the CNAE-2009 classification of later waves into 17 groups, of which seven are manufacturing (Appendix B). A *cell* is a NUTS-1 region by sector by wave; I keep cells with at least 100 workers and both exporters and non-exporters, which leaves 543 cell-waves.

For each cell I compute the exporter share ρ , the raw exporter premium $\varphi = \mathbb{E}[\ln w \mid X = 1] - \mathbb{E}[\ln w \mid X = 0]$, the premium adjusted for worker observables (the coefficient on X in a within-cell regression on sex, age, education, occupation, tenure, contract and part-time status), the within-type variances, the total variance V , the residual variance V^{res} after the same worker controls (estimated within each wave on all workers), and percentiles.

Foreign demand. Foreign demand for product k (HS 4-digit) is the imports of the other 26 EU members from the world, net of their imports from Spain, from Eurostat's Comext database, using a constant set of reporting countries; the EU absorbs about two thirds of Spanish goods exports. Its growth between EES waves, g_k , is the shift. Two share designs complete the shock.

Design A (national product mix). For each sector s and period p from t_0 to t_1 ,

$$Z_{sp} = \sum_{k \in s} \frac{X_{k,t_0}^{ES}}{\sum_{k' \in s} X_{k',t_0}^{ES}} g_k(t_0, t_1), \quad (8)$$

where X_{k,t_0}^{ES} are Spain's exports of product k in the base year. The shock varies across sectors and periods; it is common to all regions within a sector-period.

Design B (regional product mix). The Spanish Tax Agency (AEAT) publishes monthly exports by province of origin and 8-digit product from 2016. I aggregate 2016–2017 exports to NUTS-1 region $\times$ sector $\times$ HS4 and compute $Z_{rs} = \sum_k w_{k,rs} g_k(2018, 2022)$, which varies within sectors across regions, for the 2018–2022 period only. The AEAT totals reproduce official Spanish exports (254, 274, 283 and 385 billion euros in 2016, 2017, 2018 and 2022).

HS chapters are mapped to the harmonised sectors in Appendix B; agricultural chapters are excluded because agriculture is outside the EES.

5 Facts and the hump

Facts. Table 1 summarises the private-sector wage distribution. The employment share of exporters rose from 11.0% in 2006 to 16.1% in 2014, during and after the domestic slump, and has since stayed close to 15%. Exporters pay about 27 log points more than non-exporters, and about 10 points more after adjusting for the observable characteristics of their workers. Yet the between-type term $\rho(1 - \rho)\varphi^2$ accounts for only 2.8–5.1% of the variance of log wages. The variance itself *fell* throughout, from 0.225 to 0.191, with the largest drops in 2006–2010 and 2018–2022. The lower tail compressed sharply in the last period: $\ln(P50/P10)$ fell from 0.401 to 0.360. Temporary contracts fell from 27.3% to 6.2% of full-month private employment, and construction from 15% to 7%. The 2022 temporary share is well below the Labour Force Survey's figure for private employees in the fourth quarter of 2022 (15.1%); the EES excludes agriculture, is reported by employers rather than workers, and records the October after the

TABLE 1. Wages and exporters in the Spanish private sector, 2006–2022

	2006	2010	2014	2018	2022
Exporter share ρ	11.0	12.0	16.1	15.6	14.9
Raw exporter premium φ	0.254	0.272	0.258	0.281	0.269
Adjusted premium (mean over cells)	0.092	0.101	0.106	0.109	0.108
Variance of log wages V	0.225	0.210	0.206	0.203	0.191
Between-type share of V (%)	2.8	3.7	4.4	5.1	4.8
$\ln(P90/P10)$	1.182	1.154	1.153	1.132	1.100
$\ln(P50/P10)$	0.412	0.400	0.406	0.401	0.360
$\ln(P90/P50)$	0.770	0.754	0.747	0.731	0.739
Temporary contracts (%)	27.3	20.7	19.3	20.2	6.2
Firm-level agreements (%)	12.8	14.5	14.9	14.7	13.4
Construction employment (%)	15.2	9.9	6.1	7.0	7.5

Notes: EES microdata, private-sector employees paid for the full month of October, survey weights. Wages are log gross hourly earnings. Exporters are firms whose main market is the EU or the rest of the world. The 2022 temporary share reflects the 2021 labour reform; it is below the Labour Force Survey's private-sector figure (15.1% in 2022Q4) because the EES excludes agriculture and is reported by employers.

TABLE 2. The exporter hump across region–sector cells

	V (1)	V (2)	V^{res} (3)	V (4)
ρ	0.507*** (0.121)	0.213*** (0.067)	0.131*** (0.040)	0.077 (0.061)
ρ^2	−0.720*** (0.160)	−0.369*** (0.099)	−0.195*** (0.055)	−0.096 (0.078)
Implied peak ρ^*	0.35	0.29	0.34	0.40
Sector FE, year FE	yes		yes	
Sector×year FE, region FE		yes		
Cell FE, year FE				yes
Cell-waves	543	543	543	543

Notes: Weighted least squares, cell employment weights; standard errors clustered by region–sector in parentheses. *, **, ***: 10, 5, 1%. V^{res} is the variance of log wages net of sex, age, education, occupation, tenure, contract and part-time status. Column (1): cluster-bootstrap 95% interval for the peak [0.29, 0.43]; 100% of 400 replications concave.

2021 reform took effect. I use it only descriptively.

The hump. Proposition 1 predicts that, across cells, wage inequality is concave in the exporter share. Table 2 and Figure 1 confirm it. With sector and year fixed effects, the coefficient on ρ^2 is −0.72 (s.e. 0.16), and the implied peak is at $\rho^* = 0.35$, with a cluster-bootstrap 95% interval of [0.29, 0.43]; every bootstrap replication is concave. The hump survives sector-by-year and region fixed effects (peak 0.29) and, importantly, survives in the residual variance after removing sex, age, education, occupation, tenure, contract and part-time status (coefficient −0.20, s.e. 0.06; peak 0.34). It is not a composition effect of observable worker characteristics.

