

Exports and Wages: Rent Sharing, Workforce Composition, or Returns to Skills?

Mario Macis, *Johns Hopkins University, IZA, and NBER*

Fabiano Schivardi, *Bocconi University, IGER, EIEF, and CEPR*

We use linked employer-employee data from Italy to explore the relationship between exports and wages. Exploiting the 1992 devaluation of the lira, we show that exporting firms both pay a wage premium above what their workers would earn in the outside labor market (the “rent-sharing” effect) and employ workers whose skills command a higher price after the devaluation (the “skill composition” effect). The latter only emerges once we allow for the value of workers’ skills to differ in the pre- and post-devaluation periods. We also document that the export wage premium is larger for workers with more export-related experience.

I. Introduction

Since the seminal paper of Bernard, Jensen, and Lawrence (1995), a growing body of empirical work has shown that exporting firms pay higher wages

We thank Emanuele Breda, Paula Bustos, Tom Holmes, Mark Roberts, Jagadeesh Sivadasan, Christopher Stanton, Jonathan Vogel, and seminar participants at the World Bank, the CEA at the Universidad de Chile, the ERWIT workshop, the ITSG meeting in Cagliari, and the University of Michigan Ross School of Business for useful comments and suggestions. Macis gratefully acknowledges financial support from the Mitsui Life Financial Research Center and Schivardi by the Regione Sardegna, Legge 7/2007 grant CRP 26133. All errors are our own. Contact the corresponding author, Mario Macis, at mmacis@jhu.edu. Information concerning access to the data used in this article is available as supplementary material online.

[*Journal of Labor Economics*, 2016, vol. 34, no. 4]

© 2016 by The University of Chicago. All rights reserved. 0734-306X/2016/3404-0004\$10.00

Submitted August 14, 2013; Accepted March 23, 2015; Electronically published August 2, 2016

even after controlling for firm-level characteristics such as industry and size.¹ The ensuing theoretical literature has proposed two possible mechanisms. On one side, exporting firms might employ workers with higher skills, in which case the exporting wage premium is a reflection of observable and unobservable workers' characteristics—the “skill composition” effect (Yeaple 2005; Verhoogen 2008; Bustos 2011; Kugler and Verhoogen 2012). On the other side, the presence of frictions in the labor market might lead exporting firms to pay higher wages than nonexporting firms for identical workers because exporting generates rents that are shared with the employees—the “rent sharing” effect (Helpman and Itskhoki 2010; Cosar, Guner, and Tybout 2016).² Although these theoretical mechanisms are well understood, identifying their relative importance empirically has proven difficult. Traditional studies using average wages at the firm level cannot fully control for workers' skills and therefore cannot distinguish between composition or rent sharing factors. In the last few years, the literature has taken advantage of the growing availability of matched employer-employee data to address these issues, but the evidence is still not conclusive.³

In this paper, we use a matched employer-employee database including the entire workforce of a large sample of Italian manufacturing firms to study the effects of exporting on wages at the firm level. We add to the literature along at least three dimensions. First, as done by Frías et al. (2009) for Mexico and Araújo and Paz (2014) for Brazil, we exploit the sudden and large devaluation of the Italian lira in 1992 as a source of exogenous variation in the firms' incentive to export. Second, we propose an empirical framework that allows the market value of individual workers' observable and unobservable skills to vary before and after the devaluation. As we show, this is a crucial step in disentangling rent sharing from skill composition effects. Third, we document the heterogeneous effects of exporting on wages based on a measure of workers' export-specific experience.

The linked employer-employee nature of our data allows us to apply the methodology developed by Abowd, Kramarz, and Margolis (1999) to

¹ For comprehensive surveys, see Schank, Schnabel, and Wagner (2007) and Frías, Kaplan, and Verhoogen (2009).

² In addition to specificities and frictions, there are at least two other mechanisms that could imply that the firm shares the rents from exporting with their workers: efficiency wages (Frías et al. 2009; Davis and Harrigan 2011) and fair-wage considerations (Egger and Kreckemeier 2009; Amiti and Davis 2011). See Frías et al. (2009) for a more detailed survey of the various theoretical mechanisms behind export wage effects.

³ See, among others, Frías et al. (2009) and Frias, Kaplan, and Verhoogen (2012) for Mexico; Schank et al. (2007) and Baumgarten (2013) for Germany; Krishna, Poole, and Senses (2011a) and Helpman et al. (forthcoming) for Brazil. We discuss how our results compare with those of the existing literature below.

decompose individual wages into a component due to observable time-varying worker and firm characteristics, a component due to unobservable worker characteristics (“worker effects”), and a component due to firm-level unobservable characteristics (“firm effects”). As in Frías et al. (2009) and Helpman et al. (forthcoming), the firm effect is estimated separately for each firm-year, so that it can be related to changes in the firm’s export activity. The worker effect is used as a measure of workers’ unobservable (to the econometrician) skills. This procedure offers a natural decomposition of wages into rent sharing, the firm effect, and skill composition, the worker effect.

Most of the literature on export and wages treats the worker effects as being fixed over time. This implies that the skill composition can only change with changes in the workforce composition. However, because they are derived from a wage equation, the estimated worker effects measure the market value of skills rather than simply the skills themselves. This might change during a major economic episode such as a devaluation. In fact, the devaluation represented a large shock to the incentives to export and might have affected the market price of skills according to their contribution to export activity. Ideally, one would like to separate the pure skill component from the market value component. Unfortunately, by construction, the worker effects capture the joint effect of the two. To address this issue, we estimate worker and firm effects under two alternative assumptions. First, as has typically been done in the literature, we assume that the worker effects are fixed throughout the entire period. Second, we allow the individual worker effects to vary in the pre- and post-devaluation periods. By estimating separate worker effects pre- and post-devaluation, we are able to account for any change in the market value of skills at the individual level that, if not properly accounted for, would end up in the firm effect.

The Abowd-Kramarz-Margolis wage decomposition rests on an assumption of exogenous worker mobility conditional on observables. In their paper on workplace heterogeneity and wage inequality in Germany, Card, Heining, and Kline (2013) discuss several possible violations of the exogenous mobility assumption and suggest a series of tests to detect such violations. We performed the Card-Heining-Kline tests with our data, finding very similar results. We conclude that the Abowd-Kramarz-Margolis assumption is a reasonable representation of the mobility process and that the Abowd-Kramarz-Margolis model’s additively separable worker and firm effects measure to a good approximation the unobservable worker and firm components of wages in our context.

Our empirical framework is based on regressing workers’ wages and components (skill composition and rent sharing) on the share of export at the firm level in the post-devaluation period, accounting for the potential endogeneity of the export share (we discuss this issue in more detail below). Our results indicate that the increased export activity that followed the un-

expected and large devaluation of the Italian currency in 1992 led to higher wages. Our estimates imply that, other things equal, wages rose by 1.05%–1.30% (on average) for firms recording the median increase in the export share (15%). In terms of rent sharing or skill composition, we find that, when skills are assumed to be fixed throughout the period, the whole effect of the increase in exports is due to the rent sharing component. This indicates that the characteristics of the workforce have not changed systematically in relation to the export activity that firms undertook after the devaluation. In fact, when the worker effects are fixed, only changes in the workforce composition can change the skill composition at the firm level. However, when we allow the worker effects to vary before and after the devaluation, we find that the higher wages are roughly equally due to an increased firm effect, common to the entire firm workforce, and to an increase in the worker effect. Given that we found no evidence of changes in the skill composition when keeping the worker effects fixed, the increased worker effects in exporting firms must reflect an increase in the market value of skills of the workers employed by such firms. We conclude that exporting firms do share rents with their workers, which is consistent with recent models that emphasize firm heterogeneity and labor market frictions (Helpman and Itskhoki 2010). We also conclude that the market value of unobservable skills of workers employed by exporting firms increases after the devaluation. Failure to take this change into account would lead one to overestimate the rent sharing component.

To corroborate our interpretation of the results, we explore whether the export wage premium associated with the devaluation can be linked to a measure of export-specific workers' skills. We assume that past experience in exporting firms increases the level of a worker's export-specific skill, and we find that, indeed, the export effect is significantly stronger for workers with more past export experience. This result, which is robust to including an extensive set of tenure controls in the regressions, indicates that there is heterogeneity across workers in the distribution of skills in terms of how useful they are for the export activity and that the devaluation increased the demand for export-specific skills, driving their relative price up.

The devaluation generated large changes in export activity at the firm level. However, even in the context of an exogenous change in the incentive to export, one might be concerned that the most productive firms might also be those that are better equipped to take advantage of the devaluation. If this is the case, then a measured "effect" of increased export activity on wages might simply reflect the underlying heterogeneity, which could have generated both greater exports and higher wages. We have taken several steps to lessen this concern. First, we argue and provide direct evidence that in the Italian case, the concern that exporters are primarily the most productive firms is much less relevant than in other contexts. Second, in our empirical

specification, we explicitly control for predetermined conditions at the firm level. More specifically, our proposed specification allows for wages in the devaluation period to be correlated with the pre-devaluation export intensity. This allows us to establish whether the changes in wages (or wage components) that took place in the devaluation period were due to the increased export activity or simply to pre-existing heterogeneity. Finally, our results are very robust to inclusion of firm fixed effects, which control for unobservable, time-invariant firm heterogeneity.

Our results confirm the existence of a rent sharing component in the wage premium paid by exporters. Indeed, this component remains even when the worker effects are allowed to differ before and after the devaluation. At the same time, once we do that, the rent sharing component decreases substantially. This suggests that the episodes typically studied in the literature, such as devaluations or trade agreements that reduce trade barriers, will have general equilibrium effects that change the relative prices of skills. Failing to take this into account might bias the estimated rent sharing component of the wage premium.

