

Bonfiglioli, Alessandra; Crinò, Rosario; Gancia, Gino

Working Paper

Betting on Exports: Trade and Endogenous Heterogeneity

CESifo Working Paper, No. 5597

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Bonfiglioli, Alessandra; Crinò, Rosario; Gancia, Gino (2015) : Betting on Exports: Trade and Endogenous Heterogeneity, CESifo Working Paper, No. 5597, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/123231>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



Betting on Exports: Trade and Endogenous Heterogeneity

Alessandra Bonfiglioli
Rosario Crinò
Gino Gancia

CESIFO WORKING PAPER NO. 5597

CATEGORY 8: TRADE POLICY

NOVEMBER 2015

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

Betting on Exports: Trade and Endogenous Heterogeneity

Abstract

We study the equilibrium determinants of firm-level heterogeneity in a model in which firms can affect the variance of their productivity draws at the entry stage and explore the implications in closed and open economy. By allowing firms to choose the size of their investment in innovation projects of unknown quality, the model yields a Pareto distribution for productivity with a shape parameter that depends on industry-level characteristics. A novel result is that export opportunities, by increasing the payoffs in the tail, induce firms to invest in bigger projects with more spreadout outcomes. Moreover, when more productive firms also pay higher wages, trade amplifies wage dispersion by making all firms more unequal. These results are consistent with new evidence on how firm-level heterogeneity and wage dispersion vary in a panel of U.S. industries. Finally, we use patent data across U.S. states and over time to provide evidence in support of a specific mechanism of the model, namely, that export opportunities increase firm heterogeneity by fostering innovation.

JEL-Codes: F120, F160, E240.

Keywords: firm heterogeneity, productivity dispersion, wage inequality, international trade.

Alessandra Bonfiglioli
University Pompeu Fabra
Department of Economics and Business
Ramon Trias Fargas, 25-27
Spain – 08005 Barcelona
alessandra.bonfiglioli@upf.edu

Rosario Crinò
Catholic University of Milan
Department of Economics and Finance
Via Necchi 5
Italy – 20123 Milan
rosario.cрино@unicatt.it

Gino Gancia
CREI
Ramon Trias Fargas, 25-27
Spain – 08005 Barcelona
ggancia@crei.cat

This draft: October 2015

We thank Gregor Schubert and Oriol Anguera for research assistance, and Gianmarco Ottaviano, Markus Poschke and seminar participants at ERWIT (Paris, 2015), Barcelona GSE Summer Forum (2015), EEA Annual Meeting (Mannheim, 2015), INSEAD, ECARES, CREI, Ivie (Workshop on Trade and Growth) for comments. We thank Antonin Bergeaud and David Hemous for sharing some of the patent data used in Aghion et al. (2015). We acknowledge financial support from the Barcelona GSE, Spanish Ministry of Economy and Competitiveness (ECO2014-55555-P and ECO2014-59805-P) and the Fundación Ramón Areces.

1 INTRODUCTION

Current research in international trade puts firm-level heterogeneity at a center stage. As documented by a growing empirical literature, firms differ in size and productivity even within narrowly defined industries (e.g., Syverson, 2004a,b) and these differences vary systematically with trade participation (e.g., Bernard et al. 2012). In particular, exporters are bigger and more productive than non-exporters, and they pay higher wages. Firm heterogeneity also has crucial implications for macroeconomic outcomes, such as aggregate efficiency (e.g., Hopenhayn, 2014). Yet, despite the growing attention that firm-level productivity differences have attracted, we still have a limited understanding of the theoretical and empirical underpinnings of this heterogeneity.

Although the distribution of the entire population of existing firms has some common characteristics that have been documented extensively (e.g., Axtell, 2001), these aggregate statistics mask significant heterogeneity across sectors and even between countries. For example, Helpman, Melitz and Yeaple (2004) show that cross-sector variation in measures of firm heterogeneity has important effects on firm strategies. Rossi-Hansberg and Wright (2007) find that the standard deviation of establishment size is smaller in sectors with larger capital shares. Poschke (2014) and Bartelsman, Haltiwanger and Scarpetta (2009) document instead differences in the firm-size distribution across countries. Moreover, given that more productive firms pay higher wages, firm heterogeneity is likely to map into wage dispersion, and wage inequality varies significantly across countries. Besides these scant observations, systematic evidence and theoretical explanations for differences in firm heterogeneity are scarce. The primary goal of this paper is to take a step towards filling this gap.