What shapes the hump. The between-type term alone cannot generate this curvature: with $\varphi \approx 0.27$ it implies a coefficient on ρ^2 of about $-\varphi^2 \approx -0.07$, an order of magnitude smaller than the estimate. The covariance between productivity and exporting, the second term in (2), can. Fitting (3) directly, the coefficient on $-\rho \ln \rho$ is 0.43 (s.e. 0.10). The location of the peak




FIGURE 1. The exporter hump. Panel (a): variance of log hourly wages across region–sector cells against the employment share of exporters, after removing sector and year effects, in 15 employment-weighted bins; lines are fits on the same residualised data. Panel (b): Proposition 2. With Pareto productivity the hump peaks between $1/e$ and $1/2$ (dotted lines); with log-normal productivity it peaks at $1/2$. Illustrative parameters.

is suggestive but not decisive. Thin-tailed productivity puts it at one half and pure Pareto productivity between $1/e$ and one half (Proposition 2); with a Pareto tail above a log-normal body it depends on the body as well. The estimated interval $[0.29, 0.43]$ excludes one half and contains $1/e$. But the survey records the firm's *main* market and therefore understates the exporter share. If the measured share is a fraction $\mu < 1$ of the true one, the measured peak is μ times the true peak, so a true peak at one half would appear at 0.35 if the survey captured 70% of exporter employment. The bias works in favour of a spurious Pareto reading, and the peak alone cannot rule out thinner tails. The test that follows does not rely on the location of the peak.

The peak could become a quantitative test with outside information. From Proposition 2, the first-order condition at the peak is $\varphi(1 - 2\rho^*) = c(1 + \ln \rho^*)$ with $c = 2\lambda/\tilde{k}$, so the peak identifies the ratio c/φ : the pass-through of productivity to wages relative to the exporter premium, scaled by the tail parameter. The tail parameter can be estimated from the size distribution of Spanish firms (the Central Business Register reports firms by employment band), and the pass-through from the rent-sharing literature. Combined with the measured premium, they predict a peak that the survey estimate can be compared with, once the survey's understatement of exporting is corrected with firm-level export status. That comparison needs data this thesis does not use, and is left for the next version.

The tail test. The body–tail remark of Section 3.2 gives an independent test. If exporters sit in a Pareto tail, the residual wage dispersion among exporters should not vary with ρ , while dispersion among non-exporters, who span the body of the distribution, should. Regressing within-type residual variances on ρ and ρ^2 with sector and year effects, the joint p -value is 0.27 for exporters and 0.026 for non-exporters. Raw (unadjusted) within-exporter dispersion does vary with ρ ($p = 0.007$), which is what one expects if worker composition differs across cells. This resolves what looked like a tension in the pure Pareto model, whose memorylessness implies constant dispersion among *all* types: the data say the tail is Pareto and the body is not,

as in Head et al. (2014) and Nigai (2017).

Placebo. If the hump reflects international trade, it should appear when firms are classified by whether they sell abroad, not by whether they sell beyond their region. Classifying as “exporters” all firms whose main market is national, EU or world, the coefficient on ρ^2 is 0.07 (s.e. 0.07): there is no hump.

The limits of the cross-section. Within cells over time, with cell fixed effects, the coefficient on ρ^2 is -0.10 (s.e. 0.08). The exporter share moves little within a cell over four-year intervals and is measured with sampling noise, so the within estimate is uninformative rather than contradictory. But it means that the hump is, so far, a cross-sectional regularity: cells with intermediate exporter shares are more unequal than cells with few or many exporters. The next section asks whether exogenous changes in foreign demand move inequality along it.

Where Spain sits. 92% of private employment is in cells with a *measured* exporter share below $1/e$, and 90% in cells below the lower end of the confidence interval for the peak, which is estimated on the same measured shares. Because the survey understates exporting, true shares are higher; the second comparison, which uses the same measure on both sides, is the relevant one. By Proposition 3(b), in all of these cells a rise in foreign demand should raise wage inequality.

6 Foreign-demand shocks

Proposition 3 predicts that a rise in foreign demand raises the exporter share, $\Delta\rho_c = \zeta\rho_{c0}Z$, and, below the turning point, raises wage inequality, mostly in the upper tail where the exporter premium is paid. I take this to the data with the two shift-share designs of Section 4 and with a third design based on exposure to the construction bust. All three are weak; I report them in full because the reasons are informative.

Regional product mix. Panel A of Table 3 uses the AEAT design. Within a sector, regions differ in what they export, so EU demand growth over 2018–2022 differs across them (the within-sector standard deviation of Z_{rs} is 0.06). Regions whose products faced faster EU demand growth saw their exporter share rise: the coefficient is 0.37 (s.e. 0.15), with the expected sign. But the first stage is weak ($F = 6.2$), the permutation p -value that reassigns shocks across regions within sectors is 0.12, and the shock does not significantly predict regional export growth (0.48, s.e. 0.57). The reduced forms have the pattern of Proposition 3: the upper tail widens (0.34, s.e. 0.33, permutation $p = 0.14$) while the lower tail does not move (0.01), and the variance rises slightly (0.05, s.e. 0.10). None of these is statistically significant. The implied slope of the hump, instrumenting $\Delta\rho$ with Z_{rs} , is 0.12 (s.e. 0.25), which is uninformative. Two features of the period limit power. With 51 cells, a single period and shocks that differ within sectors by a few percentage points, the design has little variation to work with. And 2018–2022 includes the pandemic and the 2022 energy-price shock, which move nominal import values for reasons unrelated to demand for Spanish goods.