Our paper contributes to a small but growing body of empirical literature that uses matched employer-employee data to study the relationship between exporting activity and workers' wages. Most closely related to our research is that of Frías et al. (2009), who use matched employer-employee data from Mexico and exploit the 1994 devaluation of the peso as an exogenous shock to Mexican firms' incentives to export. Our paper too uses matched employer-employee data and a large and unexpected devaluation as a source of identification. In contrast to Frías et al. (2009), however, we do not focus on differences based on firm observable characteristics; rather, we look directly at how wages relate to changes in export shares in the devaluation period compared to the earlier years. Our paper differs from theirs also in the way we allow for the devaluation to affect workers' effects. Whereas they account for potentially time-varying effects by interacting the worker individual effects with a time effect, we estimate separate worker effects for the pre- and post-devaluation periods. When we constrain the worker effects to be fixed over time, we also find that rent sharing explains the bulk of the export wage premium. Krishna et al. (2011a) use matched employer-employee data from Brazil, and they find that, when firm-worker match controls are included in the regressions, the effect of trade openness on wages at exporting firms compared to domestic firms vanishes. In their paper, however, the firm-worker match is also fixed over time, and it is not allowed to vary with export activity. Araújo and Paz (2014) exploit a devaluation episode in Brazil to study the effect of exporting on wages; however, in their empirical specification both worker effects and match effects are again assumed to be fixed over time. Helpman et al. (forthcoming) also use linked employer-employee data from Brazil. They estimate firm-occupation-year

effects including both wage premia and unobserved worker heterogeneity. Helpman et al. (forthcoming) build a structural model of trade with heterogeneous firms to estimate the role of trade in determining wage dispersion within occupation; by contrast, we focus on disentangling skill composition from rent sharing effects.

Finally, our work includes a novel exploration of the heterogeneous effects of trade. Much of the existing literature has focused on the differential effects of trade across groups of workers, typically defined by education, occupation (blue collar, white collar, managers), and industry. However, these traditional categories are very broad, and they potentially mask substantial within-group heterogeneity. A recent exception is represented by Hummels et al. (2014), who document the heterogeneous effects of trade on workers who perform different sets of tasks (e.g., creative vs. routine tasks) or whose occupations employ different sets of knowledge (e.g., mathematics, social science, engineering). Frias et al. (2012) also contribute to the analysis of the heterogeneous effects of trade by looking at different percentiles of the within-firm wage distribution. In our paper, we document heterogeneous wage effects of exporting based on an explicit measure of export-related skills rather than use occupational categories or wage levels. To the best of our knowledge, the only other paper that considers workers' export experience explicitly is Mion and Oromolla (2011), which finds that managers with previous export experience receive a wage premium and increase the likelihood that a firm engages in export activity.

The paper proceeds as follows. In Section II, we describe the data, perform the estimation of worker and firm effects, and test the exogenous mobility assumption. In Section III, we present our main econometric analysis of the effect of exporting on wages, workforce composition, and firm-level wage premia. In Section IV, we explore the heterogeneity of the export wage premia across workers, emphasizing the role of workers' past export experience. Finally, in Section V, we conclude and offer directions for future research.

II. Data, Wage Decomposition, and Descriptive Evidence

A. Data Description

The data used in this paper were constructed from the Bank of Italy's INVIND survey of manufacturing firms. INVIND is an open panel of around 1,200 firms per year, representative of manufacturing firms with at least 50 employees. It contains detailed information on firms' characteristics, including industrial sector, year of creation, average number of employees during the year, sales, investment, and, most important for our purposes, exports. The Italian Social Security Institute (Istituto Nazionale Previdenza Sociale, INPS) provided the complete work histories of all

workers who were ever employed by an INVIND firm in the period 1980–97, including spells of employment in which they were employed in firms not listed in the INVIND survey. Overall, we have information on about 1 million workers per year, more than half of whom are employed in INVIND firms in any given year. The rest are employed in about 500,000 other firms of which we only know the unit identifier.⁴

The information on workers includes age, gender, the province where the employee works, occupational status (production, nonproduction, manager), annual gross earnings (including irregular payments such as overtime, shift work, and bonuses), number of weeks worked, and the firm identifier. We have deleted records with missing entries on either the firm or the worker identifier, those corresponding to workers younger than 15 and older than 65, those who have worked less than 4 weeks in a year, and those in the first and last percentiles of the weekly earnings distribution.

In table 1, we report summary statistics on workers' characteristics for the entire sample (col. 1), which, as explained in Section II.B below, we use to estimate worker and firm effects, as well as for workers employed in INVIND firms (cols. 2 and 3), on which we base our main analysis. Because precise information on exporting behavior for a representative sample of firms is available only for INVIND firms after 1987, we will restrict our attention to INVIND firms in the period 1987–97. For the entire sample, average gross weekly earnings at 1995 constant prices are about €378, and the average age of workers is 37 years. Almost 80% of the observations pertain to males. Production workers account for 66% of the total and nonproduction workers for 33%. The INVIND sample in years 1987–97 consists of about 4.1 million observations. The descriptive statistics for the INVIND sample are quite similar to those of the total sample; this was expected, because this sample includes the same workers but only observations of those who were employed by an INVIND firm in the period 1987–97.

Table 2 reports statistics on the firm-year level data used in our main regression analyses. A total of 1,218 unique firms are included in the INVIND sample in the period considered. The sample is unbalanced. The median INVIND firm employs about 230 workers, and it reports sales of over €31 million. Eighty-nine percent of the firms in the sample were exporters in the period considered. Conditional on exporting, the median firm exports 31% of its sales. These figures are in line with those reported in other studies on Italian firms (Castellani, Serti, and Tomasi 2010; Crinò and Epifani 2012) and are substantially higher than those found in other countries. In the United States (Bernard et al. 1995) or Mexico (Frias et al. 2009), for instance, only a small proportion of firms do export. This difference is ex-

⁴ This is the same database used by Iranzo, Schivardi, and Tosetti (2008), which the reader is referred to for further details.

Table 1
Descriptive Statistics, Workers

	Entire Sample, 1980–97 (1)	Entire Sample, 1987–97 (2)	INVIND Sample, 1987–97 (3)	INVIND Stayers, 1987–97 (4)
Weekly wage	377.6 (160.2)	401.1 (182.0)	404.2 (168.5)	397.4 (165.5)
Age	36.9 (10.0)	37.6 (10.1)	38.8 (9.9)	40.0 (10.3)
Tenure	5.11 (4.12)	6.5 (4.6)	8.1 (4.7)	10.1 (4.5)
Males	.79	.79	.78	.74
Production workers	.66	.64	.67	.68
Nonproduction workers	.34	.36	.33	.31
No. of observations	18,635,710	11,042,916	4,074,074	2,772,321

NOTE.—Entire sample refers to all workers in the data set; INVIND sample only includes workers who are currently employed by a firm that belongs to the Bank of Italy–INVIND survey and with nonmissing information on export activity. See Sec. II.A for a description of the data. Standard deviations are in parentheses.

plained by at least two factors. First, Italy's main commercial partners are countries within the European Union (EU), which are located in relative geographic proximity. Second, the INVIND sample excludes firms with fewer than 50 employees, and it is a well-known fact that small firms have

Table 2
Descriptive Statistics, Firms

	All (1)	Exporters (2)	Nonexporters (3)
Employment:			
Mean	672.3	720.9	287.4
SD	3,207.9	3,397.4	473.9
Median	228	245	147
Sales:			
Mean	112,846.0	120,801.8	49,891.3
SD	539,967.5	570,717.2	130,019.8
Median	31,240.6	33,712.8	16,500.3
Export (0/1):	.89	1	0
Export share of sales:			
Mean	.31	.34	0
SD	.27	.27	0
Median	.25	.31	0
No. of observations	7,585	6,734	851

NOTE.—The sample includes firms in the INVIND sample in the period 1987–97 and with nonmissing information on export activity. Sales are expressed in thousands of euros (in constant 1995 prices).

a much lower propensity to export compared to larger firms. In fact, comparison of columns 2 and 3 in table 2 confirms that, even within the 50+ employee firms, exporters are substantially larger than nonexporters, both in terms of employment and sales volume.

The devaluation of 1992 was substantial and had a strong impact on exports. Figure 1, *A*, reports the multilateral real exchange rate of the lira (Finicelli, Luccardi, and Sbracia 2005). This is the best indicator for our purposes because it represents a measure of competitiveness of manufacturing goods. After the initial sharp drop, the exchange rate kept devaluating (with the exception of an appreciation in the second quarter of 1993) until mid-1995, when the depreciation, compared to January 1992, was on the order of 30%. After that, the currency recorded a stable appreciation, which, by the end of the decade, brought the multilateral exchange rate back to around 85% of the January 1992 level. Figure 1, *B*, shows that the exporting behavior of the INVIND firms changed after 1992—the year of the lira devaluation. The median (mean) share of sales exported increased sharply from around 18% (26%) in the 1987–91 period to 30% (34%) in 1997. Interestingly, the share starts decreasing in 1997, arguably indicating the fading away of the competitive advantage. Indeed, this decrease further supports the view that the sharp rise in exports was linked to the devaluation itself and not to some other concomitant factor, such as the European single market, whose effects should have been permanent.