We start our analysis by documenting some little-known facts regarding how a synthetic measure of firm heterogeneity, the standard deviation of the log of sales, varies across sectors and time in the U.S. economy. We find that this measure of dispersion can differ by a factor of 10 between 6-digit NAICS industries, that it has increased on average by 11.8 per cent between 1997 and 2007, and that it correlates with industry characteristics, especially with export intensity. We also explore the robustness of these findings and argue that they are not easy to reconcile with existing models. Motivated by this evidence, we propose a novel explanation based on the idea that the observed heterogeneity stems from technological choices made when new products are introduced.

To do so, we develop a model in which endogenous investment decisions at the entry stage affect the variance of the possible realizations of productivity. We then explore the implications for the equilibrium distribution of firms and wages in closed and open economies. Leading models of heterogeneous firms take the probability dis-

tribution from which firms draw productivity as given and characterize the resulting distribution of firm-level characteristics through the dynamics of entry, exit and possibly growth. Important examples are Melitz (2003), Rossi-Hansberg and Wright (2007), Luttmer (2010) and more recently Jones and Kim (2014) and König, Lorenz and Zilibotti (2015). The aim of this paper is to take a complementary approach, namely, to recognize that firms can affect directly the expected dispersion of productivity by choosing the “riskiness” of their investment.¹

Although the success in starting a new enterprise is inherently uncertain, firms can deliberately choose between small projects with relatively safe returns and large projects with risky payoffs. Such a trade off is very familiar to anyone pursuing academic research, but is also common in the world of business. For instance, designing and assembling a new variety of laptop PCs, which mostly requires the use of established technologies, is safer and less costly than developing an entirely new product, such as tablet computers. In fact, the first tablet-like products date back to the 1980s, but did not reach success until the release of the iPad in 2010.² After decades of research, Apple’s investment was rewarded with the sale of more than 250 million units over a period of five years only.

We formalize these ideas in a multi-industry model à la Melitz (2003) in which firms can draw a random productivity level upon paying an innovation cost and there is a fixed export cost. We modify the entry stage by allowing firms to choose the size of their investment in projects of unknown quality, which in turn affects the variance of the probability distribution from which productivity is drawn. After solving for the optimal innovation size, the model yields a Pareto distribution for productivity with a shape parameter that depends on industry-level characteristics in a way consistent with the patterns found in the data. A key result is that export opportunities induce firms to draw technology from a more spread-out distribution. The reason is that trade reallocates profits in favor of the most productive firms, thereby increasing the payoffs in the tail. Hence, the chance of winning the extra prize of exporting induces firms to bet on “riskier” and bigger projects.³

¹We often use the term “riskiness” to refer to the variability of uncertain technological outcomes. To avoid confusion, we would like to stress from the onset that in our model there will be complete markets so that risk is completely diversified.

²In a 1983 speech, Steve Jobs said: “Apple’s strategy is really simple. What we want to do is we want to put an incredibly great computer in a book that you can carry around with you and learn how to use in 20 minutes... and we really want to do it with a radio link in it so you don’t have to hook up to anything and you’re in communication with all of these larger databases and other computers.” Yet, inventing the iPad required 27 years of investment constellated with failures and unforeseen spin-offs, including the development of the iPhone.

³In our model, the mean and the variance of productivity draws will be positively linked. Although this is a realistic result, we show that our main predictions still hold in an alternative model in which firms

Returning to the example of the invention of the iPad, our model suggests that globalization is what made Apple strategy so much more rewarding. Indeed, without international markets, it is difficult to imagine how Apple's revenue could have increased by a factor of 13 between 2005 and 2014. On the contrary, during the same period, the sales of less-innovative manufacturers of traditional computers, such as Dell, stagnated. There is also ample anecdotal evidence that more firms are following Apple's strategy. To name just one example, in 2010 Google started to invest in "Google X" project, a semi-secret lab dedicated to making major, high-risk, technological advancements.⁴

Next, we extend the model to show how firm heterogeneity can map into wage inequality, as in Helpman, Itskhoki and Redding (2010). When more productive firms pay higher wages, we obtain another novel result: trade amplifies wage dispersion by inducing firms to take on more technological risk *ex-ante* and hence making them more unequal *ex-post*.⁵