TABLE 3. Foreign-demand shocks, exporting and the wage distribution

Dependent variable (change)	Coef.	s.e.	Perm. p	N
<i>Panel A. Regional product mix (AEAT), 2018–2022; regressor Z_{rs}; sector FE</i>				
Regional exports, $\Delta \ln X_{rs}$	0.477	(0.573)		51
Exporter share, $\Delta \rho$ (first stage)	0.372**	(0.149)	0.12	51
Variance of log wages, ΔV	0.046	(0.097)	0.57	51
Residual variance, ΔV^{res}	0.038	(0.052)	0.46	51
Upper tail, $\Delta \ln(P90/P50)$	0.339	(0.335)	0.14	51
Lower tail, $\Delta \ln(P50/P10)$	0.013	(0.135)	0.93	51
IV: ΔV on $\Delta \rho$, instrument Z_{rs}	0.125	(0.248)		51
<i>Panel B. National product mix (Comext), four periods 2006–2022; regressor $\rho_{c0}Z_{sp}$</i>				
Relevance: Spain's sector exports, $\Delta \ln X_{sp}$ on Z_{sp}	0.038	(0.462)		24
Exporter share, $\Delta \rho$	−1.590***	(0.614)	0.17	201
Variance, ΔV	−0.273	(0.250)	0.50	201
Residual variance, ΔV^{res}	−0.185	(0.148)	0.42	201
Upper tail, $\Delta \ln(P90/P50)$	−0.662	(0.908)	0.71	201
Lower tail, $\Delta \ln(P50/P10)$	−0.916	(0.594)	0.13	201

Notes: Changes between EES waves in manufacturing and mining region–sector cells, 2018 (panel A) or base-wave (panel B) employment weights. Z is the growth of EU (ex-Spain) import demand weighted by the region–sector's 2016–17 product mix (panel A) or Spain's base-year product mix within the sector (panel B). Panel A: heteroskedasticity-robust standard errors; permutation p -values from 999 reassignments of Z_{rs} across regions within sector. Panel B: sector $\times$ period, region $\times$ period and $\rho_{c0}\times$ period controls; standard errors clustered by cell; permutation of Z_{sp} across sectors within period. The relevance row regresses Spain's export growth by sector on Z_{sp} with period effects (24 sector–periods). *, **, ***: 10, 5, 1%.

National product mix. Panel B uses Spain's national product mix within each sector and all four periods. Here the design fails at the first step. Across 24 sector–periods, EU demand growth for a sector's products does not predict the growth of Spain's exports in that sector (coefficient 0.04, s.e. 0.46). Without relevance at the sector level, the differential-exposure regressions cannot identify ζ : the coefficient on $\rho_{c0}Z$ in the exporter-share equation is negative (−1.59, s.e. 0.61) but not robust to permutation ($p = 0.17$), and the reduced forms are imprecise.

Why foreign demand does not explain Spanish exports. The failure of the national design is itself a finding, and it fits what is known about the period. Spanish exports did not grow because European demand for Spanish products grew; they grew because domestic demand collapsed and firms with capacity to spare sold abroad (Almunia et al., 2021). In the language of the model, the shock that moved ρ in 2008–2014 was a fall in the domestic market, which lowers the export threshold relative to the domestic one, not a rise in A^* . That shock was the collapse of construction, which also moves wages through composition (Proposition 4). The next design tries to use it anyway.

A venting-out design. If the slump rather than foreign demand pushed firms into exporting, the natural instrument is exposure to the slump. I combine the size of each region's housing boom, the growth of its dwelling stock between 2001 and 2007, with the share of each sector's products bought by construction in Spain's 2010 input–output table. The interaction $Z_{rs} = \text{Boom}_r \times \text{Dep}_s$ varies within sectors and within regions, and region and sector effects absorb the main terms. The boom measure is strongly related to the subsequent bust: across regions its correlation

with the 2006–2014 change in log construction employment is -0.82 . But the design has two problems. First, dependence on construction is concentrated in one sector: non-metallic minerals (cement, ceramics, glass) sell 0.44 of their output to construction, every other sector at most 0.075. Second, it has no first stage. Over 2006–2014 the coefficient of Δp on Z_{rs} is 0.061 (s.e. 0.045; $F = 1.8$), and exact randomization inference over all 5,040 assignments of the boom to regions gives $p = 0.72$; with the alternative boom measure, the 2006 construction share of employment, it is -0.025 (s.e. 0.064). The reduced form is more striking. With the construction-share measure of the boom, the upper tail widened in exposed cells even holding worker composition at its 2006 level ($\Delta \ln(P90/P50)$: 0.20, s.e. 0.06, randomization $p = 0.011$). But the effect does not run through the exporter share, and once minerals are dropped it becomes imprecise (0.59, s.e. 0.50): it rests on a single sector in six regions. I read it as a sign that the bust reshaped the wage distribution of the industries most tied to construction, plausibly through the exit of their weakest firms, and not as evidence on exporting.

What the evidence on trade amounts to. Three designs, three weak first stages. Two pieces of evidence point the same way, but neither is decisive. The cross-section shows the hump, the tail test supports its Pareto form, and in most cells the model's comparative static (rising inequality with exports) applies. The one design with a usable first stage shows the predicted pattern—upper tail widening, lower tail unchanged—without the power to reject zero. Manufacturing cells, where exporting is concentrated, have an average exporter share of 0.37, close to the peak of the hump. There the extensive margin is nearly flat ($G \approx -0.02$ at the cross-sectional fit), so the model itself predicts small effects driven by the exporter premium alone.

7 Accounting for the bust and the recovery

Exporting adds little. Table 4 applies the exact decomposition (7) to the national distribution of exporters and non-exporters, and Figure 2 plots it. In every period the change in variance is dominated by compression within both types of firm. The exporter margins raised inequality in each of the first three periods, while the exporter share rose and the premium widened, by 1.2 to 1.5 points of variance ($\times 1000$) per period. That is a tenth of the 2006–2010 fall and a third of the smaller changes in 2010–2018. The decomposition is a lower bound on the role of exporting, because the covariance term of Proposition 1 is counted within types. But the order of magnitude agrees with the master's thesis this paper replaces: international trade raised wage inequality in Spain, marginally.