B. Decomposing Wages into “Skills Component” and “Rent Sharing”

1. Wage Decomposition

Our goal is to establish whether export intensity leads to higher wages, and, if this is the case, if the higher wages simply reflect the skill composition of the workforce, including unobservable skills, or also rent sharing, defined as the excess wage that a worker obtains from working in a given firm compared to the market value of his or her skills. We exploit the matched employer-employee nature of our data to perform a decomposition of wages into two terms that capture the two potential sources of the positive correlation between exports and wages. Following Abowd et al. (1999), we decompose wages into a component due to time-variant observable individual characteristics, a “pure” worker effect, a “pure” firm effect, and a statistical residual, using the following equation:

$$w_{it} = X'_{it}\beta + \theta_i^F + \sum_j d_{ijt} \psi_j^F + \varepsilon_{it}, \quad (1)$$

where the subscript i denotes the worker, j denotes the firm, t denotes time, X'_{it} is a vector of individual time-varying controls, θ_i^F is the worker effect,

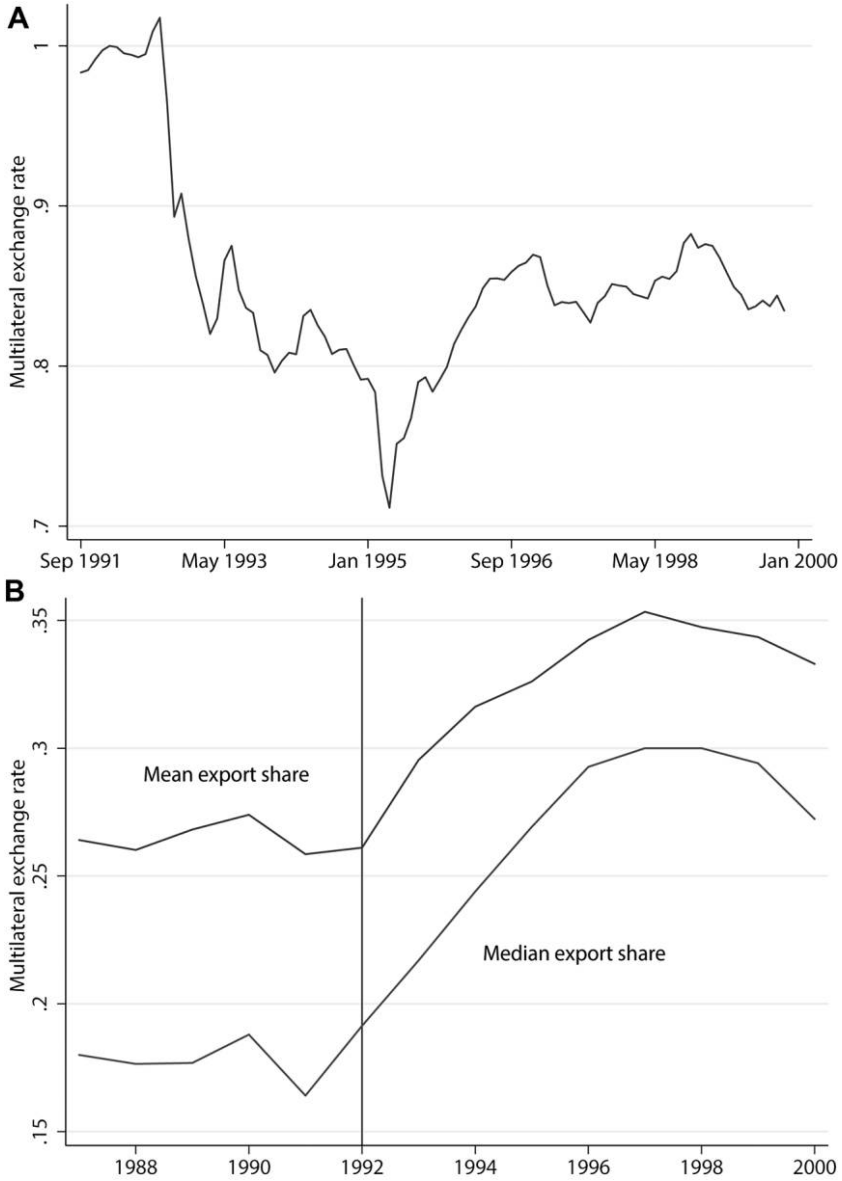


FIG. 1.—Real exchange rate and export share. *A*, Real exchange rate (January 1992 = 1). Source: Finicelli et al. 2005. *B*, Mean and median export share. Source: INVIND. A color version of this figure is available online.

d_{ijt} is a dummy equal to 1 if worker i is in firm j at time t , and ψ_{jt}^F is the firm-year effect. We use the superscript F to indicate that the worker effect is fixed over time, in order to distinguish this from the case in which we allow it to vary between the pre- and post-devaluation periods (see below). Abowd et al. (1999) show that, under the assumption of random workers' mobility across firms (conditional on the observables), equation (1) can be estimated and firm and worker effects separately identified. The identification of firm effects and worker effects is guaranteed by the substantial mobility of workers in the sample: 63% of the workers in the sample were employed by at least two different firms in the period 1982–97, and between 8.4% and 20% of workers changed employer from one year to the next.

We use the estimated worker effect $\hat{\theta}_i^F$ as our measure of the unobservable (to the econometrician) "skill component" of wages. Under the Abowd-Kramarz-Margolis assumptions (see Sec.II.B.3 for a discussion), the worker effect represents the component of wages that reflects the market value of the workers' unobservable skills independent of the characteristics of the particular firm that the individual works for and net of the workers' personal time-varying characteristics included in the controls. The firm-year effect ψ_{jt}^F represents the firm-specific contribution to wages after controlling for individual workers' characteristics. As such, it can be interpreted in terms of rent sharing. Because we are interested in relating rent sharing to the firm's export behavior, which changes over time (specifically, after the devaluation), we modified the original Abowd-Kramarz-Margolis (1999) procedure, which imposes a time-invariant firm effect, and estimate a time-varying firm effect $\hat{\psi}_{jt}^F$. Note that, thanks to the 50-employee minimum restriction and because we observe the complete workforce of all INVIND firms, for each firm-year we have at least 50 observations (the average is 672 and the median 228; see table 2), which guarantees a reasonably high precision of the estimates of the firm-year effects.⁵

The above procedure allows for the firm effect to vary over time while keeping the worker effect fixed. Although assuming that the individual worker effects do not vary over time is the standard approach in the literature (Abowd et al. 1999, Frías et al. 2009, Krishna, Poole, and Senses 2011), it might be too restrictive for the question that we are addressing. As shown

⁵ We estimate equations (1) and (2) using all available observations, and not only those of INVIND firms, because this improves the precision of the estimate of the workers' effects. However, for non-INVIND firms, for which we have, on average, 2.5 workers per year, we impose that the firm effect is fixed throughout the period. Note that these firms do not enter the subsequent analysis, because for non-INVIND firms we have no export information, so there is no advantage to recovering a time-varying measure that can be interpreted in terms of rent sharing for these firms.

in figure 1, the average increase in the export share of Italian manufacturing firms was very large. Such a shift might have induced a change in the market value of different skills. Indeed, there is evidence that the devaluation has had an impact on the demand for observable skills. In figure 2, we plot the time-series evolution of the share of production workers in the INVIND firms. It decreases regularly from .71 in 1986 to .67 in 1993, following the secular decline common to all advanced economies. When the devaluation hits, the fall stops and the share of production workers remains stable until 1998, after which it starts falling again. This is exactly the period in which the devaluation boosted the export activity (see fig. 1) and possibly the demand for production workers. More generally, some workers might be endowed with skills that are more valuable in export markets, for example, human capital specific to products that were particularly favored by the devaluation. It is indeed possible that the returns to such skills increased after the devaluation. If a firm employs workers with export-valuable skills, it might export more and pay higher wages because the market value of such skills has increased. By keeping the worker effects fixed, however, one would exclude this possibility a priori, forcing the higher wages to be picked up by the firm-year effect, thus attributing the higher wages to rent sharing. To account for changes in the market value of skills, we therefore allow the




FIG. 2.—Share of blue-collar workers. Source: INVIND. A color version of this figure is available online.

worker effects to take different values before and after the devaluation, estimating an extended version of equation (1) as follows:

$$\begin{aligned} w_{it} &= X'_{it}\beta + \theta_{it}^V + \sum_j d_{jit}\psi_{jt}^V + \varepsilon_{it}, \\ \theta_{it}^V &= (1 - DV_t)\theta_i^{\text{PRE}} + DV_t\theta_i^{\text{POST}}, \end{aligned} \quad (2)$$

where $DV_t = 1$ for $t > 1992$ and 0 otherwise, and θ_i^{PRE} and θ_i^{POST} are worker i 's effects computed separately for the pre- and post-1992 periods. Given that we have data up to 1997, θ_i^{POST} is estimated on a maximum of five observations per individual. The average number of individual-year observations in the post-devaluation period is 4.1 and the median is 5.

2. *Abowd-Kramarz-Margolis Estimates*

Prior to the estimation, we identified the groups of “connected” workers and firms. A connected group includes all of the workers ever employed by any firm in the group and all of the firms that any worker in the group has ever worked for. It is only within connected groups that worker and firm effects can be identified (Abowd, Creedy, and Kramarz 2002). By design, our sample consists of essentially one large, connected group, with 99.6% of the sample forming a single connected group. Thus, in our estimation, we focus on the largest connected group and disregard the remaining observations. In table 3, we present the results from estimation of equations (1) and (2).⁶ The dependent variable w_{it} is the natural logarithm of weekly wages. The vector X_{it} includes age and age squared (proxying for labor market experience), tenure and tenure squared,⁷ a dummy variable for nonproduction workers, and a dummy for managers (occupational status changes over time for a considerable number of workers), as well as interactions of all of these terms with a female dummy variable.

The estimated coefficients on the workers' observable characteristics, shown in table 3, deliver unsurprising results: wages appear to exhibit con-

⁶ The estimation was carried out using the conjugate gradient algorithm proposed by Abowd et al. (2002) and implemented by the Stata routine “a2reg” developed by Ouazad (2008).

⁷ Our data on tenure are right-censored because we do not have information on workers prior to 1980. To partly account for this censoring, for the affected workers, we included age-group-specific trends (based on these workers' ages in 1980) in the Abowd-Kramarz-Margolis regressions. Specifically, we start these workers' tenure as if they were joining their employer in 1980, and we interact this tenure variable with dummy variables reflecting the workers' ages in 1980 (age groups: 16–24, 25–34, 35–44, 45–54, and 55–65). The results are robust to the exclusion of these additional controls. In fact, the correlation between the worker (firm) effects with and without the tenure controls is 0.99 (0.98).

Table 3
Estimating Worker Effects and Firm Effects: Two-Way Fixed
Effects Regressions

	Abowd-Kramarz-Margolis Regressions Results	
	θ_i^F, ψ_{it}^F (1)	$\theta_{it}^V, \psi_{it}^V$ (2)
Number of observations	18,552,601	18,552,601
Number of worker/worker-period FEs	1,711,542	2,757,402
Number of firm-year FEs ^a	459,563	459,563
<i>F</i> (Prob > <i>F</i>)	49.99 (.000)	40.72 (.000)
<i>R</i> ² (adjusted <i>R</i> ²)	.87 (.85)	.90 (.87)
Root mean squared error	.124	.111
Coefficients on worker characteristics:		
Age	.0333	.0394
(Age) ²	-.0002	-.0002
Age × female	-.0153	-.0169
(Age) ² × female	.0001	.0001
Tenure ^b	.0094	.0059
(Tenure) ²	-.0004	-.0003
Tenure × female	-.0025	-.0002
(Tenure) ² × female	.0001	-.00001
White collar	.0699	.0533
Executive	.5136	.4252
White collar × female	-.0073	-.0024
Executives × female	.0191	-.0010
Comparison with match effects model ^c :		
Match effects model <i>R</i> ² ; adjusted <i>R</i> ²	.90; .87	.92; .89
Root mean squared error	.122	.115

NOTE.—The sample includes all firms and all workers in the largest connected group. The estimation was performed using the conjugate gradient algorithm proposed by Abowd et al. (2002) and implemented by the Stata code “a2reg” written by Ouazad (2008). See Sec. II.B for more details.

