EXPORTS AND WAGE PREMIUMS: EVIDENCE FROM MEXICAN EMPLOYER-EMPLOYEE DATA

Judith A. Frías, David S. Kaplan, Eric Verhoogen, and David Alfaro-Serrano*

Abstract—This paper draws on employer-employee and longitudinal plant data from Mexico to investigate the impact of exports on wage premiums, defined as wages above what workers would receive elsewhere in the labor market. We decompose plant-level average wages into a component reflecting skill composition and a component reflecting wage premiums. Using the late-1994 peso devaluation interacted with initial export propensity as a source of exogenous changes in exports, we find that exports have a significant positive effect on wage premiums and that the effect on wage premiums accounts for essentially all of the medium-term effect of exporting on plant-average wages.

I. Introduction

A growing body of research suggests, in contrast to the textbook model of perfectly competitive labor markets, that firms' wage policies matter for how much individual workers are paid. Following the seminal work of Abowd, Kramarz, and Margolis (1999, hereafter AKM), several recent studies have fit simple models with individual and firm fixed effects and have found that the firm effects account for a substantial fraction of overall wage variation (Card, Heining, & Kline, 2013; Card, Cardoso, & Kline, 2015; Abowd et al., 2018; Alvarez et al., 2018; Barth, Davis, & Freeman, 2018; Sorkin, 2018; Song et al., 2019; Lachowska, Mas, & Woodbury, 2020). Firms appear to pay different wage premiums, defined as wages above what individuals would earn elsewhere on the labor market, and the evolution of those premiums appears to be important for the broader evolution of wage structures.

Received for publication February 4, 2020. Revision accepted for publication November 2, 2021. Editor: Rema N. Hanna.

*Frías: Consejo de la Judicatura Federal; Kaplan: Inter-American Development Bank; Verhoogen: Columbia University; Alfaro-Serrano: Cornerstone Research.

This paper began many years ago, and has accumulated many debts along the way. We thank, without implicating, Rita Almeida, Paula Bustos, David Card, Bo Honoé, Leonardo Iacovone, Wojciech Kopczuk, Thomas Lemieux, Arik Levinson, Steven Lich-Tyler, Leigh Linden, Osmel Manzano, Marc Melitz, Marcelo Moreira, Marc-Andreas Muendler, Daniel Parent, Cristian Pop-Eleches, Jesse Rothstein, Matthias Schuendeln, Matthew Slaughter, John Strauss, John Sutton, Catherine Thomas, Miguel Urquiola, Till von Wachter, many seminar participants, the editor (Rema Hanna), and two anonymous referees for helpful comments. We also thank Gerardo Durand, José Antionio Mejía, Jorge Alberto Reyes, Gerardo Leyva, Abigail Durán, Adriana Ramírez, and Gabriel Romero of INEGI for help with the first versions of the paper, and Natalia Volkow and other staff of the INEGI Microdata Laboratory for assistance after a long hiatus due in part to a change in data-access regime. All analysis of confidential data was carried out on-site and in accordance with Mexican confidentiality laws. Guru Sethupathy, Hitoshi Shigeoka, Kensuke Teshima, Bruno Velloso, and Matt Wai-Poi provided excellent research assistance. E.V. is indebted to the National Science Foundation (SES-0721068) for funding. An earlier version of this paper was presented at an NBER ITI meeting in 2005, a UC Berkeley labor lunch in 2009, and a number of other seminars. The views expressed in this paper are our own and do not necessarily reflect the views of Cornerstone Research, the Instituto Mexicano del Seguro Social, Consejo de la Judicatura Federal, or the Inter-American Development Bank.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01178.

Despite the recent advances, there has been comparatively little empirical work on *why* firms pay wage premiums. What leads some firms to pay more than others to similar workers? Much of the recent literature has aimed to characterize the extent of rent sharing with different groups of workers or the contribution of firms to overall wage inequality, but not to identify the causal determinants of wage premiums.[1] The recent review by Card et al. (2018) concludes with a call for research to fill this gap.[2]

In this paper, we draw on a combination of employer-employee and longitudinal plant data from Mexico to investigate one possible determinant: firms' engagement in export markets. It is well documented that exporting firms tend to pay higher wages on average within narrow industries.[3] There is growing evidence that this relationship is causal—that exogenous increases in exporting lead firms to pay higher wages on average (Verhoogen, 2008; Álvarez & López, 2009; Kandilov, 2009; Bustos, 2011a; Brambilla, Lederman, & Porto, 2012). But it remains unclear to what extent these plant-level results reflect changes in wage premiums as opposed to changes in workforce composition. Several studies (discussed briefly below) have examined the relationship between exporting and wages in employer-employee data, but quasi-experimental evidence of the effect of exporting on wage premiums remains thin.

We proceed in two steps. First, following Card et al. (2013, hereafter CHK), we fit simple AKM-type models in different periods and decompose plant-level wages into a plant component, which we interpret as a plant-specific average wage premium, and an average person component reflecting average individual characteristics, including individual effects that capture time-invariant ability.[4] Because of data constraints, discussed below, our baseline specification focuses on the periods 1992 to 1994, 1996 to 1998, and 2000

[1] See for example, Card et al. (2015) on differences in rent sharing for different groups, and Card et al. (2013), Barth et al. (2016, 2018), and Song et al. (2019) on characterizing the role of firms in overall wage inequality. Lazear and Shaw (2018) provide a useful overview.

[2] "The field continues to rely almost exclusively on observational studies predicated on plausible, but ultimately debatable, identifying assumptions. More research is needed applying (quasi-)experimental research designs that convincingly tease out the mechanisms through which firm shocks are transmitted to workers" (Card et al., 2018, p. 24).

[3] The first papers to document this pattern are Bernard and Jensen (1995, 1999). Brambilla, Chauvin, and Porto (2017) review the literature and present evidence from 61 developing countries of a robust positive relationship between exporting and firm-level wages.

[4] In the Mexican manufacturing sector, the vast majority of firms are single-establishment. For instance, the 2009 Economic Census in Mexico reports that more than 75% of manufacturing firms with 51 or more employees and more than 90% of manufacturing firms with 11 or more employees are single-establishment. The data we use are at the plant level, and unfortunately we do not know which firm a given plant is part of. Hereafter we treat plants as decision makers.

The Review of Economics and Statistics, March 2024, 106(2): 305–321

https://doi.org/10.1162/rest_a_01178

to 2002, which we refer to, respectively, as periods 1, 2, and 3. We conduct diagnostic exercises similar to CHK and find that the data appear to be well described by a model with additive individual and plant fixed effects within several-year periods.

Second, following Verhoogen (2008), we use the late-1994 peso devaluation as a source of exogenous changes in exports in order to identify the effect of exporting on wage premiums. The key idea is that not all Mexican plants were equally well positioned to take advantage of the devaluation to increase exports. Some were already active in the export market or on the margin for entering, and others had little hope of being able to export profitably. Empirically, we need a measure of the propensity to increase exports that does not reflect plants' endogenous responses to the devaluation itself. Our preferred measure is predicted export status before the devaluation, estimated flexibly using contemporaneous plant characteristics. We interpret the interaction of initial predicted export status and an indicator for the devaluation as capturing the differential inducement to export generated by the peso devaluation. We estimate instrumental variables (IV) regressions of within-plant changes in wage premiums (and other variables) on within-plant changes in export share instrumented by this interaction. These regressions effectively compare differential changes by initial export propensity between periods 1 and 2, which span the devaluation, to differential changes between periods 2 and 3, when there was no devaluation.

Our main finding is that exporting has a robust, positive, statistically significant effect on wage premiums. The magnitude is economically significant: a 1 percentage point increase in a plant's export share (on an initial mean of 4% overall, or 15% conditional on exporting) is associated with approximately a 1% increase in plant-level wage premiums. Our estimates suggest that essentially all of the effect of exporting on average wages at the plant level is explained by the increase in wage premiums, rather than by changes in workforce composition, at least for the medium-term time frame we are able to study. These results are robust to using several different measures of export propensity and to including an earlier, pre-crisis period (1988–1990) in the analysis. We conclude that exporting is an important determinant of wage premiums among Mexican manufacturing firms.

Two broader implications seem particularly salient. First, for the growing literature on the role of firms in wage setting, our findings underline the causal role of product markets in shaping the evolution of wage distributions through their effect on firms' wage policies. Product market shocks are potentially related to, but are conceptually distinct from, shocks to productivity, which have been the focus of much of the rent-sharing literature (reviewed, for instance, by Card et al., 2018). Second, the paper contributes to a growing body of evidence about how firms' responses to trade shocks contribute to wage inequality. Previous work has argued that trade liberalization can increase wage dispersion across firms within industries in developing countries, as larger, more productive firms, which already tend to be higher wage, take greater advantage of export opportunities, which leads them to raise wages further (Verhoogen, 2008; Helpman, Itskhoki, & Redding, 2017). This paper shows that the increase in across-plant dispersion in our setting is largely driven by plants' wage policies rather than sorting on worker ability.

This paper is mainly concerned with documenting the reduced-form relationship between exporting and wage premiums, rather than arguing for a particular theoretical mechanism. Several theoretical channels are consistent with the empirical patterns. One is differential quality upgrading: in a Melitz (2003)-type context, if firms sell higher-quality goods to richer consumers on the export market and if producing higher-quality goods requires paying higher-efficiency wages, then we would expect to see wage premiums rise with exports (Verhoogen, 2008).[5] A second possibility is fair-wage concerns: if the exogenous increase in exporting is associated with an increase in profitability (as in many Melitz (2003)-type models), firms may share profits in order to induce workers to reciprocate with effort (Akerlof & Yellen, 1990; Egger & Kreickemeier, 2009, 2012; Amiti & Davis, 2012). A third possibility is that search frictions give rise to match-specific rents and firm-worker bargaining, which may lead wages to increase with exports.[6] A fourth is that workers' idiosyncratic attachments to firms give firms monopsony power in setting wages and generate a positive effect of product-demand shocks on wages and wage premiums, as in the wage-posting model of Card et al. (2018).[7] A fifth is that managers pay high wages in order to ensure themselves a quiet life rather than to maximize profits (Bertrand & Mullainathan, 2003).[8] These mechanisms (and others) are often classified under the general heading of "rent sharing." In this paper, we do not take a stand on precisely which mechanism is correct. Nevertheless, we believe that our results carry an important implication for theoretical discussions by providing evidence against the supposition (traditional in the field of international trade) of perfectly competitive, frictionless labor markets, in which any relationship between exporting and wages at the firm level is explained only by the sorting of workers by skill.

[5] Verhoogen (2008, section III.B) suggests that the need to induce higher effort through higher-efficiency wages is one of several possible reasons why producing higher-quality goods may require paying higher wages.

[6] In Helpman et al. (2010), firms employ a fixed-cost screening technology; as scale increases with increased exports, firms screen more intensively, which increases the revenue per worker to be bargained over, which in turn increases wages. In Coşar et al. (2016), firms face labor adjustment costs, and the benefit of filling a position is greater for expanding firms than shrinking ones, leading the bargained wage to be higher in such firms; to the extent that increased exporting leads firms to expand, wages will also increase. See also Davidson, Matusz, and Shevchenko (2008), Felbermayr, Prat, and Schmerer (2011), Fajgelbaum (2016), and Bellon (2017).

[7] Although the Card et al. (2018) framework does not explicitly consider exporting, it would be straightforward to extend it to do so. The model also has the advantage that it provides an explicit microfoundation for the AKM models we estimate below.