In the last section of the paper we go back to the data. Using individual-level wage data over the period 1997-2007 in the United States, we show that measures of wage dispersion covary with industry characteristics in a way that mirrors remarkably well the pattern found for the dispersion of sales. In particular, export opportunities increase significantly wage inequality at the industry level. We then provide a first attempt at testing a specific mechanism of our model, namely, that export opportunities increase firm heterogeneity by fostering investment in innovation. To do so, we follow Aghion et al. (2015) in switching to geographic data and use patent counts to build a measure of innovation intensity for a panel of U.S. states over the period 1989-2007. We also follow Autor, Dorn and Hanson (2013) in using the industry composition of manufacturing employment of each state to construct a state-level measure of export intensity. With this data, we document two sets of results. First, consistently with the model, innovation intensity is strongly correlated with export opportunities. Second, the dispersion of firms' sales, computed now at the state level, is strongly correlated with innovation intensity.

This paper is related to the vast literature aimed at explaining productivity differences across firms (see Syverson, 2011, for a survey). To the best of our knowledge, the choice of the riskiness of innovation, and its aggregate implications, has received little attention. Papers on technological change sometimes consider the distinction between radical and incremental innovation (e.g., Acemoglu and Cao, 2015). But these

can choose between productivity distributions that are a mean-preserving spread.

⁴To date, among Google X projects that have been revealed are the development of the self-driving car and Google Glasses.

⁵Dunne et al. (2004), Faggio, Salvanes and Van Reenen (2010) and Song et al. (2015) show that a large fraction of the observed wage dispersion is between firms.

types of innovations differ more in the degree to which they replace or complement existing technologies, rather than in the variance of the potential outcomes. Some exceptions are Gabler and Poschke (2013) and Bartelsman, Gautier and de Wind (2015), who study respectively how distortions and employment protection affect the choice between risky technologies. In any case, studying how different types of innovations affect the distribution of firms and wages is an underexplored and promising area of research.⁶

The large literature on trade with heterogeneous firms started by Melitz (2003) does study the implications of export opportunities for the distribution of existing firms (see Melitz and Redding, 2014, for an excellent survey). As it is well-known, trade can make firms more unequal by reallocating profits and workers from the least to the most productive firms. This effect is however very different from the one we emphasize, in that it abstracts from the possibility that trade changes the fundamental reason why firms are different, i.e., the unconditional productivity distribution.⁷ Moreover, the focus of our paper is on measures of dispersion of firms' attributes, such as the log of sales, that are scale invariant rather than other characteristics, such as average size or the productivity cutoff for exit, that have been studied more extensively. In this respect, our paper is close in spirit to a nascent strand of the literature aimed at exploring the effect of trade on higher moments of the distribution of firm characteristics (e.g., Mayer, Melitz and Ottaviano, 2015)

Some recent papers study the impact of trade on productivity via *ex-post* decisions on product scope or innovation. These include Bernard, Redding and Schott (2011), Dhingra (2013), Bustos (2011), Lileeva and Trefler (2010) and Atkeson and Burstein (2010), among others. All these models propose different channels through which trade liberalization can raise firm-level productivity, but do not focus on its dispersion. This literature has also shown that trade can help overcome the fixed cost of technology adoption through a scale effect, a result that is very different from our finding that trade induces firms to invest in projects with higher variance.⁸ Moreover, many technological choices must be made *ex-ante*, especially when new products are introduced.⁹ Hence, combin-

⁶Doraszelski and Jaumandreu (2009) show that the outcome of R&D investment is subject to a high degree of uncertainty and that R&D expenditure plays a key role in determining the differences in productivity across firms.

⁷Some papers, including Yeaple (2005) and more recently Grossman and Helpman (2014), trace productivity differences across firms to heterogeneity in ability across workers and managers. We follow the complementary approach that emphasizes the role of differences in technology rather than ability.

⁸There is a small literature on trade and risk taking, including interesting work by Vannoorenberghe (2014) and Fillat and Garetto (2015). In this paper we are instead interested in the determinants of technological variability abstracting from risk attitude.

⁹According to the NSF, in the period 2006-8 the number of firms reporting product innovation was as high as the number of firms reporting process innovation.

ing our model of innovation risk with *ex-post* decisions that can affect an initial realization of productivity seems a promising step forward to develop a comprehensive theory of how productivity differences emerge and evolve.