^a Year-specific firm effects are estimated for INVIND firms, whereas the estimated firm effects are time-invariant for non-INVIND firms.

^b To partly account for the censoring of the tenure variable for workers who appear in the data set prior to 1981, we included age group-specific trends based on these workers’ age in 1980. Specifically, we start these workers’ tenure as if they were joining their employer in 1980, and we interact this tenure variable with dummy variables based on the workers’ age in 1980 (age groups: 16–24, 25–34, 35–44, 45–54, and 55–65). These coefficients are not reported to save space.

^c Column 1 includes worker-firm match fixed effects, and col. 2 includes pre- and post-devaluation worker-firm match effects.

cave age and tenure profiles and a substantial wage premium is associated with white-collar jobs and, especially, with managerial positions. Table 4 presents the standard deviations of and the correlations between log wages (w_{it}) and the different components of wages ($\theta_i^F, \theta_{it}^V, \psi_{jt}^F, \psi_{jt}^V$). Similar to Abowd et al. (1999) and Iranzo et al. (2008), a substantial portion of the variation in earnings is due to heterogeneity in worker effects (the correlation between wages and worker effects is 0.46 when the worker effects are time-invariant and 0.39 when they are allowed to vary before and after the devaluation). Firm effects also play an important role (the correlation between

Table 4
Variance-Covariance Matrix of Workers' and Firms' Effects

	w_{it}	θ_i^F	θ_{it}^V	ψ_{it}^F	ψ_{it}^V
w_{it}	.34				
θ_i^F	.46	.24			
θ_{it}^V	.39	.95	.28		
ψ_{it}^F	.45	-.02	.01	.13	
ψ_{it}^V	.44	.01	-.04	.93	.13

NOTE.—The diagonal entry reports the standard deviation, and the other entries are correlations.

wages and ψ_{it}^F is equal to 0.45 and that with ψ_{it}^V is 0.44). The two measures of worker effects (θ_i^F and θ_{it}^V) are strongly positively correlated with each other (the correlation between θ_i^F and θ_i^{PRE} is 0.96 and that between θ_i^F and θ_i^{POST} is 0.91), and so are the two measures of firm-year effects (correlation = 0.93). The correlation between the worker and the firm effects is zero when the worker effects are time-invariant and turns negative when the worker effects are time-invariant. Finally, the pre- and post-devaluation worker effects (θ_i^{PRE} and θ_i^{POST}) are positively correlated with each other (correlation = 0.83), which was expected; in fact, even though the devaluation might have changed the returns to skills, workers who commanded a high wage before the devaluation, on average, do the same after it. Still, the correlation is substantially below 1, indicating that returns to (unobservable) skills have changed in the two periods.

3. *Tests of the Abowd-Kramarz-Margolis*
“Exogenous Mobility” Assumption

The Abowd-Kramarz-Margolis wage decomposition rests on an assumption of exogenous worker mobility conditional on observables and on worker and firm effects.⁸ Specifically, as described in greater detail in Card, Heining, and Kline (2013) and Card, Cardoso, and Kline (2016), the Abowd-Kramarz-Margolis method assumes that worker mobility is independent of (i) time-specific, firm-wide shocks (whereby workers move away from firms that

⁸ There are several forms of endogenous mobility that are not a problem for the Abowd-Kramarz-Margolis estimation, including assortative matching based on person and firm effects. As Card, Cardoso, and Kline (2016) point out, there are various mechanisms that might lead to skilled workers being more likely to end up at high-wage firms, including skilled workers being more likely to engage in on-the-job search, benefiting from networks of high-skill family members or friends, or enjoying lower rates of job destruction during downturns. The fact that all these mechanisms lead to positive assortative matching between workers and firms does not violate the Abowd-Kramarz-Margolis assumptions because the estimation conditions on time-invariant worker and firm characteristics.

are experiencing bad shocks to firms that are experiencing good shocks); (ii) firm-worker match effects (workers moving to firms where they have a large match effect and away from firms where the match effect is small); and (iii) transitory wage shocks (workers who displayed good performance being more likely to move to high-wage firms and workers who are underperforming moving to low-wage firms). Our version of the Abowd-Kramarz-Margolis model removes concern (i) because our models include year-specific firm effects. As for issues (ii) and (iii), following Card, Heining, and Kline (2013) and Card, Cardoso, and Kline (2016), we performed a series of tests. The results, which we report in detail in the appendix, available online, indicate that, similar to what was found by Card, Cardoso, and Kline in Portugal and Card, Heining, and Kline in Germany, match effects are not a primary driver of mobility in the Italian manufacturing sector context, and they also indicate that the additively separable firm and worker effects obtained from the Abowd-Kramarz-Margolis model can be taken as reasonable measures of the unobservable worker and firm components of wages.⁹

C. Exports and Wages: Descriptive Evidence

Before moving to the effects of the devaluation, we first analyze the correlation between wages and export activity and offer some suggestive evidence on the association between skill composition and rent sharing and the export wage premium. Of course, at this stage we cannot interpret this in any causal sense. As a measure of export activity, we use the share of export sales in total sales. Indeed, most firms in our sample are exporters, but they do differ considerably in their export intensity (see table 2).

In panel A of table 5, we report the wage regressions. Column 1 uses as dependent variable the log of the individual workers' weekly earnings (which we will be referring to as "log wage" for brevity), and controls for worker-level characteristics (gender, age, age squared, white collar dummy, manager dummy), firm characteristics (employment categories [<100 , $100-300$, $300-500$, and $500+$ employees], log of domestic sales, 14 industry dummies), as well as 4 geographic area dummies (South, Center, North-East and North-West) and 10 year dummies (1987-97). We find a strong, positive association between the export share (EXSH) and log wages, with a coefficient of 0.093 (SE = 0.010). Other things equal, a one standard deviation higher export share is associated with 2.5% higher wages. In column 2, we repeat the same regression using firm-level averages. We find that the results are similar, although the coefficient increases slightly to 0.108 (SE = 0.006), arguably because the workers' controls aggregated at the firm level are less precise than

⁹ We also performed three additional statistical tests proposed by Lazear, Shaw, and Stanton (2015), reaching similar conclusions. See the appendix for details.

Table 5
Export Intensity and Wages, Skill Composition, and Rent Sharing:
Cross-Sectional and Within-Firm Patterns

	(1)	(2)	(3)	(4)
Panel A: Wages				
Log Wage				
EXSH	.093*** (.010)	.108*** (.006)	.065*** (.010)	
Specification	OLS	OLS	FE	
Observations	4,074,074	7,579	7,579	
R ²	.59	.77	.96	
Panel B: Skill Composition				
	SC ^F		SC ^V	
EXSH	.027*** (.004)	-.007 (.004)	.042*** (.004)	.016* (.010)
Specification	OLS	FE	OLS	FE
Observations	7,579	7,579	7,579	7,579
R ²	.67	.96	.64	.89
Panel C: Rent Sharing				
	RS ^F		RS ^V	
EXSH	.075*** (.005)	.069*** (.010)	.061*** (.006)	.048*** (.012)
Specification	OLS	FE	OLS	FE
Observations	7,579	7,579	7,579	7,579
R ²	.55	.89	.47	.84

NOTE.—The sample includes INVIND firms, years 1987–97. EXSH = share of sales that is exported; OLS = ordinary least squares; FE = fixed effects. Panel A: an observation is a worker-year in col. 1 and a firm-year in cols. 2 and 3. The dependent variable is the log of individual weekly earnings in col. 1 and the average of log weekly earnings in the firm-year in cols. 2 and 3. All regressions include controls for worker gender, age, tenure, white collar, manager, employment categories (< 100, 100–300, 300–500, 500+ employees), log of domestic sales, industry dummies (14), geographic area dummies (4), and year dummies (10). Panel B: SC^F are firm-year averages of the time-invariant worker effects (Abowd-Kramarz-Margolis regression of table 3, col. 1), and SC^V are firm-year averages of the worker effects that were allowed to vary before and after the devaluation (table 3, col. 2). See Sec. III.B and table 3 for details. Panel C: RS^F are firm-year effects obtained from Abowd-Kramarz-Margolis regressions with time-invariant worker effects (table 3, col. 1), and RS^V are firm-year effects obtained from Abowd-Kramarz-Margolis regressions where the worker effects were allowed to vary before and after the devaluation (table 3, col. 2). See Sec. III.B and table 3 for details. Robust standard errors are reported in parentheses.

* Statistically significant at the 10% level.
 *** Statistically significant at the 1% level.

those at the individual level. Finally, in column 3, we include firm fixed effects in the regression to control for unobserved time-invariant firm heterogeneity. The estimated coefficient is equal to 0.065 (SE = 0.010), which indicates that a robust association exists between within-firm changes in exports and changes in wages.

To disentangle the effect of the skill composition from that of rent sharing, we resort to the wage decomposition described in Section II.B (similar

to Frías et al. [2009]). We define a skill composition (SC) term as the average worker effect at the firm-year level:

$$SC_{jt} \equiv \frac{1}{n_{jt}} \sum_i d_{ijt} \theta_{it}, \quad (3)$$

where, as before, d_{ijt} is a dummy equal to 1 if worker i is in firm j at time t and n_{jt} is the number of workers in firm j at time t . We measure rent sharing (RS) as the firm-year effect:

$$RS_{jt} = \psi_{jt}. \quad (4)$$

As before, the superscripts F and V will be used to denote the case in which the worker and firm effects were obtained from specification (1) (with time-invariant worker effects) and specification (2) (when the worker effects are allowed to vary before and after the devaluation), respectively.