[8] This list of possible mechanisms is not exhaustive. Another is simply that workers are unionized and are able to bargain for a share of firm profits.

In addition to the research cited above, this paper is related to several streams of literature. In an early contribution using firm-level data from collective bargaining contracts in Canada, Abowd and Lemieux (1999) employed industry-level import and export prices as instruments for quasi-rents per worker (a measure of profitability when labor is valued at its alternative wage) and found that greater quasi-rents led to higher wages. In panel data on British firms, Van Reenen (1996) related major innovations, which are arguably exogenous, to wage changes at the firm level and found that quasi-rents due to innovation were passed through to higher wages. These early papers lacked data on individual workers and hence were not able to definitively answer questions about whether the wage changes reflected changes in wage premiums or in skill composition.

A recent paper by Kline et al. (2019) also examines innovation and rent sharing, with employer-employee data and rich patent information from the United States. Under the identifying assumption that the U.S. Patent Office's initial decision on a patent application is essentially random, conditional on observable characteristics of the application and firm, the authors find that each patent-induced additional dollar of operating surplus yields a 29 cent increase in a firm's wage bill. Using a matching estimator in Finnish data, Aghion et al. (2018) find that patents are associated with higher wages for coworkers of inventors. Our paper is complementary to these papers, showing that a very different type of shock—to the set of markets a firm sells into, rather than to the acquisition of patents—can also have sizable impacts on wage premiums.

As noted above, this paper is also related to a number of papers that have related exporting to wages in employer-employee data, including Schank, Schnabel, and Wagner (2007), Munch and Skaksen (2008), Davidson et al. (2014), Baumgarten (2013), Klein, Moser, and Urban (2013), Irarrazabal, Moxnes, and Ulltveit-Moe (2013), Hummels et al. (2014), Krishna, Poole, and Senses (2014), Araújo and Paz (2014), Macis and Schivardi (2016), Helpman et al. (2017), Barth et al. (2018), and Garin and Silvério (forthcoming).[9] Of these, the papers that employ a quasi-experimental (or design-based) strategy to isolate the effect of exogenous variation in exports—which are hence closest in spirit to our paper—are Hummels et al. (2014), Araújo and Paz (2014), Macis and Schivardi (2016), and Garin and Silvério (forthcoming). Hummels et al. (2014) focus on the wage effects of offshoring (in Denmark) but also examine the wage effects of exporting, using a different instrument from ours: a weighted average of imports of particular goods by the countries a firm exports to, using the firm's initial export shares as weights.[10] Macis and Schivardi (2016) and Araújo and Paz (2014) adopt modifications of the strategy of our paper (as presented in an earlier version), but with certain limitations. Macis and Schivardi (2016) seek to implement a design similar to ours using the 1992 devaluation of the Italian lira, but they find little evidence of a differential effect of the devaluation on exporting by firm size; in the end, they regress estimates of wage premiums on current export share, without instrumenting, letting the coefficient on export share differ before and after devaluation. In Brazilian data, Araújo and Paz (2014) examine differential trends between larger and smaller firms in a period after a devaluation (2000–2004) relative to a period before a devaluation (1995–1998), but do not consider changes that span the devaluation year (1999) itself.[11] Garin and Silvério (forthcoming) construct firm-specific export shocks due to differential changes in the import demands of the countries to which Portuguese firms export, in the spirit of Hummels et al. (2014), and relates the shocks to the wage changes of workers initially employed in the corresponding firm, but does not consider the change in AKM-type firm effects as an outcome.

The next section describes our econometric strategy in more detail. Section III describes the data and briefly provides background on the peso crisis. Section IV presents results from the first step of our econometric procedure, estimating wage premiums at the plant level. Section V presents results from the second step, estimating the effect of exports on the estimated wage premiums. Section VI concludes.

II. Econometric Strategy

Our estimation strategy has two parts. We first use the administrative employer-employee data to decompose plant-level wages into a plant component due to wage premiums and an average person component due to skill composition. We then relate changes in those components to the export shock brought about by the peso devaluation, linking the employer-employee results to longitudinal survey data on manufacturing plants.

Within a given (several-year) period, we assume that the log wage of person i in time t is given by

$$w_{it} = \alpha_i + \psi_{J(i,t)} + X_{it}'\beta + \varepsilon_{it}, \tag{1}$$

where α_i is a time-invariant individual effect; $J(i,t)$ indicates the plant in which person i is employed in year t; $\psi_{J(i,t)}$ is the corresponding plant effect; X_{it} is a vector of time-varying

[9] An earlier version of the current paper (Frías, Kaplan, & Verhoogen, 2009) estimated a model that allowed more flexibility for changes over time in the return to individual ability, using an approach pioneered by Holtz-Eakin, Newey, and Rosen (1988), and found results broadly similar to those reported here. In this version, delayed in part by a change in data access regime, we focus on simple AKM-type specifications estimated in separate periods because the method is arguably more transparent and more comparable to the literature as it has developed more recently. Many of the papers cited above refer to the earlier version of this paper. In a related short paper, Frías et al. (2012) show (without controlling for individual effects) that exports increase within-plant wage inequality.

[10] As the authors acknowledge, this strategy is subject to the concern that demand shocks are correlated between Denmark and the trading partners.

[11] As a result, Araújo and Paz (2014) capture delayed responses to the devaluation by firm size, but not the short-term differential impact of the devaluation. Further differences are that they have access to a limited set of firm characteristics (lacking, for example, domestic sales and capital intensity) and do not employ an IV strategy.

observables; and ε_{it} is a mean-zero error term. This specification follows AKM and CHK.[12]

We interpret the individual effect α_i as portable ability, invariant during a given several-year period, compensated equally in all firms. The plant effect $\psi_{J(i,t)}$ reflects a premium (or discount) paid by plant $j = J(i,t)$ to all employees. In X_{it}, we include tenure and its square, year effects, and a polynomial in age. Following CHK and Card et al. (2015), we drop the linear term (which is collinear with the individual and year effects) and include the square and cube of recentered age (age −40). We define the "person" component to be the individual effect plus the contribution of the other observables: $s_{it} = \alpha_i + X_{it}'\beta$.

The identifying assumption is that the error term ε_{it} is uncorrelated with the other covariates in all years (within a given several-year period), a sufficient condition for which is

$$E(\varepsilon_{it}|X_{i1}\ldots X_{iT}, \psi_1\ldots\psi_J, \alpha_i) = 0. \tag{2}$$

In the employer-employee literature, equation (2) is referred to as a *conditional random mobility* assumption, since it requires that, conditional on observables, an individual's current-period idiosyncratic shock is uncorrelated with the plant in which she is employed. While equation (2) rules out sorting across plants on the basis of contemporaneous shocks to individual ability, it allows for many salient forms of sorting. For instance, as CHK point out, it allows for sorting on time-invariant ability (e.g., higher turnover among lower-ability workers), as well as for the possibility that workers are more likely on average to move from low- to high-wage plants than vice versa (since we condition on $\psi_1\ldots\psi_J$). In section IV, we present evidence, following CHK, that the model, equations (1) and (2), summarizes well the wage and mobility patterns in the data.

We face an important choice about how to define the several-year periods within which to estimate this model. As discussed in section III, the employer-employee data are available for 1985 to 2005, but the plant panel with export information is available on a consistent basis only for 1993 to 2003. The peso devaluation occurred in December 1994, and the crisis played out over the next several months. We need the first period to be clearly pre-devaluation, but we also need information on export status. Our preferred solution is to use 1992 to 1994 as the pre-crisis period. We do not use the crisis year 1995. To maintain equally spaced periods, we use 1996 to 1998 as the first post-crisis period and 2000 to 2002 as the second post-crisis period. We will refer to 1992 to 1994, 1996 to 1998, and 2000 to 2002 as periods 1, 2, and 3, respectively.[13]

As discussed in Abowd, Creecy, and Kramarz (2002), unique solutions for the estimates of the person and plant effects can only be obtained within a set of plants linked by worker switchers during the period (a connected set). As has become standard in the literature, we focus on the largest connected set of plants in each period. We will see that the largest connected sets capture a smaller share of plants and workers than has typically been the case in developed-country settings, but also that they capture almost all of the plants in the plant panel (for which we have export information).

Under assumption (2), for plant j' in year t', the expected plant-level wage can be expressed as the sum of two components,

$$E(w_{it}|j',t') = \psi_{j'} + E(\underbrace{\alpha_i + X_{it}'\beta}_{=s_{it}}|j',t'), \tag{3}$$

where the conditioning on $X_{i1},\ldots,X_{iT},\psi_1,\ldots,\psi_J,\alpha_i$ is omitted but should be understood. The first term is the plant component, the wage premium paid by the plant. The second term is the expected person component, which captures average workforce skill. To generate sample analogues of these components, we fit model (1) and recover the estimated plant components and a plant-year-level average of the estimated person components, which we refer to as the average person component. Computationally, we estimate equation (1) using the preconditioned conjugate gradient (PCG) algorithm in Matlab, following CHK. We estimate the model separately by (three-year) period.[14]

It is worth emphasizing that although the individual effects, α_i, are assumed to be constant within three-year periods, they are allowed to vary across periods. Hence, if the peso devaluation generated an increase in the general-equilibrium return to skill in the economy, it would be captured by an increase in the individual effects, α_i, for higher-skilled individuals relative to the α_i for lower-skilled individuals from period 1 to period 2.

Once we have recovered the plant and average person components, the next step is to determine the effect of exporting on them. As noted above, our strategy uses the late-1994 peso devaluation as a source of exogenous variation in the incentive to export. The devaluation represented a major shock to the Mexican economy and was unexpected until very shortly before it occurred. The peso lost approximately 50% of its nominal value in a few days in December

[12]Card et al. (2018) provide a theoretical justification for this specification in a model in which workers have idiosyncratic attachments to particular plants and either different worker skill groups are substitutes in the production function or the plants face the same supply elasticity from these different groups.

[13]In a robustness check using the employer-employee data alone, we will also present estimates for 1988 to 1990, which we refer to as period 0. The three-year windows are narrower than have typically been considered in the literature, but we note that others have used even narrower windows (Lachowska et al., 2023). It is important to acknowledge that because of the short windows, there may be few switchers and hence limited mobility bias. We return to this issue below.

[14]Estimating the plant components requires a normalization in each year. (Intuitively, the question is which of the many plant effects to omit from the regression.) When one seeks to compare plant effects for different subpopulations, as in Card et al. (2015), the normalization is an important choice. In our case, we do not compare across groups within year, and we will include year effects, which will absorb the normalization, in all regressions below. Our estimates are thus invariant to the choice of normalization.