Finally, several papers have shown, both theoretically and empirically, that trade impacts wage inequality because exporters pay higher wages.¹⁰ In our model, however, the effect of trade works not only through the exporters’ wage premium, but also by making the entire wage schedule steeper, with different implications. For instance, our mechanism predicts that more export opportunities will increase wage dispersion even among the group of non-exporting firms. This may help explain why the rise in inequality is often found to be “fractal”, i.e., to hold across firm size groups (Song et al., 2015).

The remainder of the paper is organized as follows. In Section 2, we document some stylized facts regarding how the dispersion of log sales varies across sectors and time in the U.S. economy. Motivated by these empirical observations, in Section 3 we propose a closed-economy model where differences in the variance of firm-level outcomes originate from technological choices at the entry stage. Section 4 adds costly trade and shows that more export opportunities induce firms to draw their productivity from riskier distributions, thereby generating more heterogeneity in equilibrium. In Section 5 we consider the implications of the model for income and wage inequality. Section 6 goes back to the data to test the predictions of the model on wage dispersion and on innovation. Section 7 concludes.

2 MOTIVATING EVIDENCE: SALE DISPERSION AND TRADE

In this section, we document how the dispersion of sales of U.S. firms varies across industries and over time, and how it correlates with a number of industry characteristics. First, we show that the dispersion of sales differs significantly across industries and has increased over time. Second, we present panel regressions suggesting that higher dispersion at the industry level is systematically associated with larger scale in terms of average sales and with higher export intensity. Finally, we provide additional evidence suggestive of a causal effect of export intensity on firm dispersion.¹¹

¹⁰See, for example, Helpman, Itskhoki and Redding (2010), Helpman et al. (2015) and other papers surveyed in section 10 of Melitz and Redding (2014).

¹¹An antecedent of this analysis is Syverson (2004b), who studies how various measures of productivity dispersion covary with industry characteristics in the U.S. manufacturing sector. Yet, his evidence is limited to the 1977 cross-section.

2.1 SALE DISPERSION ACROSS INDUSTRIES AND OVER TIME

Our main measure of dispersion is the standard deviation of the logarithm of sales per establishment.¹² We focus on sales because they are an easy-to-observe, synthetic, measure of overall size, and we take logs to make the variance scale invariant. We compute this variable using data from the ‘Statistics of U.S. Businesses’ of the U.S. Census Bureau for the years 1997, 2002 and 2007.¹³ Data on (receipts of) sales and numbers of firms, establishments and employees are available for 453 6-digit NAICS industries aggregated into sales-size categories. Since we do not have access to the underlying firm-level data, we follow Helpman, Melitz and Yeaple (2004) in assuming that all establishments falling within the same bin have sales equal to the group mean. Then, we consider each bin in a 6-digit NAICS industry as a single observation and compute the standard deviation of log establishment sales using the number of establishments as weights.¹⁴ Helpman, Melitz and Yeaple (2004) show that this methodology for computing dispersions approximates well other measures based on the entire population of firms. As an additional check, we have also computed the variance of log sales using firm-level data from Compustat, which only includes listed firms. Although the Compustat sample is too restricted for the purpose of our analysis, we find that the variance computed on that sample is highly correlated with the one we obtain from Census data.

Table 1 reports some descriptive statistics. For 3-digit manufacturing sectors, it shows the average standard deviation of log establishment sales in 2007, its minimum and maximum in each 6-digit sub-industry, and its average percentage change over the previous ten years.¹⁵ For convenience, sectors are ordered by increasing dispersion. The first column shows that the dispersion of sales varies significantly across sectors, ranging from a minimum of 1.69 (in Paper Manufacturing) to a maximum of 2.71 (in Transportation Equipment Manufacturing). The second and third columns show, however, that the main source of heterogeneity is within 3-digit sectors: among all 6-digit industries, the dispersion of sales varies by a factor of 10, as shown in the last row of the table.

¹²In what follows, we will refer indifferently to establishments and plants.

¹³In the ‘Statistics of U.S. Businesses’, information on firms’ sales is released during Census years and is currently available for the years 1997, 2002, 2007 and 2012. The substantial restructuring of the NAICS classification occurred in 2012 makes it impossible to create a mapping between the latest wave of the data and the preceding ones for many industries. We therefore use the first three waves, which also ensures that our results are not contaminated by the recent financial crisis.