Next, we explore the relationship between export intensity and SC and RS and report the results in panels B and C, respectively, of table 5. Specifically, in the first two columns, the dependent variables are SC^F and RS^F (i.e., obtained with time-invariant worker effects, respectively) and the last two SC^V and RS^V (i.e., obtained with time-varying worker effects). The OLS results reported in column 1 of the two panels indicate that the wage premium is explained by both workforce composition and rent sharing: the coefficient on the export share is positive and significant in both panels. However, the elasticity of RS^F is almost three times larger than that of SC^F (0.075 vs. 0.027). This finding is reinforced in the firm fixed effects (FE) specifications of column 2, in which case the coefficient is essentially zero for SC^F and remains positive and significant, albeit slightly smaller, for RS^F . These results suggest that some firms employ higher-skilled workers and export more (OLS results) but that changes of the export share over time are not reflected, on average, in changes in the skill composition. On the contrary, the fixed effects results indicate that there is a positive correlation between export intensity and RS^F within firms over time.

The picture changes somewhat when we use SC^V and RS^V (col. 3, OLS; col. 4, FE). The coefficient in the SC^V regression increases considerably, whereas that for RS^V decreases, compared to when SC^F and RS^F were used. The coefficient on the export share in the SC^V regression also remains marginally significant when we control for firm fixed effects. The fact that we do find a positive association between within-firm changes in export shares and SC^V is consistent with the idea that changes in export intensity associated with the devaluation might have changed the market values of workers' skills. Once this is taken into account, the correlation between export intensity and the skill composition component of wages increases and that with rent sharing decreases. We thus hypothesize that imposing a fixed skill level pre- and

post-devaluation might be too restrictive and might lead one to attribute to RS part of the effect that is instead due to a change in the market value of workers' unobserved skills. Of course, at this stage, no claims of causality can be made. We will return to this point in the next section, after describing our identification strategy.

III. Evidence from the 1992 Devaluation Episode

In this section, we tackle the issue of causality in the relationship between exporting and wages. As mentioned above, we exploit the sudden and substantial devaluation of the Italian lira that occurred in September 1992 as an exogenous shock to the incentives of Italian firms to export.

A. An Unexpected Shock to the Exchange Rate

The currency devaluation of September 1992 was largely unpredicted. The speculative attacks that led to the devaluation started after the Danish referendum of June 2, 1992, that, quite unexpectedly, rejected the Maastricht Treaty by a small margin (0.7%). The Danish referendum represented a big blow to the process of European integration. One consequence was diminished credibility of the exchange rate mechanism (ERM), which immediately led to speculative attacks against the weak currencies. The monetary authorities resisted the attacks until the end of the summer. The Italian lira devaluated by 7% during the weekend of September 12–13. On September 16, the British pound left the ERM; the lira and the Spanish peseta suspended their exchange rate agreements immediately after. Italy rejoined the ERM only on November 25, 1996. During those 4 years, the exchange rate of the lira fluctuated substantially.

Even though the depreciation was unexpected, one might argue that its effects were differentiated according to some firm characteristics, which, in turn, might be correlated with subsequent wage changes. For Mexico, Verhoogen (2008) shows that larger, more productive firms were more likely to take advantage of the peso devaluation of 1993 because the Mexican exports were directed, to a large extent, to the United States, so that exported goods were, on average, of a higher quality than were those sold domestically. Firms that increased exports to the United States were therefore those already producing high-quality goods before the devaluation. They undertook further quality upgrading, which led to an increased gap with respect to non-exporting firms.

It is not clear, however, whether the same patterns characterize the Italian case. In terms of standard indicators of development, such as income per capita or labor costs, Italy is a developed economy. However, its production structure was (and still is) specialized in medium- and low-tech activities, such as textiles, furniture, and tiles. Bugamelli, Schivardi, and Zizza (2010) argue that firms in these sectors were those that benefited most from

the devaluation.¹⁰ Their argument is based on the assumption that pure price competition is relatively more important in low-tech activities. The price advantage of a devaluation should have been therefore more pronounced for firms not at the top of the quality or technology ladder. The same type of reasoning applies within industries. For example, in the textiles sector, firms that produce low-quality shirts coexist with luxury fashion producers. The argument is that the former might have benefited more from the devaluation because the demand elasticity for such goods is higher, given the production of close substitutes in low-wage countries. It is therefore unclear *ex ante* which type of firms benefited most from the lira devaluation. In fact, such benefits might have depended on a series of factors, such as export destination, relative importance of price competition, and product composition, that are not easily linked to any specific firm characteristic.

To probe the hypothesis that changes in the export share following the devaluation were, to a large extent, exogenous with respect to pre-devaluation firm characteristics, we run a set of regressions similar to those of Verhoogen (2008):

$$\Delta \text{EXSH}_{i,t_1,t_0} = \alpha + \rho X_{i,t_0} + \text{Dummies} + \eta_i, \quad (5)$$

where $\Delta \text{EXSH}_{i,t_1,t_0}$ is the change in the firm-level share of export sales over total sales between t_0 and t_1 , X_{i,t_0} is a firm characteristic measured at t_0 , and Dummies are sector and area dummies. In the Mexican case, Verhoogen (2008) and Frías et al. (2009) use employment, sales per worker, and TFP as proxies for plant heterogeneity, and they find that the estimates of ρ for the devaluation period is substantially larger than in other periods. This is interpreted as showing that “better” firms took greater advantage of the devaluation. We run the same type of regressions and report the results in table 6. We consider three periods: the devaluation period (1991–95), the pre-devaluation (1987–91) period, and the post-devaluation (1995–99) period. Following Frías et al. (2009), we regress the change in the share of exports on the log of domestic sales, the log of employment, and the log of the ratio between investment and employment measured in the initial year.¹¹ The first two are measures of firm size, whereas the last is a proxy for the capital-labor ratio. We find no significant differences in the coefficients

¹⁰ Crinò and Epifani (2012) document only a weak relationship between exports and TFP across Italian firms.

¹¹ Note that domestic sales both enter the left-hand-side variable in the denominator and are used as a regressor. This might induce some downward bias in the estimates of ρ , for example, if they are measured with error. Notice, however, that we are not interested in the absolute value of the estimate but in its stability across time periods, which should not be affected by measurement error. More importantly, the other two proxies are immune from this potential problem and deliver the same message.

Table 6
Changes in Export and Initial Conditions

	Dependent Variable: Δ EXSH in Given Period		
	1991–87 (1)	1995–91 (2)	1999–95 (3)
Domestic sales	.011** (.005) [.45]	.016*** (.004)	.011* (.006) [.46]
Employment	-.004 (.004) [.16]	.005 (.005)	.002 (.005) [.71]
Investment/workers	-.000 (.006) [.95]	.000 (.005)	-.003 (.005) [.67]

NOTE.—Each coefficient comes from a separate regression. The dependent variable is the change in the share of export over total sales (EXPSH) over the relevant interval. Domestic sales is the log of real domestic sales, employment the log of the number of employees, and Investment/workers the log of real investment over the number of employees. The regressors are measured at the initial year of the relevant interval (i.e., Domestic sales in col. 1 is the log of real domestic sales in 1987, in cols. 2 in 1991, and so on. Standard errors are in parentheses. In square brackets, we report the *p*-value of a test of equality of the coefficient with the corresponding coefficient for the 1991–95 regression. All regressions include 17 sector and 4 area dummies.

- * Statistically significant at the 10% level.
- ** Statistically significant at the 5% level.
- *** Statistically significant at the 1% level.

between the devaluation period and the other two periods for any of the indicators. We conclude that there is no evidence that “better” firms disproportionately took advantage of the 1992 lira devaluation.

B. Empirical Strategy

We are interested in singling out the effects of a change in the export share on wages and on its components following the 1992 devaluation. For the dependent variable, we consider the wage, the firm-year average worker effect (which measures the skill composition), and the year-specific firm component (which measures rent sharing). The identifying assumption is that changes in export shares in the devaluation period are indeed attributable to the unforeseen devaluation episode and were not correlated (as shown in the previous section) with pre-existing firm attributes commonly used in the literature as proxies for firm “quality.” We take the years 1987–91 as the base period, before the devaluation occurred,¹² and we define DV as

¹² Fig. 2 shows that the average and median export shares in the INVIND sample were very stable during the pre-devaluation period. We also ran regressions for which we picked 1990 or 1991 as base years, obtaining similar results.

a dummy for the years from 1992 onward. We specify our main regression as

$$y_{jt} = \alpha + \beta [\text{EXSH}_{jt} \times (1 - \text{DV})] + \gamma (\text{EXSH}_{j\text{PRE}} \times \text{DV}) + \delta (\text{EXSH}_{jt} \times \text{DV}) + \mathbf{X}'_{jt}\boldsymbol{\theta} + \mu_j + \varepsilon_{jt}, \quad (6)$$

where j denotes firms and t years, y_{jt} is alternatively the wage w , SC and RS as defined in equations (3) and (4), respectively, EXSH_{jt} is the current export share, $\text{EXSH}_{j\text{PRE}}$ is the average share in the pre-devaluation (1987–91) period, \mathbf{X}_{jt} is a vector of additional controls of firm and workforce characteristics, and μ_j are firm fixed effects. In this specification, β measures the correlation between export share and the dependent variable in the pre-devaluation period in the same way as in the basic OLS regressions that were described in the previous section. For the devaluation period ($\text{DV} = 1$), we control for pre-existing effects of export on worker compensation by including the share of export in the pre-devaluation period interacted with the DV dummy. By doing so, we control for the possibility that firms that were exporting more in the pre-devaluation period might have also been paying higher wages, which could persist in the post-devaluation period; moreover, this also accounts for the possibility that a higher initial export share might be correlated with higher wages in the devaluation period, if, as in Frías et al. (2009), “good firms” take greater advantage of the devaluation. The coefficient γ will therefore capture any of these effects if they are indeed present. Thus, controlling for the pre-devaluation export propensity, the coefficient δ measures the effect of the current export share on wages. Despite its simplicity, this specification encompasses many different regimes, according to the values of the estimated parameters. We discuss the most interesting ones next.