The bust widened inequality through composition. The model's account of 2006–2010 is Proposition 4. In 2006 construction workers had a wage variance of 0.118, against 0.225 for all private employees, and their mean wage sat at the 55th percentile: homogeneous jobs in the middle of the distribution. By 2010 43% of those jobs were gone. The sign condition of (6) holds with room to spare ($V_D + (m - m_D)^2/(1 - d) = 0.121 < V$), and the destruction of construction jobs alone predicts a rise in variance of 0.0073. A reweighting that holds the joint distribution of sector, contract, education and age at its 2006 level (DiNardo et al., 1996) measures the full composition effect: it raised the variance by 0.019 and $\ln(P90/P10)$ by 0.064. Downward wage rigidity and dual contracts explain why this was the margin of adjustment: wages could not

TABLE 4. Exact accounting of the change in wage variance

	ΔV	Within-type dispersion	Within-type composition	Exporter extensive	Exporter intensive
2006–2010	−14.8	−16.4	0.1	0.5	1.0
2010–2014	−4.1	−6.0	0.7	2.1	−0.9
2014–2018	−3.1	−4.5	−0.1	−0.2	1.7
2018–2022	−11.7	−10.2	−0.2	−0.4	−0.9
<i>Reweighting 2006–2010 (DiNardo et al., 1996): composition vs structure</i>					
	Total	Composition	Structure		
$V (\times 1000)$	−14.8	19.0	−33.7		
$\ln(P90/P10) (\times 100)$	−2.8	6.4	−9.2		
$\ln(P50/P10) (\times 100)$	−1.3	2.5	−3.7		
$\ln(P90/P50) (\times 100)$	−1.6	3.9	−5.4		

Notes: Upper panel: equation (7) applied to the national two-type mixture (exporters and non-exporters), changes $\times 1000$; components sum exactly to ΔV . Lower panel: 2010 wages reweighted to the 2006 joint distribution of sector, contract, education and age.




FIGURE 2. Exact accounting of the change in the variance of log hourly wages, private sector. For each period, the first bar is the total change and the others are the four components of equation (7), which sum exactly to it.

fall, so jobs went, and the jobs that went were temporary and concentrated in construction.

But the wage structure compressed more. Holding composition fixed, the wage structure changed the variance by -0.034 and $\ln(P90/P10)$ by -0.092 , more than offsetting composition. The net change, -0.015 , is a fall. For hourly wages of the employed in the private sector, the premise of the 2021 thesis—that the crisis raised wage inequality—is not borne out. Inequality of *earnings*, which includes days and hours not worked, did rise (Bonhomme and Hospido, 2017); inequality of hourly *wages* did not. A plausible reading of the compression is that nominal wages were rigid downwards and sectoral agreements protected the lower part of the scale, while variable pay higher up adjusted; the decomposition itself does not distinguish these channels.

8 How much could exporting matter?

The designs above cannot identify the causal effect of exporting, but the model can bound it. Proposition 1 gives the variance of each cell as a function of its exporter share. If the

TABLE 5. How much could more exporting change wage inequality? Model-based counterfactuals, 2018

	ΔV $\times 1000$	90% interval	$\Delta \ln \frac{P_{90}}{P_{10}}$ $\times 100$	Share of $ \Delta V_{06-22} $
<i>Cross-sectional hump, $\rho(1 - \rho)$ and $-\rho \ln \rho$ (sector, year FE)</i>				
Exporter share +5 pp, every cell	12.9	[5.8, 18.6]	3.6	38%
Each cell to its sector's top regional share	9.2	[2.2, 16.6]	2.6	27%
<i>Cross-sectional hump, quadratic (sector, year FE)</i>				
Exporter share +5 pp, every cell	12.1	[5.6, 18.7]	3.4	36%
Each cell to its sector's top regional share	9.4	[2.4, 18.1]	2.7	28%
<i>Within-cell hump, quadratic (cell, year FE)</i>				
Exporter share +5 pp, every cell	2.1	[−0.4, 4.7]	0.6	6%
Each cell to its sector's top regional share	2.2	[−1.3, 5.8]	0.6	7%
Exporter premium +10% (two-type, independent of the hump)	2.3	–	0.7	7%
Exact accounting, exporter margins 2006–14, per 5 pp	2.6			

Notes: Each 2018 region–sector cell is moved along the hump estimated across all cell-waves, with sector and year effects (cross-sectional) or cell and year effects (within-cell); national variance includes the between-cell term, with cell means shifted by $\varphi_c \Delta \rho_c$. The premium counterfactual uses the two-type approximation and does not depend on the hump estimate. Intervals from 500 cluster-bootstrap replications (region–sector clusters). $\Delta \ln(P_{90}/P_{10})$ assumes log-normality. The last column compares $|\Delta V|$ with the observed fall in national variance between 2006 and 2022 (−33.7). Rows take the estimated hump as structural; they are bounds implied by the model, not causal estimates.

hump estimated across cells were structural, moving each cell along it would give the effect of more exporting on national wage inequality, including the between-cell term, since cell means also shift when more workers move to higher-paying exporters. Table 5 reports such counterfactuals for 2018.

Two answers. Taken at face value, the cross-sectional hump implies large effects. Raising the exporter share by five percentage points in every cell, roughly what happened between 2006 and 2014, would raise the variance of log wages by 12.9 points ($\times 1000$; 90% interval 5.8–18.6), or about 3.6 log points of P_{90}/P_{10} . That is 38% of the fall in variance actually observed between 2006 and 2022. Moving every cell to the highest exporter share in its sector would raise it by 9.2. Estimated within cells over time, the hump implies a fifth of that: 2.1 points (−0.4–4.7), or 0.6 log points of P_{90}/P_{10} . The exact accounting of Section 7 gives the same order of magnitude: the exporter margins added 2.6 points per five-point rise in the exporter share between 2006 and 2014. Raising the exporter premium by 10% everywhere adds 2.3.

Reading the gap. The two time-series estimates agree with each other and disagree with the cross-section by a factor of five. The gap measures how much of the cross-sectional hump reflects persistent differences between region–sector cells—in technology, firm size or bargaining—that are correlated with exporting but would not move with it. The paper cannot say which estimate is causal. It can say that the range is wide but bounded: a five-point rise in exporting, as large as the one Spain actually experienced, raises the variance of hourly wages by between about 2 and 13 points, against an observed fall of 33.7 points over 2006–2022 and a composition effect of 19 points in the bust alone. Even at the top of the range, exporting is not

what moved Spanish wage inequality. At the bottom of the range, which the accounting and the within-cell estimates both support, it is a minor force.

9 The 2019 minimum wage

In January 2019 Spain raised its minimum wage by 22.3%, from 735.90 to 900 euros a month in 14 payments, and kept raising it to 2024. The national floor bites very differently across regions and sectors. Following Card (1992), I measure the bite of each region–sector cell as the share of its 2018 workers earning less than the 2019 minimum, 6.03 euros an hour; the average is 10.4%. Proposition 5(b) predicts that cells with a larger bite compress their lower tail more, and that the upper tail does not respond.