FIGURE 1.—REAL EXCHANGE RATE

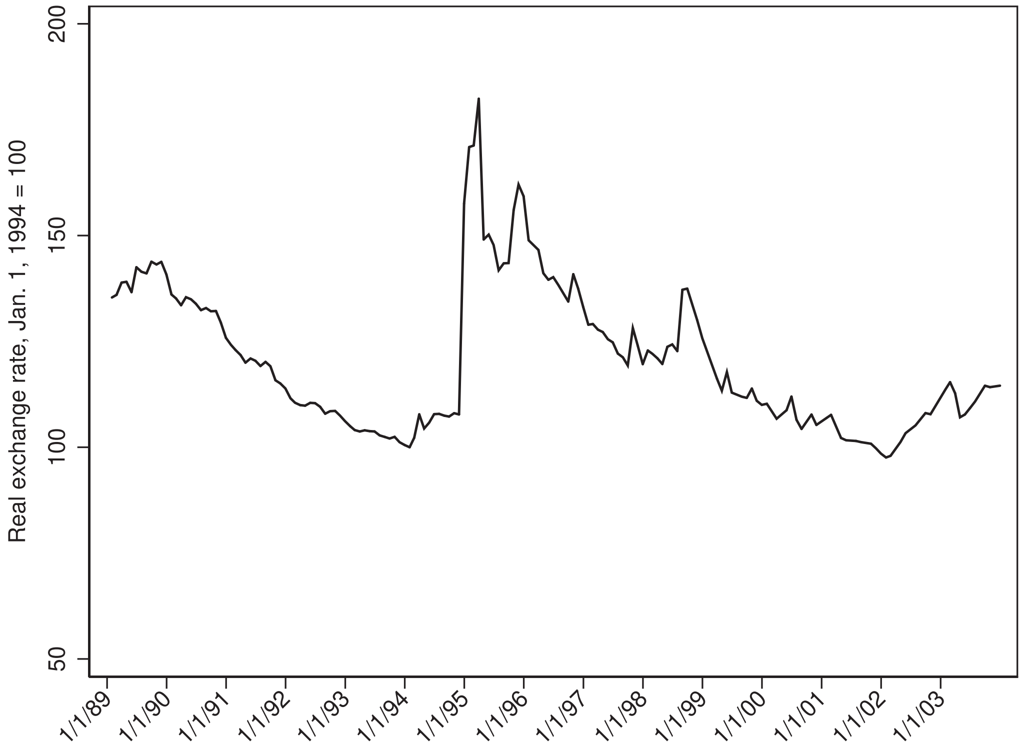

Real exchange rate calculated as $e\times$CPI(US)/CPI(Mexico), where e is peso/US\$ nominal exchange rate. Data from IMF International Financial Statistics.

1994. Figure 1 plots the real exchange rate over the 1989–2004 period. Exports rose sharply, with approximately 85% destined for the U.S. market. Using a balanced panel of manufacturing plants, the construction of which will be explained in section III, figure 2 illustrates the shift toward the export market: average export share (calculated as total exports over total sales for the panel as a whole) jumped sharply, and the share of plants with positive exports rose from approximately 30% to 45% of the sample.[15]

To use the devaluation as a source of variation in our context, we need to focus on differences in its effects at the plant level. As discussed briefly above, a number of heterogeneous-firm models in the spirit of Melitz (2003) have the feature that initially larger, more capable plants, which have a higher propensity to export, are more likely to increase exports in response to a shock such as a devaluation. In a subset of these models, the increase in exports is associated with an increase in wage premiums.[16] In such frameworks, a devaluation generates a greater increase in the incentive to export—and in exports—for more capable, larger, higher-export-propensity plants than for less capable, smaller, lower-export-propensity ones.

Motivated by these observations, we adopt an econometric specification of the following form,

$$\triangle y_{jp} = \theta \triangle e_{jp} + \gamma \widehat{\lambda}_{jp-1} + \xi_{kp} + \zeta_{rp} + u_{jp}, \tag{4}$$

where j indexes establishments and p indexes periods; $\triangle y_{jp}$ is the change in an outcome variable (e.g., the plant component or average person component) between periods $p-1$ and p; $\triangle e_{jp}$ is the change in export share between periods $p-1$ and p; $\widehat{\lambda}_{jp-1}$ is a measure of export propensity (or other proxy for plant capability), discussed below; and ξ_{kp}, ζ_{rp}, and u_{jp} are an industry-period effect, a region-period effect, and a mean-zero disturbance, respectively. For variables that vary by year, we average over years within a period. Note that equation (4) is effectively in first differences, with time-invariant firm characteristics differenced out.

There are a number of reasons why the change in export share, $\triangle e_{jp}$, might be correlated with the error term, u_{jp},

[15]The peso crisis was a much larger shock than the North American Free Trade Agreement (NAFTA), which had taken effect in January 1994. Mexico's main trade liberalization came with its entrance into the General Agreement on Tariffs and Trade in the mid-1980s, and by 1994 the vast majority of Mexican imports were covered by tariffs of 20% or less. Average U.S. tariffs on goods from Mexico were on the order of 3% to 5%. In the majority of cases, NAFTA phased out existing tariffs slowly over time. Relative to the exchange-rate devaluation, the year-by-year tariff changes were small.

[16]For instance, in the Melitz (2003)-type theoretical framework of Verhoogen (2008), plants differ in an underlying "capability" parameter, there is a fixed cost of exporting, and firm capability and worker quality are complements in determining product quality. In equilibrium, more capable firms are larger, pay higher wages, produce higher-quality goods, and are more likely to export in cross-section. In response to an exogenous inducement to export such as a devaluation, more capable plants increase exports, increase average product quality, and raise wages relative to less capable firms in the same industry. If producing higher-quality products requires employing especially motivated workers (i.e., if high-quality workers are those who supply high effort, where effort is noncontractible), which in turn requires paying higher efficiency wages, then the increase in exports is accompanied by an increase in wage premiums.

FIGURE 2.—SHIFT TO EXPORT MARKET, EIA PANEL

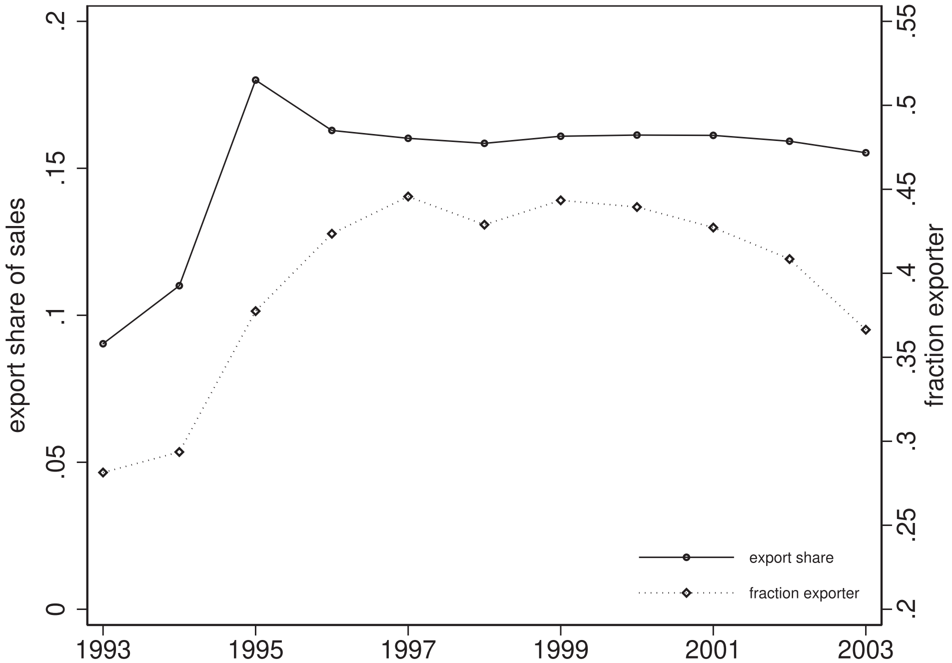

Data from 1993–2003 EIA (balanced) panel of 3,518 plants. Export percentage of sales calculated as (total exports for all plants)/(total sales for all plants). Plants with exports greater than 0 classified as exporters. See the data appendix for details.

biasing OLS estimates of the coefficient of interest, θ. On the one hand, positive productivity shocks at the plant level may lead to both greater exports and higher wages, generating a positive bias in the OLS coefficient. On the other hand, a positive labor supply shock to a plant would be expected to lead to both lower wages and greater exports, generating a negative bias in OLS. Other biases are also possible.[17]

To address the endogeneity of export changes, we instrument $\triangle e_{jp}$ with the interaction of export propensity in the previous period and an indicator for the devaluation shock. Since the devaluation occurred in late 1994, that is, between our periods 1 and 2, the indicator is "on" for $p = 2$ (explaining changes from $p = 1$ to $p = 2$) and 0 otherwise. Letting P_2 denote this indicator, the instrument is $\widehat{\lambda}_{jp-1} \times P_2$. The instrument can be interpreted as capturing the differential inducement to export created by the peso devaluation, which we take to be exogenous. Any stable differential trends for plants with different export propensities (or capabilities) will be captured by the uninteracted $\widehat{\lambda}_{jp-1}$ term in equation (4). The IV estimate of θ will reflect only the differential changes in export share between period 1 and period 2 (spanning the late-1994 devaluation), but will not capture stable trends.[18]

Note that this strategy requires that the differential increase in the incentive to export be larger from period 1 to period 2 (when the devaluation intervened) than from period 2 to period 3 (a period of real-exchange-rate appreciation); it does not require that the differential change be 0 between periods 2 and 3. Interestingly, we will see that there was little differential change in export behavior in the later years, despite the appreciation of the peso. This fact is hinted at already in figure 2, which shows that the average export share was roughly constant from 1996 through 2003, despite the peso appreciation.[19] The asymmetry in response to exchange-rate changes may have to do with the facts that exporting requires upfront, sunk investments and that firms, once they have paid these sunk costs, will strive to maintain export contacts and destination market share in order

[17]For instance, a positive demand shock specific to the domestic market would likely lead to greater investment, a higher capital-labor ratio, higher wages, and a lower export share (since domestic sales appear in the denominator of export share), generating a negative bias in OLS.

[18]In addition to the the IV specification, equation (4), we also report the corresponding reduced form,

$$\triangle y_{jp} = \varphi(\widehat{\lambda}_{jp-1} * P_2) + \widetilde{\gamma}\,\widehat{\lambda}_{jp-1} + \widetilde{\xi}_{kp} + \widetilde{\zeta}_{rp} + \widetilde{u}_{jp}, \tag{5}$$

where $\widetilde{\xi}_{kp}$, $\widetilde{\zeta}_{rp}$, and $\widetilde{u}_{jp}$ are again an industry-period effect, a region-period effect, and a mean-zero error. This reduced-form specification is similar to the approach of Verhoogen (2008). That paper estimated OLS models of the form

$$\triangle y_{jp} = \mu + \widehat{\lambda}_{jp-1}\pi_p + \widetilde{\xi}_k + \widetilde{\zeta}_r + \nu_j \tag{6}$$

separately by period, with $p = 1993$, 1997, 2001, and then compared π_{1997} to π_{2001}. Except for the fact that here we consider three-year periods (1992–1994, 1996–1998, and 2000–2002) in place of individual years (1993, 1997, and 2001), the difference in coefficients $\pi_{2001} - \pi_{1997}$ in that paper maps directly to the OLS estimate of φ in our reduced form, equation (5). Verhoogen (2008) used log domestic sales as the proxy, $\widehat{\lambda}_{jp-1}$, rather than export propensity. We present results using the same proxy in a robustness check in section V.

[19]The fraction of plants with positive exports behaved similarly, with a slower increase at the beginning and tailing off a bit more at the end, but the basic story is the same.

to avoid having to pay the sunk costs again in the future (Roberts & Tybout, 1997; Das, Roberts, & Tybout, 2007). The key point for our purposes is that there was a sharp, exogenously driven, differential increase in exports from period 1 to 2 and not from period 2 to 3.