$\delta = 0, \gamma > 0$: In this case, the relation between wages and export activity is a fixed firm attribute. The pre-devaluation export share captures this attribute and constitutes a sufficient statistic to predict the effects of the devaluation on wages, whereas the actual share has no effect.

$\delta > 0, \gamma = 0$: This configuration would indicate that export propensity is not a fixed firm attribute: a devaluation might entail changes in export that cannot be predicted on the basis of pre-existing conditions. Moreover, controlling for the current propensity, pre-existing conditions play no role in determining the impact of export propensity on wages.

$\delta > 0, \gamma > 0$: In this case, there is a role for both a predetermined fixed component and for current conditions.

C. Main Results

We estimate equation (6) using log wages, SC and RS as dependent variables, and report the results in table 7. In panel A, we present OLS results;

Table 7
Devaluation Regressions

	Dependent Variable				
	Log W (1)	SC^F (2)	RS^F (3)	SC^V (4)	RS^V (5)
Panel A. Ordinary Least Squares					
δ : $EXSH \times DV$.085*** (.015)	.005 (.008)	.080*** (.014)	.032*** (.012)	.054*** (.015)
γ : $EXSH_{PRE} \times DV$.037*** (.015)	.026*** (.008)	.007 (.015)	.028** (.012)	.006 (.016)
β : $EXSH \times (1 - DV)$.100*** (.008)	.041*** (.005)	.047*** (.007)	.033*** (.005)	.057*** (.007)
Observations	6,328	6,328	6,328	6,328	6,328
R^2	.77	.68	.55	.66	.48
Panel B. Firm Fixed Effects					
δ : $EXSH \times DV$.070*** (.011)	-.014** (.006)	.083*** (.012)	.031*** (.012)	.041*** (.014)
γ : $EXSH_{PRE} \times DV$.034** (.017)	.008 (.007)	.025 (.016)	.005 (.016)	.025 (.019)
β : $EXSH \times (1 - DV)$.082*** (.016)	.007 (.006)	.071*** (.015)	.013 (.015)	.066*** (.018)
Observations	6,328	6,328	6,328	6,328	6,328
R^2	.95	.95	.87	.88	.82

NOTE.—The sample includes INVIND firms, years 1987–97. One observation is a firm-year. $EXSH$ is the share of sales that is exported; $EXSH_{PRE}$ is the average export share in the pre-devaluation period; DV is a dummy variable equal to 1 after 1992 (devaluation period). Controls include firm-year-level workforce characteristics (average age, percentage of females in the workforce, percentage of white-collar employees), time-varying firm characteristics (employment categories, log of domestic sales) as well as year, industry, and region dummies. SK^F and RS^F are firm-year-level worker effects and firm effects obtained from the Abowd-Kramarz-Margolis regressions described in table 3, col. 1, where the worker effect is time-invariant, and SK^V and RS^V were obtained from the Abowd-Kramarz-Margolis regressions described in table 3, col. 2, where the worker effects were allowed to take different values in the pre- and post-devaluation periods. See Sec. III.B and table 3 for more details. Robust standard errors are reported in parentheses.

** Statistically significant at the 5% level.
*** Statistically significant at the 1% level.

in panel B, we adopt a fixed effects specification, with the fixed effects defined at the level of the firm.

1. Assuming Time-Invariant Worker Skills

We begin by describing the OLS results. All regressions include firm-year-level workforce characteristics (average age, average tenure, percentage of females in the workforce, percentage of white-collar employees), and time-varying firm characteristics (employment size categories, log of domestic sales), as well as region, industry, and year dummies. For wages, we find in column 1 that $\delta = 0.085$, statistically significant at the 1% level, and $\gamma = .037$, statistically significant at the 5% level. This implies that, controlling for the pre-devaluation export share, a higher post-devaluation share is associated with higher wages. Moreover, given that γ is less than half δ in

magnitude, the current export share is what matters the most for wage determination, contrary to the idea that some pre-determined condition, captured by the pre-devaluation share, is the main driver of the post-devaluation outcomes. In terms of magnitude, the effect is not negligible: given that the median export share has increased by approximately 15 percentage points during the devaluation episode, workers in the median firm recorded a wage 1.3% higher following the devaluation. Finally, the coefficient on the export share in the pre-devaluation period, β , is 0.1, highly significant and not too dissimilar to what we found in the OLS regressions presented in table 4. Columns 2 and 3 report the results for SC^F and RS^F (i.e., the worker and firm effects estimated while assuming the worker effects are constant over time) as the dependent variables. For SC^F , we obtain an estimated δ essentially equal to zero (0.005, not statistically significant). The coefficient on the pre-share value γ is instead positive (0.026) and statistically significant at the 1% level. This suggests that the workforce composition is indeed a quasi-fixed attribute, so that, controlling for the pre-devaluation export share, changes in the share after the devaluation do not affect the firms' skill composition. Stated differently, the higher post-devaluation wages do not seem to be due to a workforce composition effect. The relationship between export intensity and workers' skills is confirmed by the estimate of β (0.027, statistically significant at the 1% level): "normal" (pre-devaluation) export activity is associated with higher skills. For RS^F , our measure of firm-year wage premia, the estimated δ is positive, whereas the estimated γ is small (0.007) and not statistically significant. This indicates that the higher wages related to increases in the export shares are mostly due to a time-varying firm effect, which is common to all workers in the firm. This is consistent with the idea that the firm and the workers share the surplus from the increase in exports following the devaluation. Moreover, the rent sharing component does not seem to be a fixed firm attribute in that it is not related to the average export share in the pre-devaluation period. The coefficient estimate on the export share in the pre-devaluation period, β , is instead sizable (0.047) and statistically significant; the fact that its magnitude is smaller than that of δ suggests that a larger share of the surplus is enjoyed by workers after the devaluation compared to "normal" times.

We now turn to the fixed effects specifications (panel B of table 7). This exploits only within-firm variation, and it ensures that we are controlling for any firm-specific, time-invariant unobservable characteristics. In fact, one could argue that controlling for the pre-devaluation average export share and other firm-level controls such as employment and domestic sales is not sufficient to account for potential firm heterogeneity. Once we do that, we obtain a slightly smaller estimate of δ for wages (0.070) in column 1; this is still highly statistically significant, which indicates that the effects of export on wages do not simply reflect a fixed firm attribute. We also still obtain that when we constrain the worker effects to be time-invariant, all of the wage

effect is attributable to changes in wage premia rather than changes in workforce composition. In fact, when the dependent variable is SC^F (col. 2), the coefficient estimate of δ is negative (-0.014), in line with the idea that the devaluation might have favored low-skill firms, and the estimated γ and β are small and not statistically significant. This is consistent with the view that the skill composition is a rather fixed attribute so that, once we control for firm fixed effects, the within-firm variation in exports has very little effect on the skill composition. Instead, when the dependent variable is the rent sharing component RS^F (col. 3), the estimated coefficient δ remains positive and strongly significant and essentially unchanged in magnitude (0.083). This indicates that within-firm variations in the share of exports over time are strongly reflected in the rent sharing component of wages.

2. *Allowing Worker Skills to Vary Pre- and Post-Devaluation*

The picture that emerged from columns 1, 2, and 3 of table 7 panels A and B is that the export activity stimulated by the devaluation of the lira led to higher wages, and the increased wages were entirely due to rent sharing with little evidence of changes in skill composition. However, SC^F and RS^F were estimated under the assumption that the worker effect is fixed over time. Such an effect captures the combination of two elements in the wage determination: the worker's unobservable skills and the price that the labor market assigns to such skills. Although it seems plausible that there is a fixed component of workers' skills, such as education and other cognitive skills, and noncognitive ability, it is less obvious that the market value of these skills is unchanged following such a strong shock as the devaluation that we are analyzing. Specifically, workers might be heterogeneous in terms of export-specific skills. For example, the devaluation might have been particularly advantageous for some products more traded on international markets. Then, if part of the human capital is product-specific, workers with skills that are more useful for the exporting activity might observe an increase in the market value of their skills. Consider now the case of a firm with abundant export-specific skills that, after the devaluation, increases its export share substantially. If the market value of its workers' skills has increased, the firm will increase compensation accordingly, not to share rents but rather to meet the higher market value of its workers' skills. Failure to allow the worker effects to change over time would imply that the increased market value of those skills would be absorbed by the time-varying firm component of wages (RS^F), thus overestimating the rent sharing component of the correlation between export and wages.¹³

¹³ Measurement error in the export share could in principle bias our estimates. There are several reasons why we believe this is not a concern in our setting. First, the quality of the INVIND survey data we use is very high. The survey is run by

This conjecture is confirmed by the results reported in columns 4 and 5 of table 7, where we report the results for SC^V and RS^V , that is, the average firm-year worker effect and rent sharing component estimated while allowing the individual worker effects to take different values in the pre- and post-devaluation periods (model (2)). In fact, when SC^V is the dependent variable, the estimated δ is positive and statistically significant, in both the OLS (col. 4 of panel A) and fixed effects (col. 4 of panel B) specifications (0.032 and 0.031, respectively, compared to 0.005 and -0.014 when using SC^F). The estimated δ remains large and significant also for RS^V (col. 5 of panels A and B), although its magnitude is reduced compared to when RS^F was used (0.054 vs. 0.085 in the OLS specification and 0.041 vs. 0.083 in the fixed effects specification). Thus, it appears that the increased wage associated with exporting is due to both a firm-level component, plausibly related to rent sharing, and a component attributable to a change in the market value of workers' unobservable skills. Specifically, the results from the specifications with firm fixed effects (panel B, cols. 4 and 5) indicate that the two components contribute roughly equally to explaining the effect of export intensity on wages. Not allowing the worker effects to vary over time would have led to incorrectly concluding that the wage premium was entirely explained by rent sharing, with workers' skills not playing any role. We corroborate this interpretation in Section IV.C, where we relate the export wage premium to a measure of export-specific worker experience.

D. Accounting for Effort and Productivity

We address two potential concerns that might arise with the analysis above and our interpretation of the results: the estimated coefficient δ might reflect increased worker effort in response to the extra demand or it might reflect increased productivity that is related both to higher wages and to higher exports.