¹⁴The raw data are disaggregated into 10 bins in 1997, 8 bins in 2002 and 18 bins in 2007. The lowest bin contains firms with revenue below 50 thousand US\$ while the highest bin contains firms with revenue above 100 millions US\$. In most of the analysis, we aggregate the data into six bins consistently defined throughout the period. In a robustness check, we show that our results are unchanged when using the raw bins.

¹⁵All reported averages are simple averages. Weighting by sales does not affect the qualitative results.

Table 1: Descriptive Statistics on the Dispersion of Sales in U.S. Manufacturing

NAICS code	Industry Description	Std. Dev. of Log Establishment Sales				# of Est. Mean
		Mean	Min.	Max.	% Change	
322	Paper Manufacturing	1.69	0.96	2.27	0.07	252
327	Nonmetallic Mineral Products Manufacturing	1.89	0.97	3.56	0.03	728
333	Machinery Manufacturing	1.97	0.48	3.21	-0.03	498
313	Textile Mills	2.00	1.17	2.97	-0.11	258
331	Primary Metal Manufacturing	2.07	1.44	3.10	0.09	203
326	Plastics and Rubber Products Manufacturing	2.07	1.36	3.08	0.08	1024
315	Apparel Manufacturing	2.09	1.00	2.94	0.05	546
332	Fabricated Metal Products Manufacturing	2.12	1.22	3.33	0.10	1387
324	Petroleum and Coal Products Manufacturing	2.17	0.75	4.48	0.47	482
316	Leather and Allied Products Manufacturing	2.20	0.89	3.26	0.06	139
339	Miscellaneous Manufacturing	2.24	0.95	3.67	0.23	1571
325	Chemical Manufacturing	2.24	0.55	4.01	0.06	394
337	Furniture and Related Products Manufacturing	2.26	0.92	3.18	0.18	1753
334	Computer and Electronic Products Manufacturing	2.26	0.88	3.76	-0.01	499
321	Wood Products Manufacturing	2.43	1.46	3.25	0.32	1187
335	Electrical Equipment, Appliance, and Component Manufacturing	2.47	1.79	3.66	0.05	279
312	Beverage and Tobacco Products Manufacturing	2.52	1.98	3.12	0.31	452
323	Printing and Related Support Activities	2.54	1.33	3.31	0.37	2773
311	Food Manufacturing	2.59	1.07	4.57	0.41	549
314	Textile Products Mills	2.59	1.45	3.34	0.32	649
336	Transportation Equipment Manufacturing	2.71	1.72	4.22	0.02	404
	Total	2.23	0.48	4.57	0.12	720

Notes: The standard deviation of log establishment sales is computed for each 6-digit NAICS industry on six sales-size bins homogeneous for all years. Mean, minimum and maximum refer to the year 2007; percentage changes are computed between 1997 and 2007. Statistics are computed across all 6-digit industries belonging to a given 3-digit code.

The fifth column reports the average number of establishments in each 3-digit sector. Comparing the first and fifth columns reassures that the dispersion of sales in a sector is not mechanically driven by sample size. Finally, the fourth column shows that the dispersion of sales has increased remarkably between 1997 and 2007, on average by 11.8 per cent (28.5 per cent if we weight industries by sales). Although this rise in dispersion is not a well-known stylized fact, it is consistent with the evidence in Dunne et al. (2004), who find that inequality in productivity across U.S. manufacturing plants increased between 1975 and 1992, and in Faggio, Salvanes and Van Reenen (2010), who find similar results for the United Kingdom between 1984 and 2001.