Table 6 and Figure 3 confirm both predictions. With sector and region fixed effects, the coefficient of $\Delta \ln(P50/P10)$ on the bite is -0.66 (s.e. 0.17): a cell with 10 percentage points more of its workers below the new floor compressed its lower tail by about 6.6 log points more. The upper tail does not respond (0.20, s.e. 0.28). This is the sharpest result in the paper. It is consistent with the large within-type compression between 2018 and 2022 in Table 4, although that period also saw the 2021 labour reform cut temporary employment, and the accounting does not separate the two.

The employment coefficient is -0.27 (s.e. 0.19): negative but imprecise, and the EES is not designed to measure employment. Proposition 5(d) offers a way to recover the employment loss from wage data alone, since the shift in the median of the observed distribution identifies $eF(m)$ once the pre-increase distribution is known. Implementing it requires separating the effect of the floor from general wage growth between 2018 and 2022, for example with the annual EES tables or Social Security records, and is left for future work.

The minimum wage also interacts with the hump. It binds almost entirely among non-exporters: in 2018, 11.6% of non-exporter employees but only 4.0% of exporter employees earned less than the 2019 floor, so 94% of those affected worked for non-exporters. Because of this, it lowers their dispersion and narrows the exporter premium (Proposition 5(e)). The level of inequality falls, but the peak of the hump moves to the right, so the range over which exporting raises inequality widens. Testing this requires more variation in exporter shares over time than four years of survey data provide.

TABLE 6. The 2019 minimum-wage increase and the wage distribution, 2018–2022

	No FE	Sector FE	Sector + region FE
$\Delta \ln(P50/P10)$	−0.446*** (0.070)	−0.601*** (0.141)	−0.659*** (0.175)
$\Delta \ln(P90/P50)$	−0.067 (0.111)	−0.154 (0.192)	0.198 (0.277)
$\Delta \ln \text{employment}$	0.009 (0.129)	−0.079 (0.154)	−0.274 (0.192)

Notes: Coefficient on the share of 2018 workers in the cell earning less than the 2019 minimum wage (6.03 euros per hour). 116 region–sector cells, 2018 employment weights, standard errors clustered by cell. Employment is the weighted EES headcount and is only indicative.




FIGURE 3. Bite of the 2019 minimum wage and compression of the lower tail across region–sector cells. The line has the slope of the employment-weighted estimate without fixed effects in Table 6, drawn through the origin for reference.

10 Discussion

What remains of the thesis. The 2021 master’s thesis regressed a dispersion index of wages (σ^2/μ) on trade openness, $(X + M)/\text{GDP}$, across Spain’s 17 regions over 2004–2018, with regional fixed effects. It found that openness raised inequality, marginally; that unemployment raised it substantially; and, puzzlingly, that a larger manufacturing share lowered it. The model and the microdata reinterpret each result.

Thesis result	Reading in this paper
Openness raises inequality, marginally	Most cells are on the rising side of the hump (Prop. 3); the exact accounting puts exporting at 1–1.5 points of variance per period. Openness measured relative to GDP also rises mechanically when GDP falls.
Unemployment raises inequality	Composition (Prop. 4): the jobs destroyed were homogeneous mid-distribution jobs in construction. The wage structure compressed more.
Manufacturing lowers inequality	Manufacturing cells have exporter shares near or beyond the peak of the hump, where more exporting no longer raises inequality.
Infrastructure raises inequality	No role in the model; dropped.

The thesis also used annual regional wage data for 2004–2007 that do not exist in the official

survey, and conflated inequality within regions with inequality between them (Rodríguez-Pose, 2012). The present paper uses the survey's microdata, defines the object of interest as the within-cell wage distribution, and derives the relationship it estimates.

Limitations. Five limitations matter. (i) The survey identifies a firm's main market, not whether it exports; this understates exporter shares and biases the location of the hump to the left. (ii) Regions are NUTS-1 aggregates, so cells are large and few. (iii) The evidence for the hump is cross-sectional; within cells over time it is uninformative. (iv) None of the three designs delivers a strong first stage, so the causal effect of exporting on inequality is not identified here; Proposition 3 is supported by the cross-section and by the sign pattern of the weak regional design, and its quantitative importance is bounded by the model rather than estimated. (v) The minimum-wage design compares cells with different bites but cannot rule out that low-wage cells were on different trends for other reasons, such as the 2021 reform of temporary contracts, which also concentrated on low-wage jobs.

Next steps. Matched employer–employee data from Social Security records (the *Muestra Continua de Vidas Laborales*) linked to customs records would identify exporters directly, allow AKM firm effects, and give provinces rather than NUTS-1 regions. With firm-level exporter status, the venting-out design of Section 6 could be run at the level of firms and provinces, where exposure to construction varies far more than across seven regions and one sector. And Appendix C shows that the real incidence of regional trade shocks depends on the housing-supply elasticity, which links the wage effects studied here to the housing market studied in Working Paper No. 1 (Menéndez-Pidal, 2026).

11 Conclusion

Whether trade raises or lowers wage inequality within a local labour market depends on how many of its workers are employed by exporters. With heterogeneous firms, inequality is hump-shaped in that share for any productivity distribution, the location of the peak is informative about the tail of the distribution, and foreign demand raises inequality everywhere up to and beyond the peak. In Spain's wage microdata the hump is there, in the cross-section of region–sector cells, with evidence that exporters occupy a Pareto tail of the productivity distribution, and most employment is on its rising side.

That does not make trade the main force behind Spanish wage inequality. Hourly wage inequality in the private sector fell between 2006 and 2022. The 2008 bust widened it through the composition of employment, as the destruction of mid-wage construction jobs predicts, but compressed it more through the wage structure. The 2019 minimum-wage increase compressed the lower tail in proportion to its bite. Exporting added a little inequality in each period. How much more it could add is bounded: a five-point rise in the exporter share raises the variance of log wages by between 2 and 13 points, against a fall of 34 points over the period. That is the conclusion of the thesis this paper replaces, but now with a mechanism, a measure, bounds, and an account of what else was going on.