The first concern arises because our wage measure is total weekly earnings, and we have no information on hours worked at the individual level in the social security data. If employees in firms that increase the export share after the devaluation are working more hours per week to meet the extra demand, we would be capturing an effect on hours and not directly

the Italian Central Bank, which puts substantial effort in ensuring data reliability. In particular, the 100 provincial branches of the Bank of Italy are responsible for the questionnaires of the firms located in their province, so that each questionnaire is double checked (e.g., the firms are recontacted when anomalies are spotted). Second, if measurement error is i.i.d (and we see no obvious reason to assume that it is not), it is not correlated between the two variables, as γ is the coefficient of the export share in the pre-devaluation period and δ in the post-devaluation period. If this is the case, some attenuation bias could be present, and our estimates would be, in the worst of cases, a lower bound on the effect of export on wages.

on the wage rate. Fortunately, the INVIND firm survey does report the total hours worked at the firm-year level, from which we can recover a measure of average per capita hours. The results from including this additional control in the regressions are reported in panel A of table 8.¹⁴ The table presents results from fixed effects specifications, for which the fixed effects are defined at the firm level. As can be seen in columns 1, 3, and 5, hours worked are positively and significantly correlated with wages as well as with the firm effects (both RS^F and RS^V). However, our main coefficient of interest, δ , is still positive, sizable, and statistically significant, and its magnitude is only slightly reduced with respect to the results reported in table 6. The estimates for the worker effects (both SC^F and SC^V) are essentially unaffected. This indicates that the effects on the total compensation is not solely attributable to an increase in the number of hours worked.

The second possible issue is that the firm might become more productive as a consequence of expanding its export activities due to “learning by exporting” (De Loecker 2011). Indeed, evidence from other contexts does suggest that labor productivity and TFP (total-factor productivity) increase when firms begin exporting or when they expand their exports (Park et al. 2010). Thus, this would be a different mechanism for export to affect wages other than skill composition and rent sharing. To account for this possibility, we include TFP in the regression, computed using the Iranzo et al. (2008) procedure (see Olley and Pakes [1996] for the details). Because computing TFP requires further data, available only for a subset of firms, our sample size is reduced to 3,858 firm-year observations. Despite this, our results are robust to the inclusion of TFP in the controls, as shown in panel B of table 8.¹⁵

IV. Export Wage Premium and Workers’ Past Export Experience

Our results imply that workers enjoy higher wages when their firm increases its exports. The export wage premium is explained both by a firm-year factor, RS , which we interpret as rent sharing, and by a skill composition effect, SC , that emerges only if we allow the returns to (unobservable) skills to differ in the pre- and post-devaluation periods. The latter result can be explained if (a) there is heterogeneity in the distribution of skills in terms of usefulness for the export activity and (b) the devaluation increases the demand for those skills, driving their relative price up. In this section, we

¹⁴ Note that, because of some missing values in hours worked, our estimation sample is reduced from 6,328 to 6,219 observations.

¹⁵ Our findings and interpretation are further corroborated by an IV analysis that we performed, and that is described in the appendix. Our instrument is a measure of foreign demand constructed using data on exports by country of destination at the sector-province level (unfortunately, we do not have this information at the level of the firm). Even though the diagnostic tests indicate that the instrument is weak, the IV results deliver patterns in line with those presented in the text.

Table 8
Robustness Tests

	Dependent Variable				
	Log <i>W</i> (1)	SC ^F (2)	RS ^F (3)	SC ^V (4)	RS ^V (5)
Panel A. Controlling for Hours Worked					
δ : EXSH \times DV	.058*** (.011)	-.015*** (.006)	.071*** (.012)	.029** (.011)	.029** (.014)
Hours worked	.062*** (.010)	.002 (.003)	.065*** (.009)	.005 (.009)	.061*** (.011)
Observations	6,219	6,219	6,219	6,219	6,219
R ²	.95	.95	.88	.88	.82
Panel B. Controlling for Hours Worked and TFP					
δ : EXSH \times DV	.058*** (.014)	-.030*** (.006)	.089*** (.015)	.032** (.014)	.031* (.017)
Hours worked	.069*** (.009)	.000 (.003)	.073*** (.009)	.002 (.007)	.071*** (.010)
TFP	.032*** (.005)	.002 (.002)	.031*** (.005)	.003 (.004)	.029*** (.005)
Observations	3,858	3,858	3,858	3,858	3,858
R ²	.96	.98	.91	.89	.86

NOTE.—The sample includes INVIND firms, years 1987–97. One observation is a firm-year. All results are from fixed effects regressions, where the fixed effect is defined at the level of the firm. See the note to table 7 for the definitions of the dependent variables, the explanatory variables, and the list of control variables. Panel A: the controls include average hours worked (total house/employees) at the firm-year level. Panel B: the controls include average hours worked and total-factor productivity, TFP. See Sec. IV.D.1 for details. Robust standard errors are reported in parentheses.

* Statistically significant at the 10% level.

** Statistically significant at the 5% level.

*** Statistically significant at the 1% level.

corroborate this interpretation by testing whether the export wage premium associated with the devaluation can be linked to a measure of export-specific workers' skills. Measures of such skills are typically unobservable in administrative data or labor force surveys. Our data, however, allow us to construct measures of export-related skills based on workers' past experience in exporting activities. If producing for foreign markets requires a certain degree of specificity, then it seems plausible that a worker employed by an exporting firm can actually accumulate export-specific human capital.¹⁶ We use two measures of cumulated export experience. The first measure is the total number of years that a worker was employed at an exporting firm. Although this is perhaps the most natural way to measure export experience, this measure presents some limitations. In our data the vast majority of the firms are exporters (89% of all firm-year observations), which limits the variability

¹⁶ Mion and Opromolla (2011) find evidence of this mechanism for managers in Portuguese firms.

across workers in export experience thus defined. Moreover, for workers who never switch firms, such a measure would be perfectly collinear with tenure, which again is problematic. Our second, and preferred, measure of export experience exploits the substantial variation in the share of total revenue exported by firms and is constructed as $EXPER_{it} = \sum_{k=1}^5 EXSH_{j(i,t-k)}$, where i denotes workers, t denotes time, $j(i, t)$ denotes the firm where worker i was employed at time t , and $EXSH_{j(i,t-k)}$ is firm j 's export share at time t . The index $EXPER_{it}$ can take values between 0 (if a worker was employed by firms that produced only for the domestic market in the previous 5 years) and 5 (if the worker was employed by firms that exported their entire output in the previous 5 years). This measure of export experience breaks the collinearity between export experience and tenure; also, it differentiates between workers who spent the same number of years in firms that exported the majority of their output and those who were in firms that exported only a small share. Given that we are interested in the effects of an individual characteristic (export experience) on wages, we perform this analysis on the individual workers' data rather than on firm-year averages. Because we have information on exports for INVIND firms only, the export experience variables can be computed only for workers who have been employed at INVIND firms throughout the 1987–97 period. This subsample consists of 58% of the total INVIND workers' sample. As shown in column 4 of table 1, the characteristics of these workers are very similar to those in the full sample. Because export data are available only starting in 1984, a 5-year export experience index can be computed only beginning in 1989. Thus, the sample is reduced to slightly fewer than 1,200,000 person-year observations. In 1991, the year before the devaluation, workers in our sample had, on average, 5.5 years prior experience with exporting firms. As for our second measure, its mean (standard deviation) was equal to 1.21 (1.06).¹⁷ In table 9, we present results from estimating the following equation:

$$\begin{aligned} w_{it} = & \alpha + \beta [EXSH_{j(i,t)} \times (1 - DV)] + \gamma (EXSH_{j(i,t)PRE} \times DV) \\ & + \delta_0 (EXSH_{j(i,t)} \times DV) + \zeta EXPER_{it} + \xi (EXPER_{it} \times DV) \quad (7) \\ & + \delta_1 (EXSH_{j(i,t)} \times EXPER_{it} \times DV) + \theta' X_{ij(i,t)} + \mu_j + \varepsilon_{it}, \end{aligned}$$

which is a version of equation (6) augmented with $EXPER_{it}$ and its interaction with DV and $EXSH_{j(i,t)} \times DV$. By interacting past experience with the

¹⁷ We have also performed all of the regressions with $EXPER3Y_{it}$, defined in an analogous way but considering only the previous 3 years of employment, which allows us to use the entire 1987–97 period. We obtained very similar results. In 1991, the mean (median) $EXPER3Y$ was equal to 0.81 (0.68).

Table 9
Export Wage Premium and Workers' Export Experience

	Dependent Variable: Log <i>W</i>		
	(1)	(2)	(3)
Panel A			
δ_0 : EXSH \times DV	.114*** (.019)	.082*** (.021)	.083*** (.026)
δ_1 : EXSH \times DV \times EXPER	.006*** (.001)	.008*** (.002)	.002 (.002)
EXPER	.004*** (.001)	.008*** (.001)	.007*** (.001)
EXPER \times DV	-.002 (.001)	-.002** (.001)	.004** (.001)
Tenure dummies and interactions	No	Yes	Yes
Observations	1,176,688	1,176,688	267,301
R^2	.60	.60	.62
Panel B			
δ_0 : EXSH \times DV	.152*** (.024)	.154*** (.025)	.063*** (.021)
δ_1 : EXSH \times DV \times EXPER	.013** (.005)	.013** (.006)	.022*** (.007)
EXPER	-.003 (.003)	-.005 (.003)	-.008*** (.003)
EXPER \times DV	-.001 (.005)	.000 (.006)	.007 (.006)
Tenure dummies and interactions	No	Yes	Yes
Observations	1,176,688	1,176,688	267,301
R^2	.65	.65	.67

NOTE.—INVIND panel, years 1989–97. One observation is a worker-year. All regressions include firm fixed effects. EXSH is the share of sales that is exported. In panel A, $EXPER_{it}$ measures the total number of years a worker was employed at exporting firms as of year t . In panel B, $EXPER_{it}$ measures 5-year cumulated past export experience for worker i in year t as described in detail in Sec. IV. Additional controls include age and age squared, gender, white-collar indicator, manager indicator, log employment, log of domestic sales, and indicators for year, industry, and region. Standard errors, clustered by firm-year, are reported in parentheses.