2.2 EXPLORING THE DATA: CORRELATIONS

To further explore the data, we now exploit variation across 6-digit industries and over time, and study how the dispersion of sales correlates with a number of industry characteristics. Among the covariates, we consider average sales per establishment, export intensity, total employment, the number of establishments, the intensities in skill, physical capital and raw materials, and the mean and standard deviation of log work-

ers' educational attainment. Export intensity is the ratio of exports to total shipments, constructed with export data from Schott (2008) and shipment data from the NBER-CES Manufacturing Industry Productivity Database, and captures engagement in global markets.¹⁶ Total employment is the number of employees at the industry level sourced from the U.S. Census Bureau. Skill, capital and material intensities are computed as in Romalis (2004) with data from the NBER-CES Manufacturing Industry Productivity Database, and are equal to the ratios of non-production workers' wage bill, capital compensation and material expenditure, respectively, over the sum of value added and material costs. These variables are meant to capture technological characteristics of industries. Finally, the mean and standard deviation of workers' education are computed with data from the CPS Merged Outgoing Rotation Groups. Including these variables can help account for the role of sorting between firms and workers based on observables.¹⁷

We regress the standard deviation of log establishment sales on these industry characteristics. To account for cyclical conditions, we also control for the growth rate of nominal GDP over the two years prior to each observation.¹⁸ We estimate specifications without industry fixed effects, so as to exploit the whole variation across industries and years, as well as with industry fixed effects, so as to control for unobserved industry heterogeneity and identify the coefficients through within-industry variation over time. Finally, we also estimate specifications with industry fixed effects and variables in first differences, so as to control for industry-specific time trends.

Table 2 reports the baseline OLS results. Columns (1)-(4) show coefficients from pooled-OLS specifications, columns (5) and (6) control for industry fixed effects, and column (7) reports the estimates from the first-difference specification. We work with a consistent sample of observations for which we have data on all covariates, and correct the standard errors for clustering within 6-digit industries. In columns (1) and (2) we start by regressing sale dispersion on establishment sales and export intensity, respectively. We find that the dispersion of sales is strongly positively correlated with average plant size and exports across industries and time. In column (3) we include establishment sales and export intensity jointly with other industry characteristics. The correlation of sale dispersion with plant size and export intensity is unchanged. Moreover, the dispersion of sales is positively correlated with the dispersion of workers' education, although this result is not robust across specifications as shown below. In column (4) we

¹⁶We will also control for import penetration. However, given the collinearity between the two variables, we focus mostly on export intensity, which is found to have statistically stronger effects.

¹⁷See the Appendix for more details on variables definitions and data sources.

¹⁸The growth rate of nominal GDP is computed with data from the World Development Indicators. All variables except for GDP growth and standard deviations are expressed in logarithms.

Table 2: Dispersion of Sales and Industry Characteristics: Baseline Estimates

	Pooled	Pooled	Pooled	Pooled	Industry Fixed Eff.	Industry Fixed Eff.	Industry FE + First Diff.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log establishment sales	0.203*** [0.020]		0.194*** [0.056]	0.210*** [0.057]	0.587*** [0.110]	0.678*** [0.111]	0.519** [0.248]
Log export intensity		0.076*** [0.015]	0.091*** [0.015]	0.089*** [0.015]	0.130*** [0.047]	0.075* [0.046]	0.171** [0.070]
Log number of establishments			0.036 [0.070]	0.053 [0.071]	0.449** [0.176]	0.556*** [0.175]	0.959** [0.391]
Log employment			0.130* [0.069]	0.109 [0.070]	-0.159 [0.108]	-0.320*** [0.110]	-0.252 [0.336]
Log capital intensity			0.064 [0.108]	0.041 [0.108]	-0.037 [0.220]	-0.181 [0.222]	0.074 [0.192]
Log skill intensity			-0.094* [0.054]	-0.082 [0.054]	0.202 [0.124]	0.339*** [0.126]	-0.172 [0.199]
Log material intensity			-0.218 [0.156]	-0.248 [0.157]	-0.420 [0.344]	-0.759** [0.354]	0.190 [0.422]
Log average education			0.283 [0.431]	0.258 [0.432]	-0.626 [0.907]	-0.442 [0.874]	-0.853 [1.222]
S.D. log education			1.032*** [0.363]	1.021*** [0.362]	0.205 [0.538]	0.129 [0.515]	-0.314 [0.658]
GDP growth				2.091*** [0.463]		3.464*** [0.497]	
Obs.	1,036	1,036	1,036	1,036	1,036	1,036	670
R ²	0.17	0.03	0.29	0.30	0.67	0.69	0.46

Notes: The dependent variable is the standard deviation of log establishment sales. It is computed for each 6-digit NAICS industry on six sales-size bins homogeneous for all years. All variables except for GDP growth are observed at the 6-digit NAICS industry level in the years 1997, 2002 and 2007. All industry-level controls are contemporaneous to the dependent variable. GDP growth is computed over the two years prior to each observation. All specifications are estimated with Ordinary Least Squares. Standard errors are clustered by 6-digit industry and reported in brackets. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

add GDP growth. The coefficient on this variable is positive and precisely estimated, suggesting that the dispersion of sales is higher in periods of economic expansion. The dispersion of sales is still positively correlated with average plant size and export intensity.