An important step in the implementation of our approach is to choose the measure of export propensity, $\widehat{\lambda}$, which serves as a proxy for plants' likelihood of responding to the exchange-rate devaluation by increasing exports. In this paper, our preferred measure is predicted export status. We first estimate the following linear-probability model in a preliminary stage:

$$1(e_{jt} > 0) = Z'_{jt}a_k + b_{rt} + c_{kt} + v_{jt}, \tag{7}$$

where k indexes four-digit industries; $1(e_{jt} > 0)$ is an indicator for whether plant j has positive exports in period t; Z_{jt} is a vector of plant characteristics; a_k is a vector of coefficients on those characteristics, which we allow to vary across four-digit sectors; b_{rt} and c_{kt} are state-year and sector-year fixed effects; and v_{jt} is a mean-zero disturbance. In our preferred specification, the vector Z_{jt} contains third-degree polynomials in log employment (hours) and log sales and linear terms for log capital-labor ratio, log value-added per worker, log investment per worker, the white-collar share of employment (hours), and an indicator for whether the plant has 10% or more foreign ownership. Note that only contemporaneous characteristics are included as covariates. The predicted values from this regression, which depend only on contemporaneous observables and the coefficient estimates from equation (7), are an estimate of export propensity. We average these at the period level; the average for plant j in period $p-1$ is the $\widehat{\lambda}_{jp-1}$ term included in equation (4).

To check robustness, we present results using several other proxies. One is predicted export share rather than predicted export status, constructed using export share, e_{jt}, as the outcome variable in equation (7). A potential concern is that measurement error in sales could lead to a spurious positive correlation between average sales in $p-1$ and the change in average export share (which has total sales in the denominator) from $p-1$ to p. Similarly, measurement error in employment (from the EIA survey) could generate a spurious positive correlation between average employment in $p-1$ and the change in the average capital labor ratio or the average hourly wage in the EIA survey from $p-1$ to p (which are used as outcomes in equation (4) in some specifications). One would expect predicted export status to be less subject to these concerns than predicted export share (and indeed is our preferred measure of export propensity for this reason), but as further checks, we present results in which sales or employment (and variables derived from them) are omitted from the preliminary stage. To facilitate comparison with previous work, we also consider plant size and total factor productivity as proxies for plant capability.

We conduct three further robustness checks. First, we include an earlier period, 1988 to 1990 (which we refer to as period 0), for which only the employer-employee data are available, in addition to the three periods we have already identified as periods 1, 2, and 3. This allows us to compare a "placebo" change between two preshock periods (periods 0 and 1) to the change between periods that span the shock (periods 1 and 2). Because we do not observe export share in the employer-employee data, we use the reduced-form model, equation (5) in note 18, for this check, as described below. Second, we estimate the effect of exporting on the wages of stayers, workers who are continuously employed in a given firm over two periods. This specification does not use the AKM methodology and does not require the conditional random mobility assumption, equation (2). It is analogous to estimating a model with job-spell fixed effects. Third, we check the assumption of a linear relationship between predicted export status and changes in export share and other outcomes embedded in our model, equation (4). In the presence of fixed costs of exporting, the relationship may not be linear or even monotonic.[20] We present estimates similar to our baseline model but include indicators for quartile of the predicted export status distribution in place of the linear term. All three of these further robustness checks yield results similar to our baseline estimates.

Our causal interpretation of the results relies on the assumption that the devaluation affected wage outcomes differentially within industries only through its impact on the incentives of plants to export. Verhoogen (2008) considered a number of reasons why this assumption might be violated, if, for instance, the devaluation affected larger and smaller firms within industries differently for reasons unrelated to exporting. Readers are referred to that paper for more extensive discussion, but two points are worth reemphasizing. First, the assembly-for-export (*maquiladora*) sector in Mexico provides a sort of additional placebo test. Maquiladoras exported essentially all of their output both before and after the devaluation, so there was not a differential within-industry shock to exporting in this broad sector. But if the macro shock had a differential within-industry impact through a channel other than exporting, one would expect it to show up among maquiladoras as well as nonmaquiladoras. Consistent with our interpretation, there was no differential change in plant-level wages between larger and smaller *maquiladora* plants during the peso-crisis period. Second, it does not appear that differential access to credit markets can explain the empirical patterns. While there is evidence that exporting plants faced a lower cost of capital than nonexporters, likely due to greater access to foreign capital, they also had a greater share of dollar-denominated loans before the crisis and hence their balance

[20]For instance, Bustos (2011b) finds that Argentinian firms in the third quartile of the firm-size distribution respond more than those in the highest quartile to a multilateral trade liberalization.

TABLE 1.—SUMMARY STATISTICS, IMSS INDIVIDUAL-LEVEL DATA, LARGEST CONNECTED SETS

					avg. daily wage (raw, 2002 pesos)		avg. daily wage (winsorized, 2002 pesos)	
year	# individuals	# establishments	avg. age	fraction male	mean	std. dev.	mean	std. dev.
1988	4,813,280	219,015	31.44	0.72	152.90	727.68	119.24	63.05
1989	5,474,761	235,765	30.97	0.70	158.50	640.48	130.00	74.59
1990	6,202,204	253,573	30.76	0.69	152.04	427.17	128.59	79.09
1992	7,007,431	298,743	30.81	0.68	169.47	271.40	149.05	95.53
1993	6,929,394	301,639	31.11	0.68	190.29	254.97	159.40	106.09
1994	7,109,067	302,213	31.29	0.68	201.12	266.48	162.78	110.33
1996	7,260,643	296,641	31.58	0.67	147.44	193.03	116.93	81.88
1997	7,816,910	309,762	31.54	0.67	148.37	201.05	118.07	84.10
1998	8,205,521	316,307	31.66	0.67	150.85	182.16	120.80	86.50
2000	9,396,001	360,718	31.93	0.65	160.49	185.94	130.33	92.57
2001	9,261,411	374,332	32.47	0.65	169.14	191.87	138.23	97.93
2002	9,303,348	380,401	32.84	0.65	171.96	193.72	140.95	99.00

Sample is from IMSS employer-employee records after cleaning steps 1–7 in appendix A.1. Sample includes individuals who earn a positive wage in an establishment for which municipality and industry are observed, are ages 14 to 64, and are employed in an establishment in the largest connected graph of establishments in their respective period (among 1988–1990, 1992–1994, 1996–1998 and 2000–2002). Winsorization is at the 10th and 90th percentiles. Wages are reported in both "raw" (i.e., pre-winsorized) and winsorized form. See section III and appendix A.1 for further details. Average 2002 exchange rate: 9.60 pesos/US$1.

sheets were more adversely affected by the devaluation. These effects appear to have offset. There was no differential within-industry change in the cost of capital that would explain the differential wage changes.[21]

III. Data and Background

The employer-employee data we use are from the administrative records of the Instituto Mexicano del Seguro Social (IMSS), the Mexican social security agency. In principle, all private Mexican employers are required to report wages for their employees to IMSS and to pay social security taxes on the basis of their reports. In practice, only about half of private sector-remunerated employees are reported to IMSS (and hence considered formal), and half are unreported (and considered informal); see appendix table B1. While the size of the informal sector seems large by developed-country standards, it is not out of line for countries at Mexico's income level.[22] At the level of individuals, the IMSS data contain information on age, sex, daily wage (including benefits), and state and year of the individual's first registration with IMSS. Unfortunately, they do not contain information on education levels or other individual characteristics. At the establishment level, the data contain only industry and location. We have access to the IMSS data from 1985 to 2005. Our sample-selection and cleaning procedures for the IMSS data are described in appendix A.1. After cleaning, we observe between 5.26 and 10.15 million individuals per year over the period. Summary statistics for the cleaned IMSS data are reported in appendix table B2.

The plant-level data are from the *Encuesta Industrial Anual (EIA)* [Annual Industrial Survey], conducted by the Instituto Nacional de Estadísticas y Geografía (INEGI), the Mexican statistical agency. For reasons discussed in appendix A.2, we focus on a balanced panel for the 1993–2003 period, which we refer to as the "EIA panel." After cleaning, there are 3,518 plants with complete EIA information in every year over this period. Summary statistics by export status for the EIA panel for 1993 are in appendix table B3. The differences between exporters and nonexporters are similar to those documented in other data sets: exporters are larger, more capital intensive, and higher-wage than nonexporters, and they make up a minority of plants in each industry.

The EIA data have been linked to the IMSS employer-employee data using establishment name, location (municipality and state), and street address. Although in principle all plants appearing in the EIA should also appear in the IMSS data, it has only been possible to link approximately 2,800 of the 3,518 plants from the EIA panel to the IMSS data in each year.

As noted above, we estimate wage premiums using the largest connected set in each period. Table 1 reports summary statistics from the IMSS data for the largest connected

[21] A further possible concern is that plants (within industries) differ in the extent to which they import inputs and that they were differentially affected by the devaluation for this reason. But the import share of inputs is positively correlated with plant size and export status, so this effect, if present, would tend to suggest that higher-export-propensity plants were more negatively affected by the devaluation than lower-export-propensity plants were, making it more difficult for us to find the positive (relative) wage effects that we find. Another possible concern is that we do not allow the return to ability to change within periods, and the peso devaluation might have affected the general-equilibrium return to skill. But note that we estimate the model, equation (1), separately by periods before and after the devaluation, allowing the individual effect to differ across periods. A one-time effect of the devaluation on the return to ability will be captured by the period-specific individual effects. Also, a previous version of this paper (Frías et al., 2009) allowed for year-by-year changes in the return to ability (see note 9) and found broadly similar results.

[22] See Schneider and Enste (2000). The employment figures for manufacturing in the IMSS data in 1993 differ by less than 10% from independently reported figures in the 1993 Industrial Census, suggesting that underreporting of employment is not especially severe in the social security data (Kaplan et al., 2004, 2005). Samaniego de la Parra and Fernández Bujanda (forthcoming) takes advantage of newly available information on the employers of respondents to a household survey and documents substantial informal employment in formal establishments but relatively modest rates (less than 5%) in larger plants (100 or more employees) in manufacturing, which are our focus here.

sets in each period—period 1 (1992–1994), period 2 (1996–1998), and period 3 (2000–2002), as well as period 0 (1988–1990), which will be used in our robustness check below. Comparing to appendix table B2, we see that the largest connected sets typically include roughly half the number of establishments in the full IMSS data set in each year, a smaller share than has typically been the case in developed countries. This is in part because of the high rate of informality in Mexico, in part because we focus on three-year periods (in contrast to, for instance, AKM and Abowd et al., 2002, who focus on a single twelve-year period, and CHK, who focus on seven-year periods), and in part because the establishment size distribution in Mexico is skewed to the left relative to, for instance, the United States (Hsieh & Klenow, 2014). But the limited coverage of the largest connected sets does not appear to be a severe problem for our purposes, for two reasons. First, although they cover a minority of establishments, the largest connected sets cover a large majority of workers—approximately 90%. Second, and perhaps more important, the largest connected sets include almost all of the larger plants in manufacturing that appear in the EIA panel, for which we observe exporting behavior. Appendix table B4 reports the number of EIA panel plants that can be linked to the IMSS data, by connected set status, for each of our three main periods. Of the EIA plants that can be linked to the IMSS data, 98% or more are in the largest connected set in each year.