Next steps. Three additions would turn the bounds into estimates. Firm-level export status, from the Central de Balances or commercial registers, would correct the survey's understatement of exporting and remove the main bias in the location of the peak. Destination-specific shocks give sharper first stages than aggregate EU demand: the 2016 Brexit vote for cells exposed to the United Kingdom, the 2018 US tariffs on Spanish olives and steel, and the 2014 Russian embargo on food imports for exposed agricultural sectors. And the size distribution of Spanish firms would give an independent estimate of the tail that the peak can be tested against. The natural outlet for that version is a regional or Spanish journal such as *SERIEs* or *Papers in Regional Science*.

References

- John M. Abowd, Francis Kramarz, and David N. Margolis. High wage workers and high wage firms. *Econometrica*, 67(2):251–333, 1999.
- Rodrigo Adão, Michal Kolesár, and Eduardo Morales. Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010, 2019.
- Miguel Almunia, Pol Antràs, David Lopez-Rodriguez, and Eduardo Morales. Venting out: Exports during a domestic slump. *American Economic Review*, 111(11):3611–3662, 2021.
- Mary Amiti and Donald R. Davis. Trade, firms, and wages: Theory and evidence. *Review of Economic Studies*, 79(1):1–36, 2012.
- David H. Autor, David Dorn, and Gordon H. Hanson. The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.
- Robert L. Axtell. Zipf distribution of U.S. firm sizes. *Science*, 293(5536):1818–1820, 2001.
- Samuel Bentolila, Pierre Cahuc, Juan J. Dolado, and Thomas Le Barbanchon. Two-tier labour markets in the great recession: France versus Spain. *Economic Journal*, 122(562):F155–F187, 2012.
- Andrew B. Bernard and J. Bradford Jensen. Exporters, jobs, and wages in U.S. manufacturing: 1976–1987. *Brookings Papers on Economic Activity. Microeconomics*, pages 67–119, 1995.
- Andrew B. Bernard, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. Firms in international trade. *Journal of Economic Perspectives*, 21(3):105–130, 2007.
- Stéphane Bonhomme and Laura Hospido. The cycle of earnings inequality: Evidence from Spanish social security data. *Economic Journal*, 127(603):1244–1278, 2017.
- Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213, 2022.
- David Card. Using regional variation in wages to measure the effects of the federal minimum wage. *Industrial and Labor Relations Review*, 46(1):22–37, 1992.

- David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics*, 128(3):967–1015, 2013.
- David Card, Ana Rute Cardoso, Jörg Heining, and Patrick Kline. Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70, 2018.
- Doruk Cengiz, Arindrajit Dube, Attila Lindner, and Ben Zipperer. The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454, 2019.
- John DiNardo, Nicole M. Fortin, and Thomas Lemieux. Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach. *Econometrica*, 64(5):1001–1044, 1996.
- Vicente Donoso, Víctor Martín, and Asier Minondo. Do differences in the exposure to Chinese imports lead to differences in local labour market outcomes? An analysis for Spanish provinces. *Regional Studies*, 49(10):1746–1764, 2015.
- Hartmut Egger and Udo Kreickemeier. Firm heterogeneity and the labor market effects of trade liberalization. *International Economic Review*, 50(1):187–216, 2009.
- Jesús Fernández-Villaverde, Luis Garicano, and Tano Santos. Political credit cycles: The case of the Eurozone. *Journal of Economic Perspectives*, 27(3):145–166, 2013.
- Judith A. Frías, David S. Kaplan, and Eric Verhoogen. Exports and within-plant wage distributions: Evidence from Mexico. *American Economic Review: Papers & Proceedings*, 102(3):435–440, 2012.
- Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.
- Keith Head, Thierry Mayer, and Mathias Thoenig. Welfare and trade without Pareto. *American Economic Review: Papers & Proceedings*, 104(5):310–316, 2014.
- Elhanan Helpman, Oleg Itskhoki, and Stephen Redding. Inequality and unemployment in a global economy. *Econometrica*, 78(4):1239–1283, 2010.
- Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J. Redding. Trade and inequality: From theory to estimation. *Review of Economic Studies*, 84(1):357–405, 2017.
- David S. Lee. Wage inequality in the United States during the 1980s: Rising dispersion or falling minimum wage? *Quarterly Journal of Economics*, 114(3):977–1023, 1999.
- Jorge Menéndez-Pidal. The incidence of population shocks in local housing markets: Housing supply elasticity, migration and spatial adjustment. Working Paper 1, Project Frontier, 2026.
- Enrico Moretti. Local labor markets. In David Card and Orley Ashenfelter, editors, *Handbook of Labor Economics*, volume 4B, pages 1237–1313. Elsevier, Amsterdam, 2011.
- Sergey Nigai. A tale of two tails: Productivity distribution and the gains from trade. *Journal of International Economics*, 104:44–62, 2017.

- Jennifer Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90(6): 1257–1278, 1982.
- Andrés Rodríguez-Pose. Trade and regional inequality. *Economic Geography*, 88(2):109–136, 2012.
- Sherwin Rosen. Wage-based indexes of urban quality of life. In Peter Mieszkowski and Mahlon Straszheim, editors, *Current Issues in Urban Economics*, pages 74–104. Johns Hopkins University Press, Baltimore, 1979.
- Stephanie Schmitt-Grohé and Martín Uribe. Downward nominal wage rigidity, currency pegs, and involuntary unemployment. *Journal of Political Economy*, 124(5):1466–1514, 2016.
- Eric A. Verhoogen. Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector. *Quarterly Journal of Economics*, 123(2):489–530, 2008.

A Proofs

All statements below are also verified symbolically and numerically in the replication file `01_model_checks.py`.