** Statistically significant at the 5% level.

export share, we allow for exports to have heterogeneous effects across workers according to their export experience: in the post-devaluation period, $\partial w_{it} / \partial EXSH_{j(i,t)} = \delta_0 + (\delta_1 \times EXPER_{it})$. In panel A of table 9, we present results using the first measure of export experience, and in panel B we use the second measure based on export shares. All results are from the firm fixed effects regressions, including the usual set of controls; we cluster the standard errors at the firm-year level. Column 1 in both panels A and B shows that our main coefficient of interest, δ_1 , is positive and statistically significant, indicating that the export wage premium increases with a worker's past export experience. Within firms, $EXPER_{it}$ varies both cross-sectionally (because workers vary in their tenure at the firm) and longitudinally for

workers with the same tenure but who were hired in different years. Even though our regressions do include a quadratic tenure term, it is possible that $EXPER_{it}$ is picking up a tenure effect if the quadratic term does not fully capture workers' tenure profiles. In particular, for workers at their first job, $EXPER_{it}$ will grow with tenure (unless the firm has zero export). Thus, it might then be the case that workers with longer tenure receive a larger share of the extra rent generated by the increased exports during the devaluation period. To account for this possibility, in column 2, we include a full set of tenure dummies and interactions of each of these dummies with $EXSH_{j(i,t)} \times DV$; that is, we allow for the post-devaluation export share to affect workers with different tenure at the firm differently. Comparing columns 1 and 2, we see that the estimates of δ_1 are virtually unchanged in both panels A and B. In the INPS-INVIND data, tenure is measured precisely for workers who joined their firm in 1981 or subsequent years, but it is censored for those who were in the firm's workforce as of 1980, the first year in the data set.¹⁸ In column 3, we report results from the same specification as in column 2 but limiting the sample to the cohorts of workers who joined their current employer after 1980, for whom tenure is precisely measured. In panel A, our coefficient of interest becomes small and statistically insignificant. However, in this case tenure and export experience defined by number of years spent at an exporting firm are very collinear. In panel B, where we use our preferred measure of export experience that, among other advantages, breaks this collinearity, the estimated coefficient is positive and strongly statistically significant, and its magnitude is even larger compared to the previous estimates. A one standard deviation increase in export experience increases wages by about 1.1% for a worker in a firm with an export share equal to the sample mean.

Thus, we find that the export wage premium is larger for workers with higher cumulated past export experience. This finding, and its robustness to controlling for tenure, corroborates our interpretation of the previous results that part of the export wage premium is due to an increase in the market value of workers' skills specifically related to export activity.

V. Conclusions

We exploited the large and unexpected devaluation of the Italian lira in 1992 to study the effect of firms' exporting activity on wages. Our results indicate that the increase in the export share of sales induced by the 1992 devaluation did cause an increase in wages and that this effect was due to

¹⁸ Columns 1 and 2 included a dummy variable for workers who were already employed as of 1980, as well as trends based on these workers' ages in 1980 as described in footnote 9.

both exporting firms paying a wage premium and to changes in the market value of workers' unobservable skills. A novel contribution of this paper is to show that this result depends crucially on whether one allows the returns to individual workers' unobservable skills to vary over time.

The "rent sharing" effect is consistent with theoretical models that emphasize the role of firm heterogeneity and frictions and that predict an effect of trade on wage dispersion across occupations, industries, and firms (Helpman and Itskhoki 2010). The "skill composition" effect, together with our finding that the workers who benefited the most were those with more export-related past experience, documents the importance of export-specific skills, which are typically not observed in traditional data sets. Moreover, our finding that the change in the export share in the devaluation period had a significantly stronger effect on workers with more past export experience indicates that "workforce composition effects" are in part due to increases in the market value of export-related skills. This is another novel contribution of this paper, and the result further supports a causal interpretation of the effects of export on wages.

In addition to providing new evidence on the relationship between exporting and wages, our paper has implications for both future empirical and theoretical analyses. On the empirical front, researchers have only recently started exploring the heterogeneous effects of export shocks within industries, occupations, and firms (Frias et al. 2012; Hummels et al. 2014; Helpman et al. forthcoming); future research should aim at obtaining more precise measures of export-specific skills. On the theoretical side, our findings suggest that labor market frictions and export-specific skills should be essential ingredients in models of the effects of international trade on wages.

References

- Abowd, John M., Robert H. Creedy, and Francis Kramarz. 2002. Computing person and firm effects using linked longitudinal employer-employee data. Technical paper no. TP-2002-06, US Census Bureau, Suitland, MD.
- Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67, no. 2:251–333.
- Amiti, Mary, and Donald R. Davis. 2011. Trade, firms, and wages: Theory and evidence. *Review of Economic Studies* 79, no. 1:1–36.
- Araújo, Bruno César, and Lourenço Paz. 2014. The effects of exporting on wages: An evaluation using the 1999 Brazilian exchange rate devaluation. *Journal of Development Economics* 111, no. C:1–16.
- Baumgarten, Daniel. 2013. Exporters and the rise in wage inequality: Evidence from German linked employer-employee data. *Journal of International Economics* 90, no. 1:201–17.

- Bernard, Andrew, J. Bradford Jensen, J., and Robert Z. Lawrence. 1995. Exporters, jobs, and wages in US manufacturing, 1976–1987. In *Brookings Papers on Economic Activity: Microeconomics*, 67–119. Washington, DC: Brookings Institute Press.
- Bugamelli, Matteo, Fabiano Schivardi, and Roberta Zizza. 2010. The euro and firm restructuring. In *Europe and the euro*, ed. Alberto Alesina and Francesco Giavazzi, 99–138. Chicago: University of Chicago Press.
- Bustos, Paula. 2011. Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms. *American Economic Review* 101, no. 1:304–40.
- Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics* 131, no. 2: 633–86.
- Card, David, Jörg Heining, and Patrick Kline. 2013. Workplace heterogeneity and the rise of West German wage inequality. *Quarterly Journal of Economics* 128, no. 3:967–1015.
- Castellani, Davide, Francesco Serti, and Chiara Tomasi. 2010. Firms in international trade: Importers and exporters heterogeneity in Italian manufacturing industry. *World Economy* 33, no. 3:424–57.
- Cosar, A. Kerem, Nezhil Guner, and James Tybout. 2016. Firm dynamics, job turnover, and wage distributions in open economies. *American Economic Review* 106, no. 3:625–63.
- Crinò, Rosario, and Paolo Epifani. 2012. Productivity, quality, and export behaviour. *Economic Journal* 122, no. 565:1206–43.
- Davis, Donald R., and James Harrigan. 2011. Good jobs, bad jobs, and trade liberalization. *Journal of International Economics* 84:26–36.
- De Loecker, Jan. 2011. Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity. *Econometrica* 79, no. 5:1407–51.
- Egger, Hartmut, and Udo Kreickemeier. 2009. Firm heterogeneity and the labor market effects of trade liberalization. *International Economic Review* 50, no. 1:187–216.
- Finicelli, Andrea, Alessandra Liccardi, and Massimo Sbracia. 2005. A new indicator of competitiveness for Italy and the main industrial and emerging countries. In *Bank of Italy, supplements to the statistical bulletin: Methodological notes* 15, no. 66, Bank of Italy.
- Frías, Judith A., David S. Kaplan, and Eric Verhoogen. 2009. Exports and wage premia: Evidence from Mexican employer–employee data. Unpublished manuscript, Columbia University.
- . 2012. Exports and within-plant wage distributions: Evidence from Mexico. *American Economic Review: Papers and Proceedings* 102, no. 3: 435–40.

- Helpman, Elhanan, and Oleg Itskhoki. 2010. Labour market rigidities, trade and unemployment. *Review of Economic Studies* 77, no. 3:1100–37.
- Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen Redding. Forthcoming. Trade and inequality: From theory to estimation. *Review of Economic Studies*.
- Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang. 2014. The wage effects of offshoring: Evidence from Danish matched worker-firm data. *American Economic Review* 104, no. 6:1597–1629.
- Iranzo, Susana, Fabiano Schivardi, and Elisa Tosetti. 2008. Skill dispersion and firm productivity: An analysis with employer-employee matched data. *Journal of Labor Economics* 26, no. 2:247–85.
- Krishna, Pravin, Jennifer P. Poole, and Mine Zeynep Senses. 2011. Trade liberalization, firm heterogeneity, and wages: New evidence from matched employer-employee data. World Bank Policy Research Working Paper no. 5711, World Bank, Washington, DC.
- . 2014. Wage effects of trade reform with endogenous worker mobility. *Journal of International Economics* 93, no. 2:239–52.
- Kugler, Maurice, and Eric Verhoogen. 2012. Prices, plant size, and product quality. *Review of Economic Studies* 79, no. 1:307–39.
- Lazear, Edward, Kathryn Shaw, and Christopher Stanton. 2015. The value of bosses. *Journal of Labor Economics* 33, no. 4:823–61.
- Mion, Gioridano, and Luca David Opromolla. 2011. Managers' mobility, trade status, and wages. CEPR Discussion Paper no. DP8230, Centre for Economic Policy Research, London.
- Olley, G. Steven, and Ariel Pakes. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64, no. 6:1263–97.
- Ouazad, Amine. 2008. A2REG: Stata module to estimate models with two fixed effects, Statistical Components S456942. Boston College, Department of Economics.
- Park, Albert, Dean Yang, Xinheng Shi, and Yuan Jiang. 2010. Exporting and firm performance: Chinese exporters and the Asian financial crisis. *Review of Economics and Statistics* 92, no. 4:822–42.
- Schank, Thorsten, Claus Schnabel, and Joachim Wagner. 2007. Do exporters really pay higher wages? First evidence from German linked employer-employee data. *Journal of International Economics* 72, no. 1:52–74.
- Verhoogen, Eric A. 2008. Trade, quality upgrading, and wage inequality in the Mexican manufacturing sector. *Quarterly Journal of Economics* 123, no. 2:489–530.
- Yeaple, Stephen. 2005. A simple model of firm heterogeneity, international trade, and wages. *Journal of International Economics* 65, no. 1:1–20.