Once the EIA panel has been linked to the IMSS data, we impose the requirement that plants be in the largest connected set in periods 1 to 3. The resulting panel contains 2,621 plants. We refer to this sample as the EIA-IMSS panel. Table 2 reports summary statistics for the EIA-IMSS panel for 1993. Comparing to appendix table B3, this panel is on average quite similar to the EIA panel. As mentioned above, once the plant and person components have been calculated in the IMSS data, we collapse the EIA-IMSS panel to the plant-period level, averaging across years within plant period.[23]

An important caveat to our study is that our estimates of wage premiums are valid only conditional on individuals being formally employed in formal establishments, and our estimates of the effect of exporting on those premiums are valid only conditional on plants appearing in the EIA-IMSS panel, in which large firms are overrepresented. As suggested above, many individuals in Mexico are not formally employed, even if they work in formal establishments, and many establishments are either not formally registered, are outside of manufacturing, or are too small to appear in the EIA. In these senses, our results should not be considered representative of the Mexican economy as a whole. At the same time, our data are arguably representative of larger manufacturing plants, the subpopulation of plants for which exporting is a realistic possibility.

TABLE 2.—SUMMARY STATISTICS, EIA-IMSS PLANT PANEL, 1993

	(1) nonexporters	(2) exporters	(3) all plants
Sales	139.37	281.03	181.69
	(6.53)	(16.71)	(6.89)
Employment (workers)	188.55	333.29	231.79
	(5.56)	(13.42)	(5.74)
Employment (hours)	451.84	782.59	550.65
	(13.46)	(31.87)	(13.72)
White-collar share	0.30	0.33	0.31
	(0.00)	(0.01)	(0.00)
K/L	139.96	177.47	151.17
	(5.35)	(9.20)	(4.66)
Hourly wage	42.79	58.12	47.37
	(0.55)	(1.15)	(0.53)
Foreign ownership indicator	0.08	0.29	0.14
	(0.01)	(0.02)	(0.01)
Export share of sales		0.15	0.04
		(0.01)	(0.00)
Import share of material purchases	0.14	0.30	0.19
	(0.01)	(0.01)	(0.01)
N	1838	783	2621

The table reports statistics using 1993 data for the EIA-IMSS (balanced) panel of plants. Standard errors of means in parentheses. Exporter defined as export sales above 0. Employment and wage information is from EIA data. Hours are in thousands per year. White-collar share is hours of employment for white-collar workers (*empleados*, nonproduction) over hours for blue-collar workers (*obreros*, production). Foreign ownership indicator takes the value 1 if foreign ownership share is more than 10%, 0 otherwise. Export share is a fraction of total sales derived from exports. Sales are measured in millions of 2002 Mexican pesos, capital-labor ratio in thousands of 2002 pesos, and average daily wage in 2002 pesos. Average 2002 exchange rate: 9.60 pesos/US$1. For further details, refer to section III and appendix A.2.

IV. Estimating Wage Premiums

As described in section II, the first step of our approach is to estimate the AKM-type model, equation (1), in the IMSS individual-level data, separately for periods 1, 2, and 3. Before presenting the estimates, we provide evidence for the validity of the conditional random mobility assumption, equation (2). Following CHK, we show simple event-study plots of average wage changes between firms in different quartiles of the distribution of wages paid to coworkers. For these plots, to ensure that we can observe two years before and two years after a job transition (again following CHK), we focus on four-year periods, 1992 to 1995 and 2000 to 2003.[24] Figures 3 and 4 show the mean real wages of movers; to reduce clutter, we focus on workers leaving firms in quartiles 1 and 4. As in CHK, two features stand out. First, the wage trends before and after job switches are parallel across the different types of transitions; there are no Ashenfelter (1978)-type dips or rises before transitions. If the individual time-varying productivity shocks, ε_{it}, were determining job transitions, we would expect positive wage changes prior to movement to a higher-quartile plant and negative changes prior to movement to a lower-quartile plant. Second,

[23]For the plant and person components, we average over 1992 to 1994, 1996 to 1998, and 2000 to 2002; for the EIA variables, we average over 1993 to 1994, 1996 to 1998, and 2000 to 2002, since the EIA data we use begin in only 1993. In the robustness check using IMSS data to extend back to 1988 to 1990, we include the plants in the EIA-IMSS panel for which IMSS information is also available in 1988 to 1990. There are 2,314 such plants.

[24]The requirement that we observe two years before and two years after means that we must focus on transitions between the middle two years of each period, 1993 to 1994 or 2001 to 2002.

FIGURE 3.—MOVERS' MEAN REAL WAGES, IMSS DATA, 1992–1995

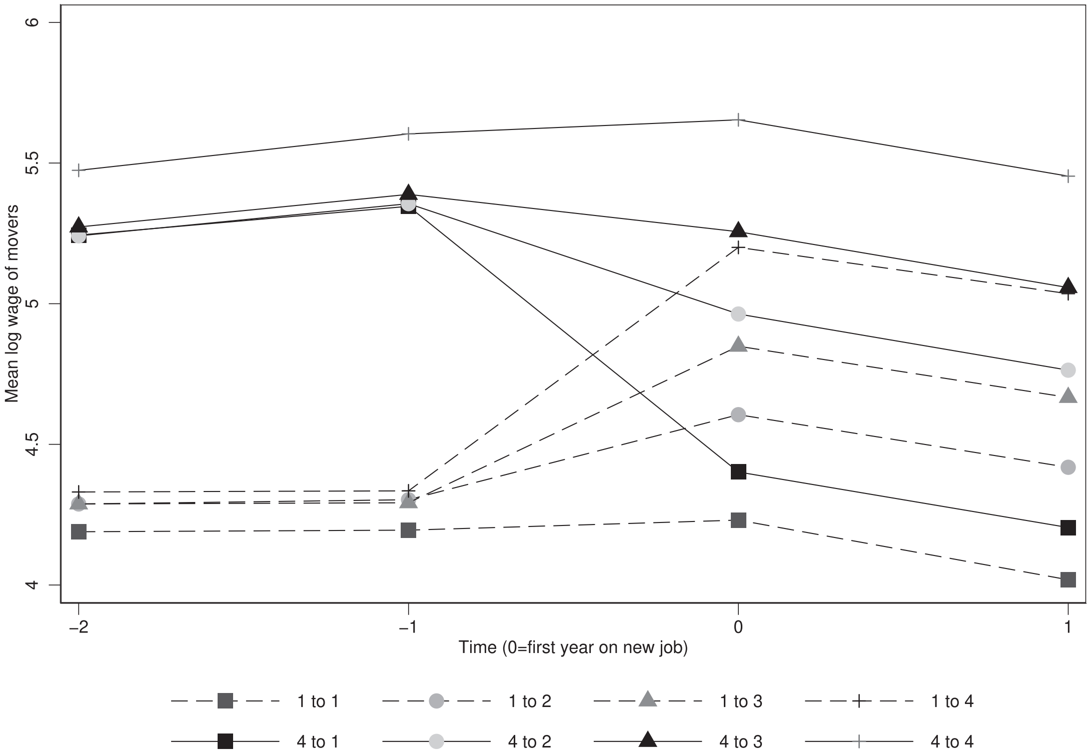

Sample is all workers observed over 1992 to 1995 in the IMSS database (after cleaning steps 1 to 6 described in appendix A.1) who changed job between 1993 and 1994 and held both the preceding and new job for at least two years. Each line corresponds to a transition between types of firms classified by quartiles of the average coworkers' wage.

FIGURE 4.—MOVERS' MEAN REAL WAGES, IMSS DATA, 2000–2003

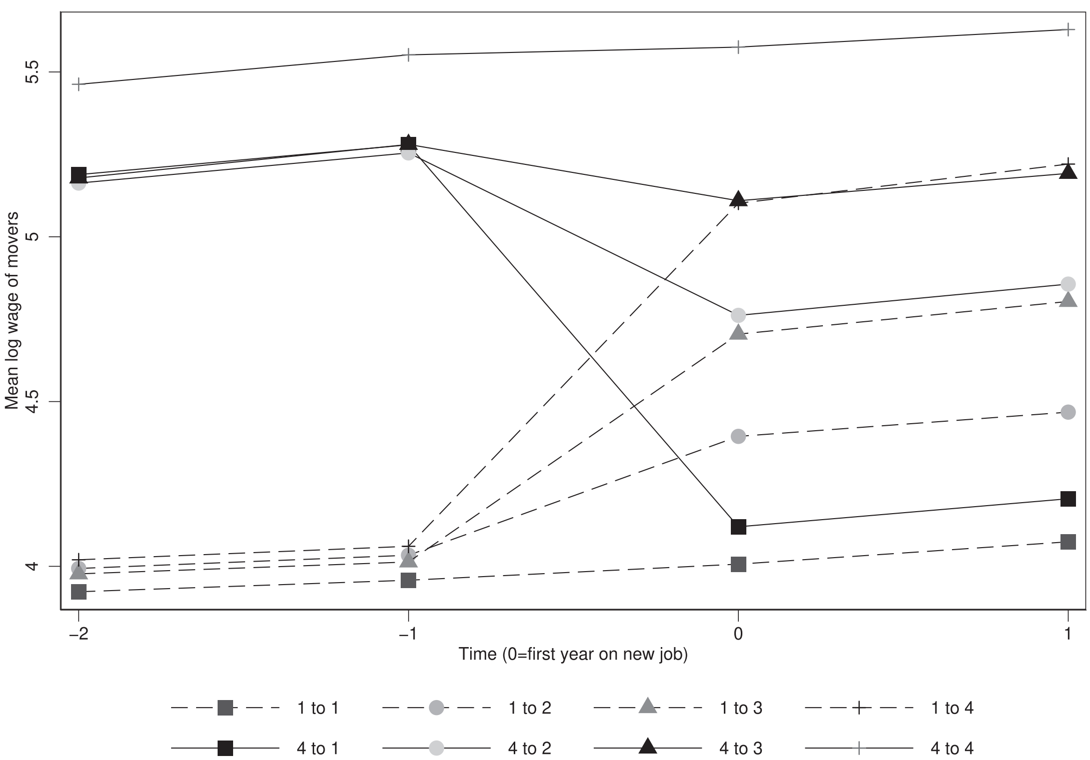

Sample is all workers observed over 2000 to 2003 in the IMSS database (after cleaning steps 1–6 described in appendix A.1) who changed job between 2000 and 2001 and held both the preceding and new job for at least two years. Each line corresponds to a transition between types of firms classified by quartiles of the average coworker's wage.

TABLE 3.—ESTIMATION RESULTS FOR AKM-TYPE MODEL, PER PERIOD

	(1) Period 1 1992–1994	(2) Period 2 1996–1998	(3) Period 3 2000–2002
Number of individuals	10,121,284	11,155,022	13,137,161
Number of plants	394,672	402,213	479,088
Summary of parameter estimates			
Std. dev. individual effect	0.451	0.441	0.445
Std. dev. plant effect	0.362	0.390	0.423
Std. dev. Xb	0.144	0.196	0.164
Corr. individual/plant effects	0.007	0.060	0.054
Corr. individual effect/Xb	0.145	0.019	0.068
Corr. plant effect/Xb	0.116	0.112	0.092
RMSE of AKM residual	0.184	0.189	0.189
Adjusted R-squared	0.916	0.919	0.923
Additional statistics			
Std. dev. of log wages	0.637	0.665	0.682
N	21,045,892	23,283,074	27,960,760

The table shows statistics from the estimation of equation (1) for periods 1992–1994, 1996–1998, and 2000–2002 separately. Sample is described in notes to table 1. Individual effects are estimates of α_i, plant effects are estimates of $\psi_{J(i,t)}$, and Xb are estimates of $X_{it}'\beta$ in equation (1). Covariates included in X_{it} are age squared, age cubed (both recentered at 40), tenure, tenure squared, and year effects.

the wage change for workers moving from one quartile to another has approximately the same magnitude as (and opposite sign of) the change for workers moving in the opposite direction.[25] If the ε_{it} contained an employer-employee match-specific effect that also affected workers' mobility, we would expect the gain for workers moving to a higher quartile to be larger than the loss for those making the opposite move.[26] Both features suggest that the data are well described by a model with a fixed individual component, a fixed plant component, time-varying observables, and a random shock uncorrelated with mobility.