Proposition 1. From (1), with a independent of (x, X) , $V = \sigma_a^2 + \lambda^2 \text{Var}(x) + \varphi^2 \text{Var}(X) + 2\lambda\varphi \text{Cov}(x, X)$. $\text{Var}(X) = \rho(1-\rho)$ and, by the law of total covariance, $\text{Cov}(x, X) = \rho(1-\rho)(\mathbb{E}[x | X = 1] - \mathbb{E}[x | X = 0]) = \rho(1-\rho)\Delta(\rho)$. Since X selects the upper tail, $\Delta > 0$. As $\rho \rightarrow 0$, $\rho\Delta(\rho) = \rho(\mathbb{E}[x | x \geq x_0] - \mathbb{E}[x])/(1-\rho) \rightarrow 0$ whenever $\mathbb{E}[x] < \infty$, and symmetrically as $\rho \rightarrow 1$. With $\lambda = 0$ and type-specific variances, the law of total variance gives $V = \rho\sigma_X^2 + (1-\rho)\sigma_N^2 + \rho(1-\rho)\varphi^2$, whose derivative $\sigma_X^2 - \sigma_N^2 + \varphi^2(1-2\rho)$ vanishes at $\rho^* = \frac{1}{2} + (\sigma_X^2 - \sigma_N^2)/(2\varphi^2)$. $\square$

Proposition 2. If $x \sim \text{Exp}(\tilde{k})$, $\rho = e^{-\tilde{k}x_0}$ and $\mathbb{E}[x | x \geq x_0] = x_0 + 1/\tilde{k}$ by memorylessness, so $\text{Cov}(x, X) = \rho(\mathbb{E}[x | X = 1] - \mathbb{E}[x]) = \rho x_0 = -\rho \ln \rho / \tilde{k}$, and $\text{Var}(x) = 1/\tilde{k}^2$. This gives (3). Differentiating, $V' = \varphi^2(1-2\rho) + (2\lambda\varphi/\tilde{k})(-\ln \rho - 1)$ and $V'' = -2\varphi^2 - 2\lambda\varphi/(\tilde{k}\rho) < 0$. At $\rho = 1/e$ the second term of V' vanishes and $V' = \varphi^2(1-2/e) > 0$; at $\rho = 1/2$ the first vanishes and $V' = (2\lambda\varphi/\tilde{k})(\ln 2 - 1) < 0$. By strict concavity the unique maximum lies in $(1/e, 1/2)$. If $x \sim N(\mu, s^2)$, $\text{Cov}(x, X) = s \phi(z)$ with $z = \Phi^{-1}(1-\rho)$, where ϕ is the standard normal density; $d\phi/dz = -z\phi$ vanishes at $z = 0$, i.e. $\rho = 1/2$, and $\varphi^2\rho(1-\rho)$ also peaks there. $\square$

Remark on body and tail. If x has a Pareto tail above $\bar{x}$ with mass $\bar{\rho}$ and $\rho \leq \bar{\rho}$, exporters are in the tail and $\text{Var}(x | X = 1) = 1/\tilde{k}^2$ independently of ρ . $\mathbb{E}[x | X = 1] = x_0 + 1/\tilde{k}$ with $x_0 = \bar{x} - \ln(\rho/\bar{\rho})/\tilde{k}$, so $\text{Cov}(x, X) = \rho(\mathbb{E}[x | X = 1] - \mathbb{E}[x]) = -\rho \ln \rho / \tilde{k} + \kappa\rho$, with $\kappa = \bar{x} + (1 + \ln \bar{\rho})/\tilde{k} - \mathbb{E}[x]$ a constant determined by the body. $\square$

Proposition 3. Holding F fixed, $\ln A^*$ moves ρ (through x_0) and φ . By the chain rule $dV/d \ln A^* = V_\rho \zeta\rho + V_\varphi \Lambda$. From (2), $V_\rho = 2\varphi\rho(1-\rho) + 2\lambda \text{Cov}(x, X) = 2\rho(1-\rho)(\varphi + \lambda\Delta) > 0$, which proves (a). At ρ^* , $V_\rho = 0$ and $dV/d \ln A^* = \Lambda V_\varphi > 0$; since the extensive term is positive below ρ^* , any zero of $dV/d \ln A^*$ lies above it, proving (b). Under a Pareto tail, substituting

(3) gives $\rho H(\rho)$ with H as stated; $H' = -(\Lambda + \varphi\zeta)(c + 2\varphi\rho)/\rho < 0$, $H(\rho) \rightarrow +\infty$ as $\rho \rightarrow 0$ because of the $-\ln \rho$ terms, and $H(1) = -\zeta\varphi(\varphi + c) < 0$, so the zero is unique, proving (c). Finally $E = \rho H$ satisfies $E'' = 2H' + \rho H'' = -(\Lambda + \varphi\zeta)(c/\rho + 4\varphi) < 0$, $E(0^+) = 0$ and $E'(0^+) = H(0^+) + \lim \rho H' = +\infty$, proving (d). With $c = 0$ the zero of H is $\tilde{\rho} = \omega/2 + (1 - \omega)$, $\omega = \zeta\varphi/(\zeta\varphi + \Lambda)$, the two-type benchmark. $\square$

Proposition 4. Write the pre-shock distribution as a mixture of surviving jobs (weight $1 - d$, mean m_R , variance V_R) and destroyed jobs (weight d , mean m_D , variance V_D): $V = (1 - d)V_R + dV_D + d(1 - d)(m_R - m_D)^2$ and $m = (1 - d)m_R + dm_D$. Substituting $m_R - m_D = (m - m_D)/(1 - d)$ and solving for V_R gives (6). $\square$

Rigidity and duality. With labour demand $\ln L = \ln D - \varepsilon \ln w$ and supply $\ln L = \nu \ln w$, a shock $d \ln D = -x$ lowers the flexible wage by $x/(\varepsilon + \nu)$. If wages can fall by at most $\omega < x/(\varepsilon + \nu)$, employment is demand-determined, $d \ln L = -x + \varepsilon\omega$, and rationing is $x - (\varepsilon + \nu)\omega$; a positive shock of size x moves the wage by $x/(\varepsilon + \nu)$. With firing costs on permanent contracts only, separations fall on temporary workers first, at rate $\min\{1, d/\tau\}$. $\square$

Proposition 5. After the floor, the employed mass is $1 - eF(m)$, with a mass $(1 - e)F(m)$ at m and the latent density above m . (a) The observed q -quantile equals m iff $q(1 - eF(m)) \leq (1 - e)F(m)$; with $q = 0.1$ this is $F(m) \geq 0.1/(1 - 0.9e)$. (d) For $y > m$ the observed CDF is $(F(y) - eF(m))/(1 - eF(m))$; setting it to $\frac{1}{2}$ gives $F(y) = \frac{1}{2} + \frac{1}{2}eF(m)$. (b) and (c) follow from (a) and (d): when the floor pins $P10$, a higher bite raises $P10$ to m and disemployment raises only the median; when it does not, disemployment truncates the bottom and raises $P10$ more than the median. These are verified numerically for log-normal latent wages over 150 parameter draws. (e) Truncating a log-concave distribution from below lowers its variance and raises its mean; for normal latent log wages the same holds for the mixture of censoring and truncation that the floor induces (verified numerically over 60 draws). If the floor binds only among non-exporters, σ_N^2 falls and φ falls, so $(\sigma_X^2 - \sigma_N^2)/(2\varphi^2)$ rises whenever $\sigma_X^2 \geq \sigma_N^2$. $\square$

Equation (7). The two-type variance $V = \rho\sigma_X^2 + (1 - \rho)\sigma_N^2 + \rho(1 - \rho)\varphi^2$ is a sum of products; the Shapley decomposition of each product ab is $\Delta a \bar{b} + \bar{a} \Delta b$, which is exact. $\square$

B Data construction

Sources. INE, *Encuesta de Estructura Salarial*, microdata for 2006, 2010, 2014, 2018 and 2022 (public FTP, ine.es/ftp/microdatos/salarial). INE, EES annual tables by region (table 28191, used only for context). Eurostat Comext, dataset DS-045409, imports by HS4 of each EU member from the world and from Spain, and Spain's exports by HS4. Agencia Tributaria, monthly territorial foreign-trade files, 2016–2018 and 2022 (exports by province of origin and CN8). Replication code: `code/02_download_ees.py` to `code/08_tables.py`.

Sample. Private-sector employees (CONTROL=2) paid for the whole of October (2006: AFECMES=6; 2010–2022: SIESPM2=6 and 31 days of employment in October). Public administration (CNAE-

2009 section O, absent in 2006) is excluded. Hourly wages below the 0.1st and above the 99.9th percentile of each wave are dropped.

Harmonised sectors.

Sector	CNAE-93 groups (2006)	CNAE-2009 groups (2010–2022)
Food, textiles, leather	DA, DB, DC	C1
Wood, paper, printing	DD, DE	C2, C3
Chemicals, refining, rubber, plastics	DF, DG, DH	C4
Non-metallic minerals	DI	C5
Metals	DJ	C6
Machinery, electrical, electronic	DK, DL	C7
Transport equipment, furniture, other	DM, DN	C8
Mining, energy, water	CA, E	B, D, E
Construction	F	F
Trade	G	G1, G2
Hospitality	H	I
Transport and communications	I	H1, H2
Finance	J	K
Business services	K	J, L, M, N
Education	M	P
Health and social work	N	Q
Other services	O	R, S

The 2006 and later classifications differ most in information and communication (telecommunications were in transport in CNAE-93 and are in J in CNAE-2009); this affects two service sectors with almost no exporters.

Products to sectors. HS chapters 15–24, 41–43 and 50–67 to food, textiles and leather; 44–49 to wood, paper and printing; 27–40 to chemicals; 68–70 to minerals; 71–83 to metals; 84–85 and 90–92 to machinery; 86–89 and 93–96 to transport equipment and other manufacturing; 25–26 to mining. Chapters 01–14 (agriculture) and 97–99 are excluded.

Regions. NUTS-1: Northwest (Galicia, Asturias, Cantabria); Northeast (Basque Country, Navarre, La Rioja, Aragon); Madrid; Centre (Castile and León, Castile-La Mancha, Extremadura); East (Catalonia, Valencia, Balearic Islands); South (Andalusia, Murcia, Ceuta, Melilla); Canary Islands. AEAT provinces of origin are mapped accordingly.

Foreign demand. For product k , $M_k = \sum_j (M_{jk}^{\text{world}} - M_{jk}^{\text{ES}})$ over the EU members other than Spain that report in every year used. Growth rates are log differences bounded at ± 2 to limit the influence of tiny product lines; the shares give such lines negligible weight in any case.

C Spatial extension: real-wage incidence

The model holds the population of each cell fixed. When a regional shock raises local wages, workers move and housing costs respond, so the nominal gains measured in Section 6 overstate

real gains. This appendix gives the incidence in a linear spatial equilibrium in the tradition of Rosen (1979), Roback (1982) and Moretti (2011), and links it to the housing-supply elasticity estimated in Project Frontier Working Paper No. 1.

A region receives a labour-demand shock a (in logs). Local wages satisfy $d \ln w = a - d \ln N / \varepsilon_L$; housing, with expenditure share α and supply elasticity η , clears at $d \ln p = (d \ln N + d \ln w) / (1 + \eta)$; and mobility responds to real wages with elasticity κ , $d \ln N = \kappa(d \ln w - \alpha d \ln p)$.

Proposition 6 (Real-wage incidence). *The share of the shock received by workers as a real wage gain is*

$$\Gamma \equiv \frac{d \ln w - \alpha d \ln p}{a} = \frac{\varepsilon_L(1 + \eta - \alpha)}{(1 + \eta)(\varepsilon_L + \kappa) + \alpha\kappa(\varepsilon_L - 1)} \in (0, 1), \quad (9)$$

increasing in η and decreasing in κ , with $\Gamma \rightarrow \varepsilon_L / (\varepsilon_L + \kappa)$ as $\eta \rightarrow \infty$. Landlords receive $\alpha d \ln p > 0$. With two skill groups of mobility κ_H and κ_L , the high-skill share of the region rises if and only if $\kappa_H(d \ln w_H - \alpha d \ln p) > \kappa_L(d \ln w_L - \alpha d \ln p)$; the housing term $-(\kappa_H - \kappa_L)\alpha d \ln p$ pushes out the more mobile group first.

Working Paper No. 1 estimates that Spain's housing-supply elasticity fell from about 2 in 1998–2008 to about 0.45 in 2015–2026. By the proposition, the export-led recovery after 2013 in the most open regions should have translated into real wage gains less than the earlier booms did, and into higher housing costs more. If so, nominal wage comparisons across regions overstate real ones in exactly the period when exports grew most. Measuring this requires regional price levels, which the present paper does not use.