We now turn to estimation of the AKM-type model, equation (1). We estimate plant and person components for all individuals in the largest connected sets described by table 1, separately by period, for periods 1, 2, and 3. Table 3 reports key statistics. There are several points to notice. First, the model has a high in-sample predictive power, with adjusted R^2 above 90%, as in other studies in the literature. Second, we see that the standard deviations of the individual effects remain relatively constant, but the standard deviations of the plant components increase over time, foreshadowing our results below that the plant-level response to the export shock is explained primarily by changes in wage premiums.[27] Third, there is little evidence of negative correlation between estimated person and plant effects.

One might be concerned that the use of three-year periods and the high rate of informality in our context might limit the number of observed plant-to-plant switches and exacerbate what Abowd et al. (2004) call "limited mobility bias." Small numbers of switchers can generate a negative bias in the correlation between individual and plant effects (Andrews et al., 2008) which can be substantial (Maré & Hyslop, 2006; Andrews et al., 2012). One way to address this concern is to conduct the "leave-out" estimation of variance components recently proposed by Kline et al. (2020). This entails estimating our model in each period in the leave-one-out connected set, that is, the set of plants that remain connected once any single worker has been removed from the sample. Appendix table B5 reports the results of the variance decomposition, for both the standard AKM estimates and the Kline et al. (2020) leave-out estimates, for this restricted sample. As in Kline et al. (2020), the leave-one-out connected sets are substantially smaller than the largest connected sets used above. But reassuringly, the variance estimates are quite similar using the two methods, with the variances due to individual and plant effects just slightly smaller than the standard AKM estimates, suggesting that limited mobility bias is not severe in our context. It is also worth noting that if the limited mobility bias is constant across periods, it will be differenced out when we look at within-firm changes in the plant and average person components. Given the similarity of the Kline et al. (2020) and AKM results for the leave-one-out connected sets, our preferred approach is to continue using the standard AKM estimates, which allow us to maintain a larger sample size in the second-step estimation of the effect of exporting.

V. Estimating the Effect of Exporting on Wage Premiums

This section estimates the relationship between exporting and the estimated plant and average person components using the IV specification in equation (4). As a preliminary step, we construct the measures of export propensity discussed in section II. Given that we are allowing the coefficients on the covariates in Z_{jp} in equation (7) to vary by four-digit sector, of which there are 50, it is challenging to communicate the estimates in a concise way. Rather than report all parameters (1,596 total in the full model, 550 if we do not report industry-year or state-year effects), appendix table B6 simply averages the coefficients and standard errors across the four-digit industry interactions. Appendix table B7 shows that, as expected, the measures of export propensity constructed in the preliminary stage are highly correlated in cross-section with plant size (measured by log employment) as well as the main outcomes we consider: capital intensity,

[25] Appendix figures B2 and B3 make the symmetry clear visually. Appendix figures B4 and B5 are similar to figures 3 and 4, but use nominal wages instead of real wages. The patterns are similar.

[26] As Card et al. (2018) point out, models with a worker-firm match component tend to predict that workers who move will see wage increases, irrespective of whether they move to higher or lower quartiles.

[27] The results in table 3 indicate that the plant effects explain about 32% of the overall wage variance in period 1, and the share rises over time. This share is larger than has been observed in developed-country contexts (e.g., in Germany; Card et al., 2013) or Italy (Kline et al., 2020), which both find that plant or firm effects account for approximately 20% of the overall wage variance. But Kline et al. (2020) find that for younger workers, firm effects make up about 31% of overall wage variance. We believe that it is plausible that the Mexican labor market, where it is arguably more difficult for firms to observe employment track records than in developed-country settings, is similar to labor markets for younger workers in developed countries on this dimension. Also, Engbom and Moser (2022) find that plant effects explain between 26% and 30% of the overall wage variance in Brazil for 1994 to 1998, shares that are similar to what we find.

TABLE 4.—BASELINE ESTIMATES, EFFECT OF EXPORTING

	(1) △ export share	(2) △ log K/L	(3) △ log avg. hourly wage (EIA)	(4) △ avg. log daily wage (IMSS)	(5) △ plant component	(6) △ avg. person component
A. OLS						
△ export share		0.010	0.017	0.014	0.041	−0.028
		(0.083)	(0.038)	(0.023)	(0.026)	(0.024)
init. pred. export status		0.198***	0.145***	0.112***	0.095***	0.019
		(0.043)	(0.020)	(0.013)	(0.017)	(0.015)
B. First stage and reduced form						
init. pred. export status × devaluation	0.123***	0.061	0.125***	0.137***	0.133***	0.007
	(0.017)	(0.086)	(0.041)	(0.026)	(0.033)	(0.030)
init. pred. export status	−0.005	0.169***	0.086***	0.047***	0.033	0.015
	(0.010)	(0.058)	(0.025)	(0.016)	(0.024)	(0.021)
F-stat	55.549					
C. IV						
△ export share		0.495	1.015***	1.111***	1.080***	0.053
		(0.667)	(0.347)	(0.249)	(0.290)	(0.233)
init. pred. export status		0.172***	0.091***	0.052***	0.038	0.015
		(0.053)	(0.025)	(0.018)	(0.024)	(0.019)
6-digit industry × period effects	Y	Y	Y	Y	Y	Y
region (state) × period effects	Y	Y	Y	Y	Y	Y
N (plants)	2621	2621	2621	2621	2621	2621
N (obs)	5242	5242	5242	5242	5242	5242

The dependent variables (at top) are changes between periods ($p-1$ and p). The excluded instrument is the interaction between predicted export status (predicted using equation (7), reported in appendix table B6) in period $p-1$ and an indicator for the devaluation that equals 1 for period 2, and 0 otherwise. Panel A reports OLS regressions, panel B reports the first stage in column 1 and reduced-form results in columns (2)–(6), and panel C reports the corresponding IV regressions. Sample is the EIA-IMSS panel. Export share is fraction of total sales derived from exports. Robust standard errors in parentheses. *10% level, **5% level, and ***1% level.

plant-average wages from the EIA and IMSS data, and the plant and average person components.

Table 4 presents our baseline estimates of equation (4), using our preferred specification for predicted export status (from column 1 of appendix table B6) as the proxy, $\widehat{\lambda}$. Panel A presents OLS estimates; the correlation of export increases and wages changes is close to 0 and not statistically significant. Column 1 of panel B reports the first stage: the relationship between the excluded instrument (initial predicted export status, $\widehat{\lambda}_{jp-1}$, interacted with a period 2 indicator, P_2) and the change in export share is strong, with an F-statistic of 55.5. Columns 2 to 6 of panel B present the reduced form corresponding to our IV specification, and we see a strong relationship between the excluded instrument and changes in plant-level average wages and the plant components, but not in capital intensity or average person components. Panel C presents the IV estimates, which display an economically and statistically significant positive effect of changes in export share on changes in plant-average wages and plant components. A 1% increase in export share is associated with approximately a 1.1% increase in plant-average wages (in either the EIA or IMSS data) and in plant components. Strikingly, there is a near-zero effect of exporting on average person components. Essentially all of the increase in plant-average wages is explained by increases in wage premiums, not by changes in the composition of the workforce. The estimates in columns 3 to 5 of panel C are notably larger than the corresponding OLS estimates in panel A. This disparity is consistent with the idea that there is an omitted-variable bias in OLS due to labor-supply shocks: if the labor supply to a plant increased idiosyncratically for exogenous reasons, we would expect the plant to pay lower wages and to increase exports. (Because its costs would be lower, it would be more competitive on the export market.) This would generate a negative bias in the OLS estimates. Other exogenous shocks to labor costs would have similar effects. The panel C estimates also indicate stable differential trends in capital intensity and plant-average wages between plants with different initial export propensities, as captured by the uninteracted initial predicted export status term. The coefficients for changes in plant components and average person components are also positive but not significant; the point estimates suggest that the lion's share of the differential trend in plant-average wages is accounted for by trends in plant components. But the key point is that even controlling for such differential trends, the devaluation-induced increase in exporting between period 1 and period 2 had a significant effect on plant-average wages and wage premiums.

To explore robustness to the method of constructing export propensity, table 5 presents similar specifications using three additional measures. To save space, we focus on the IMSS wage outcomes and do not report OLS specifications.[28] In panel A, which uses predicted export share in place of predicted export status, the first stage is a bit weaker than in our baseline estimates, with an F-statistic of 23.62, but the IV estimates are statistically indistinguishable from those in table 4. The same is true for the results in panels B and C, which drop either employment (and variables that depend on it) or sales (and variables that depend on it) from the preliminary-stage estimation of export propensity.

[28]The results for the other outcomes and OLS are similar to those in table 4.

TABLE 5.—EFFECT OF EXPORTING, USING ALTERNATIVE MEASURES OF EXPORT PROPENSITY

	(1) △ export share	(2) △ avg. log daily wage	(3) △ plant comp.	(4) △ avg. person comp.
A. Predicted export share as proxy				
1. First stage and reduced form				
init. pred. export share × devaluation	0.279***	0.307***	0.304***	0.011
	(0.057)	(0.075)	(0.081)	(0.077)
init. pred. export share	0.026	0.090**	0.050	0.042
	(0.033)	(0.039)	(0.056)	(0.051)
F-stat	23.619			
2. IV				
△ export share		1.099***	1.088***	0.040
		(0.332)	(0.345)	(0.265)
init. pred. export share		0.062	0.021	0.041
		(0.057)	(0.069)	(0.054)
B. Predicted export status as proxy, employment omitted from preliminary stage				
1. First stage and reduced form				
init. pred. export status, no emp. × devaluation	0.127***	0.163***	0.161***	0.004
	(0.018)	(0.028)	(0.036)	(0.034)
init. pred. export status, no emp.	−0.010	0.051***	0.046*	0.005
	(0.011)	(0.016)	(0.026)	(0.023)
F-stat	48.383			
2. IV				
△ export share		1.285***	1.267***	0.034
		(0.280)	(0.322)	(0.254)
init. pred. export status, no emp.		0.063***	0.058**	0.005
		(0.019)	(0.025)	(0.020)
C. Predicted export status as proxy, sales omitted from preliminary stage				
1. First stage and reduced form				
init. pred. export status, no sales × devaluation	0.116***	0.134***	0.136***	0.001
	(0.017)	(0.026)	(0.034)	(0.030)
init. pred. export status, no sales	−0.002	0.041***	0.029	0.013
	(0.011)	(0.016)	(0.024)	(0.021)
F-stat	46.052			
2. IV				
△ export share		1.153***	1.168***	0.009
		(0.272)	(0.322)	(0.250)
init. pred. export status, no sales		0.044**	0.031	0.013
		(0.019)	(0.025)	(0.020)
6-digit industry × period effects	Y	Y	Y	Y
region (state) × period effects	Y	Y	Y	Y
N (plants)	2,621	2,621	2,621	2,621
N (obs)	5,242	5,242	5,242	5,242

Specifications similar to panels B–C, columns 1, 4–6, of table 4. See sections II and V for details of export propensity measures. Robust standard errors in parentheses. *10% level, **5% level, and ***1% level.

Overall, the empirical patterns appear not to depend on the details of the construction of export propensity.

For purposes of comparison to previous work, we also report results using plant size and TFP as proxies for plant capability in appendix B. In part because we effectively constrain the relationship between the proxies and export propensity to be the same across sectors, the first stage is weaker in these specifications, but the results are qualitatively consistent with the results using predicted export status or share.

As an additional robustness check, we compare changes between periods 1 (1992–1994) and 2 (1996–1998) to changes between similarly spaced periods before the devaluation. As discussed in appendix A.2, we are not able to use EIA plant-level data prior to 1993 and hence do not observe the change in export share from period 0 to period 1, which prevents us from estimating our IV specification. However, we are able to estimate a corresponding reduced form, as follows. We define period 0 to be 1988 to 1990 and estimate the AKM-type model, equation (1), for the largest connected set for this period. We then select establishments in the EIA-IMSS panel for which AKM estimates are also available in period 0, yielding a balanced panel of 2,316 plants observed over the four periods, 0 to 3. Using data from periods 1 to 3, we estimate a preliminary-stage linear probability model, equation (7), using only a third-degree polynomial in log employment in the vector of covariates, Z_{jp}.[29] Using the coefficients from this regression and IMSS information on employment, we generate predicted export status for all four periods, 0 to 3. Panel A of table 6 reports cross-sectional correlations of this measure of export propensity and the wage outcomes; the results are broadly similar to those in the EIA-IMSS panel (see appendix table B7). Panel B of table 6 reports regressions period-by-period, of changes in the dependent variables on the initial level of export propensity. As

[29] In the IMSS data, employment is a head count measure; hours are not observed.

TABLE 6.—DIFFERENTIAL EFFECT OF DEVALUATION, IMSS 1988–2002 PANEL

	(1) avg. log daily wage (IMSS)	(2) plant comp.	(3) person comp.
A. Cross-sectional correlations, period 1			
pred. export status, period 1	0.371***	0.319***	0.052*
	(0.051)	(0.038)	(0.029)
	△ avg. log daily wage (IMSS)	△ plant comp.	△ person comp.
B.1 Outcomes: changes from period 0 to 1			
pred. export status, period 0 (π_1)	0.043*	0.038	0.006
	(0.025)	(0.026)	(0.025)
B.2 Outcomes: changes from period 1 to 2			
pred. export status, period 1 (π_2)	0.175***	0.152***	0.022
	(0.027)	(0.034)	(0.031)
B.3. Outcomes: changes from period 2 to 3			
pred. export status, period 2 (π_3)	0.059***	0.000	0.060**
	(0.022)	(0.035)	(0.029)
C. Differences in coefficients			
$\pi_2 - \pi_1$	0.131***	0.114***	0.017
	(0.036)	(0.043)	(0.040)
$\pi_3 - \pi_2$	−0.115***	−0.151***	0.038
	(0.035)	(0.048)	(0.043)
6-digit industry effects	Y	Y	Y
region (state) effects	Y	Y	Y
N	2,314	2,314	2,314

Sample is plants in EIA-IMSS panel that are also in largest connected set in period 0 (1988–1990). Variables are from IMSS data. Panel A reports three separate cross-sectional regressions of the dependent variable at top on average log employment in period 1, analogous to table B.7. Panels B1 to B3 report three separate regressions each, of the dependent variable at the top (i.e., change for indicated period) on average log employment in the previous period. See discussion in section V for details. All regressions include industry and state effects and have sample size indicated at bottom. Panel C reports differences in coefficients from panel B and the standard errors on the differences. Robust standard errors in parentheses. *10% level, **5% level, and ***1% level.

discussed in note 18, if only periods 1 to 3 were included, the difference in coefficients between panels B.2 and B.3 would be precisely the estimate of the coefficient φ in the reduced-form model, equation (5).[30] The advantage of viewing the results period-by-period in this way is that we can compare the coefficient in panel B.1 (for the change from period 0 to period 1) to the coefficient in panel B.2 (for the change from period 1 to period 2), separately from the comparison between the coefficients in panels B.2 (π_2) and B.3 (π_3). Panel C reports these differences in coefficients and their standard errors. The coefficients for (changes in) both plant-average log wages and the plant component are significantly larger for the periods spanning the peso crisis (periods 1 and 2, π_2) than for either the two earlier periods (periods 0 and 1, coefficient π_1) or the two later periods (periods 2 and 3, π_3). The magnitudes of the differences in coefficients are similar when making either the earlier or later placebo comparison (with sign flipped, as expected given that we preserve chronological order when taking the differences). The magnitudes are also similar to the reduced-form estimates in our baseline specification, in panel B of table 4. Overall, the table is consistent with our previous findings.

As a further robustness check, we consider the effect of the export shock on the wages for stayers, employees continuously employed in a given plant for consecutive periods. Table 7 reports first-stage, reduced-form, and IV results similar to columns 1 and 4 of table 4, but where the change in the average log daily wage of stayers is the outcome of interest. As mentioned in section II, this is analogous to estimating a model with job-spell fixed effects, as in several existing papers.[31] The estimates are identified by variation in wages within firm-worker matches in response to changes in export share at the plant level. Despite the very different approach, the reduced-form and IV coefficients in column 2

TABLE 7.—DIFFERENTIAL EFFECT OF DEVALUATION, IV ESTIMATES, STAYERS ONLY

	(1) △ export share	(2) △ avg. log daily wage, stayers
A. First stage and reduced form		
init. pred. export status × devaluation	0.124***	0.115***
	(0.017)	(0.022)
init. pred. export status	−0.005	0.085***
	(−0.005)	(0.085)
F-stat	56.293	
B. IV		
change in export share		0.927***
		(0.206)
init. pred. export status		0.090***
		(0.016)
6-digit industry × period effects	Y	Y
region (state) × Period effects	Y	Y
N	2,619	2,619

In column 2, the plant-level average wage has been computed using only workers who did not change jobs between periods (stayers). Sample is plants in EIA-IMSS panel for which a change in average log wage for stayers could be constructed both between periods 1 and 2 and between periods 2 and 3. (Of the 2,621 EIA-IMSS panel plants, two did not have stayers between one pair of these periods.) Export share is fraction of total sales derived from exports. Robust standard errors in parentheses. *10% level, **5% level, and ***1% level.

[30]The coefficients in panels B.1 to B.3 correspond to estimates of π_p in equation (6).

[31]This specification is not exactly equivalent to an individual-year-level specification with job-spell effects, because we have collapsed the data at the period level and define as a stayer any worker who appears in a given plant in at least one year of two consecutive periods.

are very similar to the baseline estimates for changes in wage premiums in table 4, column 5. It is reassuring that the estimates for stayers are similar to those using the AKM methodology, which is based on switchers between plants.

As a final check, we present estimates where we include indicators for quartile of the predicted export status distribution in place of the linear term (appendix table B11). Although this coarsening of the export propensity proxy weakens the first stage somewhat, the estimates are consistent with our approach above. The relationship between quartile of export propensity interacted with the devaluation indicator and changes in export share (column 1), plant average wages from the employer-employee data (column 4), and the plant component (column 5) are monotonic and seem plausibly approximated by a linear relationship.

VI. Conclusion

This paper has investigated the role of exporting in shaping Mexican plants' wage policies, in particular the payment of wage premiums, defined as wages above what individuals would earn elsewhere on the labor market. We have estimated simple AKM-type models separately by three-year periods, to decompose plant-average wages into a plant component (interpreted as an average wage premium) and an average person component. We have argued that the peso devaluation of late 1994 generated a greater export inducement for plants with higher initial export propensity and have found that increased exporting led plants to increase wage premiums. The increase in the plant component can explain essentially all of the differential increase in plant-average wages arising from the export shock over the period we study. The findings highlight the causal importance of trade shocks, and product-market shocks more generally, in determining wage premiums at the firm level.

Three questions seem particularly worthy of further research. First, how persistent are the effects on wage premiums? In theoretical models in which an initial rise in wage premiums with increased exporting is due to search frictions, one would expect wage premiums to decline gradually over time. We have found that the effects persist over six to eight years, but it remains an open question whether firms gradually reduce premiums over a longer period. The fact that we find differences of similar magnitude when comparing to a predevaluation "placebo" as to a postdevaluation one suggests that the effects were persistent. But answering this question definitively will require a longer data series following a trade shock.

Second, to what extent did the differential increase in wage premiums in response to the peso devaluation contribute to aggregate wage inequality in Mexico? Our paper suggests a mixed response to this question. On one hand, the export shock led plants with greater initial export propensity, which were already larger and paid higher wages and wage premiums, to increase wages and wage premiums further. In this sense, the effect of exports on wage premiums tended to increase within-industry inequality. On the other hand, the fact that we find limited effects on skill composition suggests that the effect of the export shock on the general-equilibrium return to skill may have been limited. A fuller accounting of the role of firms' wage policies in overall inequality in Mexico is a task for future work.

Third, to what extent do the income levels (or other characteristics) of destination markets matter for how export demand shocks get transmitted into wage premiums? There is increasing evidence that destination income, as opposed to the volume of exports per se, matters for plant-average wages (Brambilla et al., 2012) and other input prices (Bastos, Silva, & Verhoogen, 2018). But since more than 80% of Mexican exports go to the United States, we do not have the variation to investigate the effect of destination income on wage premiums in our setting. This question awaits an application in a setting with greater diversity of export destinations.

REFERENCES

Abowd, John, Robert Creecy, and Francis Kramarz, "Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data," LEHD technical paper 2002-06 (2002).
Abowd, John, Francis Kramarz, Paul Lengermann, and Sébastien Pérez-Duarte, "Are Good Workers Employed by Good Firms? A Test of a Simple Assortative Matching Model for France and the United States," unpublished paper, Cornell University (January 2004).
Abowd, John, Francis Kramarz, and David Margolis, "High Wage Workers and High Wage Firms," *Econometrica* 67:2 (1999), 251–333. 10.1111/1468-0262.00020
Abowd, John, and Thomas Lemieux, "The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada," *Quarterly Journal of Economics* 108:4 (1999), 983–1014. 10.2307/2118457
Abowd, John M., Kevin L. McKinney, and Nellie L. Zhao, "Earnings Inequality and Mobility Trends in the United States: Nationally Representative Estimates from Longitudinally Linked Employer-Employee Data," *Journal of Labor Economics* 36:S1 (2018), S183–S300.
Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen, "On the Returns to Invention within Firms: Evidence from Finland," *American Economic Review Papers and Proceedings* 108:2 (2018), 208–212. 10.1257/pandp.20181108
Akerlof, George, and Janet Yellen, "The Fair Wage–Effort Hypothesis and Unemployment," *Quarterly Journal of Economics* 105:2 (1990), 255–283. 10.2307/2937787
Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser, "Firms and the Decline in Earnings Inequality in Brazil," *American Economic Journal: Macroeconomics* 10:1 (2018), 149–189. 10.1257/mac.20150355
Álvarez, Roberto, and Ricardo A. López, "Skill Upgrading and the Real Exchange Rate," *World Economy* 32:8 (2009), 1165–1179.
Amiti, Mary, and Donald R. Davis, "Trade, Firms, and Wages: Theory and Evidence," *Review of Economic Studies* 79:1 (2012), 1–36. 10.1093/restud/rdr016
Andrews, Martyn J., Leonard Gill, Thorsten Schank, and Richard Upward, "High Wage Workers and Low Wage Firms: Negative Assortative Matching or Limited Mobility Bias?" *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 171:3 (2008), 673–697. 10.1111/j.1467-985X.2007.00533.x
——— "High Wage Workers Match with High Wage Firms: Clear Evidence of the Effects of Limited Mobility Bias," *Economics Letters* 117:3 (2012), 824–827. 10.1016/j.econlet.2012.07.033
Araújo, Bruno César, and Lourenço S. Paz, "The Effects of Exporting on Wages: An Evaluation Using the 1999 Brazilian Exchange Rate Devaluation," *Journal of Development Economics* 111 (2014), 1–16.
Ashenfelter, Orley, "Estimating the Effect of Training Programs on Earnings," this REVIEW 60:1 (1978), 47–57.

Barth, Erling, Alex Bryson, James C. Davis, and Richard B. Freeman, "It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States," *Journal of Labor Economics* 34:S2 (2016), S67–S97.

Barth, E., James C. Davis, and Richard B. Freeman, "Augmenting the Human Capital Earnings Equation with Measures of Where People Work," *Journal of Labor Economics* 36:S1 (2018), S71–S97.

Bastos, Paulo, Joana Silva, and Eric Verhoogen, "Export Destinations and Input Prices: Evidence from Portugal," *American Economic Review* 108:2 (2018), 353–392. 10.1257/aer.20140647

Baumgarten, Daniel, "Exporters and the Rise in Wage Inequality: Evidence from German Linked Employer-Employee Data," *Journal of International Economics* 90:1 (2013), 201–217. 10.1016/j.jinteco.2012.10.001

Bellon, Matthieu, "Trade Liberalization and Inequality: A Dynamic Model with Firm and Worker Heterogeneity," unpublished paper, Columbia University (December 2017).

Bernard, Andrew B., and J. Bradford Jensen, "Exporters, Jobs, and Wages in U.S. Manufacturing: 1976–1987," *Brookings Papers on Economic Activity: Microeconomics* (1995), 67–112.

——— "Exceptional Exporter Performance: Cause, Effect, or Both?" *Journal of International Economics* 47 (February 1999), 1–25. 10.1016/S0022-1996(98)00027-0

Bertrand, Marianne, and Sendhil Mullainathan, "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences," *Journal of Political Economy* 111:5 (2003), 1043–1075. 10.1086/376950

Brambilla, Irene, Nicolas Depetris Chauvin, and Guido Porto, "Examining the Export Wage Premium in Developing Countries," *Review of International Economics* 25:3 (2017), 447–475. 10.1111/roie.12231

Brambilla, Irene, Daniel Lederman, and Guido Porto, "Exports, Export Destinations and Skills," *American Economic Review* 102:7 (2012), 3406–3488. 10.1257/aer.102.7.3406

Bustos, Paula, "The Impact of Trade Liberalization on Skill Upgrading: Evidence from Argentina," unpublished paper, Universitat Pompeu Fabra (July 2011a).

——— "Trade Liberalization, Exports and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *American Economic Review* 101:1 (2011b), 304–340. 10.1257/aer.101.1.304

Card, David, Ana Rute Cardoso, and Patrick Kline, "Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women," *Quarterly Journal of Economics* 131:2 (2015), 633–686. 10.1093/qje/qjv038

Card, David, Jörg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics* 36:S1 (2018), S13–S70.

——— "Workplace Heterogeneity and the Rise of West German Wage Inequality," *Quarterly Journal of Economics* 128:3 (2013), 967–1015. 10.1093/qje/qjt006

Coşar, A. Kerem, Nezih Guner, and James Tybout, "Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy," *American Economic Review* 106:3 (2016), 625–663.

Das, Sanghamitra, Mark Roberts, and James Tybout, "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica* 73:3 (2007), 837–873. 10.1111/j.1468-0262.2007.00769.x

Davidson, Carl, Fredrik Heyman, Steven Matusz, Fredrik Sjöholm, and Susan Chun Zhu, "Globalization and Imperfect Labor Market Sorting," *Journal of International Economics* 94:2 (2014), 177–194. 10.1016/j.jinteco.2014.08.001

Davidson, Carl, Steven J. Matusz, and Andrei Shevchenko, "Globalization and Firm Level Adjustment with Imperfect Labor Markets," *Journal of International Economics* 75:2 (2008), 295–309. 10.1016/j.jinteco.2008.02.004

Egger, Hartmut, and Udo Kreickemeier, "Firm Heterogeneity and the Labour Market Effects of Trade Liberalisation," *International Economic Review* 50:1 (2009), 187–216. 10.1111/j.1468-2354.2008.00527.x

——— "Fairness, Trade, and Inequality," *Journal of International Economics* 86:2 (2012), 184–196. 10.1016/j.jinteco.2011.10.002

Engbom, Niklas, and Christian Moser, "Earnings Inequality and the Minimum Wage: Evidence from Brazil," *American Economic Review* 112:12 (2022), 3803–3847.

Fajgelbaum, Pablo D., "Labor Market Frictions, Firm Growth, and International Trade," unpublished paper, UCLA (2016).

Felbermayr, Gabriel, Julien Prat, and Hans-Jörg Schmerer, "Globalization and Labor Market Outcomes: Wage Bargaining, Search Frictions, and Firm Heterogeneity," *Journal of Economic Theory* 146:1 (2011), 39–73. 10.1016/j.jet.2010.07.004

Frías, Judith A., David S. Kaplan, and Eric A. Verhoogen, "Exports and Wage Premia: Evidence from Mexican Employer-Employee Data," unpublished paper, Columbia University (2009).

——— "Exports and Within-Plant Wage Distributions: Evidence from Mexico," *American Economic Review Papers and Proceedings* 102:3 (2012), 435–440. 10.1257/aer.102.3.435

Garin, Andrew, and Filipe Silvério, "How Responsive Are Wages to Firm-Specific Changes in Labor Demand? Evidence from Idiosyncratic Export Demand Shocks," *Review of Economic Studies* (forthcoming).

Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J Redding, "Trade and Inequality: From Theory to Estimation," *Review of Economic Studies* 84:1 (2010), 357–405. 10.1093/restud/rdw025

Helpman, Elhanan, Oleg Itskhoki, and Stephen Redding, "Inequality and Unemployment in a Global Economy," *Econometrica* 78:4 (2017), 1239–1283.

Holtz-Eakin, Douglas, Whitney Newey, and Harvey S. Rosen, "Estimating Vector Autoregressions with Panel Data," *Econometrica* 56:6 (1988), 1371–1395. 10.2307/1913103

Hsieh, Chang-Tai, and Peter J. Klenow, "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics* 129:3 (2014), 1035–1084. 10.1093/qje/qju014

Hummels, David, Rasmus Jørgensen, Jakob R. Munch, and Chong Xiang, "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data," *American Economic Review* 104:6 (2014), 1597–1629. 10.1257/aer.104.6.1597

Irarrazabal, Alfonso, Andreas Moxnes, and Karen Helene Ulltveit-Moe, "Heterogeneous Firms or Heterogeneous Workers? Implications for Exporter Premiums and the Gains from Trade," this REVIEW 95:3 (2013), 839–849.

Kandilov, Ivan, "Do Exporters Pay Higher Wages? Plant-Level Evidence from an Export Refund Policy in Chile," *World Bank Economic Review* 23:2 (2009), 269–294. 10.1093/wber/lhp004

Kaplan, David S., Gabriel Martinez Gonzalez, and Raymond Robertson, "Worker- and Job-Flows in Mexico," unpublished paper, ITAM (2004).

——— "What Happens to Wages after Displacement?" *Economía* 5:2 (2005), 197–242.

Klein, Michael W., Christoph Moser, and Dieter M. Urban, "Exporting, Skills and Wage Inequality," *Labour Economics* 25b (2013), 76–85. 10.1016/j.labeco.2013.03.005

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar, "Who Profits from Patents? Rent-Sharing at Innovative Firms," *Quarterly Journal of Economics* 134:3 (2019), 1343–1404. 10.1093/qje/qjz011

Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten, "Leave-Out Estimation of Variance Components," *Econometrica* 88:5 (2020), 1859–1898. 10.3982/ECTA16410

Krishna, Pravin, Jennifer P. Poole, and Mine Zeynep Senses, "Wage Effects of Trade Reform with Endogenous Worker Mobility," *Journal of International Economics* 93:2 (2014), 239–252. 10.1016/j.jinteco.2014.03.003

Lachowska, Marta, Alexandre Mas, Raffaele Saggio, and Stephen Woodbury, "Do Firm Effects Drift? Evidence from Washington Administrative Data," *Journal of Econometrics* 233:2 (2023), 375–395. 10.1016/j.jeconom.2021.12.014

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury, "Sources of Displaced Workers' Long-Term Earnings Losses," *American Economic Review* 110 (2020), 3231–3266. 10.1257/aer.20180652

Lazear, Edward P., and Kathryn L. Shaw, "Introduction: Firms and the Distribution of Income: The Roles of Productivity and Luck," *Journal of Labor Economics* 36:S1 (2018), S1–S12.

Macis, Mario, and Fabiano Schivardi, "Exports and Wages: Rent Sharing, Workforce Composition, or Returns to Skills?" *Journal of Labor Economics* 34:4 (2016), 945–978. 10.1086/686275

Maré, David C., and Dean R. Hyslop, "Worker-Firm Heterogeneity and Matching: An Analysis Using Worker and Firm Fixed Effects Estimated from LEED," unpublished paper (2006).

Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica* 71:6 (2003), 1695–1725. 10.1111/1468-0262.00467

Munch, Jakob Roland, and Jan Rose Skaksen, "Human Capital and Wages in Exporting Firms," *Journal of International Economics* 75:2 (2008), 363–372. 10.1016/j.jinteco.2008.02.006

Roberts, Mark J., and James R. Tybout, "The Decision to Export in Colombia: An Empirical Model of Entry with Sunk Costs," *American Economic Review* 87:4 (1997), 545–564.

Samaniego de la Parra, Brenda, and León Fernóndez Bujanda, "Increasing the Cost of Informal Workers: Evidence from Mexico," *American Economic Journal: Applied Economics* (forthcoming).

Schank, Thorsten, Claus Schnabel, and Joachim Wagner, "Do Exporters Really Pay Higher Wages? First Evidence from German Linked Employer-Employee Data," *Journal of International Economics* 72 (2007), 52–74. 10.1016/j.jinteco.2006.08.004

Schneider, Friedrich, and Dominik H. Enste, "Shadow Economies: Size, Causes, and Consequences," *Journal of Economic Literature* 38:1 (2000), 77–114. 10.1257/jel.38.1.77

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter, "Firming Up Inequality," *Quarterly Journal of Economics* 134:1 (2019), 1–50. 10.1093/qje/qjy025

Sorkin, Isaac, "Ranking Firms Using Revealed Preference," *Quarterly Journal of Economics* 133:3 (2018), 1331–1393. 10.1093/qje/qjy001

Van Reenen, John, "The Creation and Capture of Rents: Wages and Innovation in a Panel of U.K. Companies," *Quarterly Journal of Economics* 111:1 (1996), 195–226. 10.2307/2946662

Verhoogen, Eric, "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector," *Quarterly Journal of Economics* 123:2 (2008), 489–530. 10.1162/qjec.2008.123.2.489