In columns (5) and (6) we add industry fixed effects. The coefficients on establishment sales and export intensity remain positive and statistically significant. Both coefficients are larger than before, suggesting that the correlations are stronger within industries and that unobserved industry heterogeneity tends to bias these coefficients downward. The dispersion of sales also remains positively correlated with GDP growth. Finally, in column (7) we express all variables in first differences and include industry fixed effects to control for industry-specific trends. We still find sale dispersion to be positively correlated with establishment sales and export intensity.¹⁹

¹⁹GDP growth is expressed in differences over the previous two years, so we exclude this variable from the first-differenced specifications to avoid double differencing, which would not have a clear interpreta-

Table 3: Dispersion of Sales and Industry Characteristics: Robustness Checks

	Dep. Var.: S.D. of Log Establishment Sales				Dep. Var.: S.D. of Log Labor Productivity		
	Industry Fixed Effects and First Differences						
	Controlling for Import Penetr.	Excluding Smallest Bins	Excluding Largest Bin	Using Raw Bins	Pooled	Industry Fixed Eff.	Industry FE + First Differences
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log establishment sales	0.519** [0.249]	0.367** [0.175]	0.106 [0.177]	0.544** [0.230]	0.171*** [0.029]	0.282*** [0.067]	0.262*** [0.095]
Log export intensity	0.169** [0.069]	0.094* [0.049]	0.142* [0.079]	0.157** [0.068]	0.028*** [0.008]	0.110*** [0.033]	0.149*** [0.038]
Log number of establishments	0.955** [0.393]	0.356 [0.268]	0.811** [0.383]	0.867** [0.354]	0.251*** [0.035]	0.399*** [0.107]	0.758*** [0.223]
Log employment	-0.252 [0.336]	-0.164 [0.231]	-0.067 [0.285]	-0.295 [0.322]	-0.206*** [0.032]	-0.283*** [0.073]	-0.432*** [0.144]
Log capital intensity	0.076 [0.193]	-0.191 [0.143]	0.429** [0.205]	0.044 [0.169]	-0.009 [0.063]	-0.009 [0.148]	0.112 [0.175]
Log skill intensity	-0.178 [0.202]	-0.189 [0.132]	-0.101 [0.195]	-0.081 [0.184]	-0.084*** [0.030]	0.036 [0.087]	0.007 [0.087]
Log material intensity	0.191 [0.426]	-0.224 [0.287]	0.817** [0.414]	0.069 [0.390]	-0.097 [0.093]	-0.249 [0.273]	0.150 [0.257]
Log average education	-0.818 [1.231]	-0.023 [0.586]	0.127 [1.288]	-1.065 [1.133]	0.119 [0.215]	0.107 [0.559]	0.077 [0.619]
S.D. log education	-0.293 [0.657]	-0.154 [0.347]	-0.134 [0.630]	-0.492 [0.615]	0.420** [0.207]	0.027 [0.312]	0.060 [0.316]
GDP growth					0.410* [0.209]	0.659** [0.308]	
Log import penetration	0.026 [0.082]						
Obs.	670	660	646	670	1,029	1,029	664
R ²	0.46	0.54	0.48	0.46	0.16	0.48	0.49

Notes: The dependent variable, indicated in the heading of each column, is computed for each 6-digit NAICS industry on six sales-size bins homogeneous for all years, except for: column (2), where it is computed excluding observations for the two smallest sales-size bins for each industry; column (3), where it is computed excluding observations for the largest sales-size bin for each industry; and column (4), where it is computed on all sales-size bins available for each industry and year. In columns (5)-(7), the dependent variable is the dispersion of log labor productivity, defined as sales per worker. All variables except for GDP growth are observed at the 6-digit NAICS industry level in the years 1997, 2002 and 2007. All industry-level controls are contemporaneous to the dependent variable. GDP growth is computed over the two years prior to each observation. All specifications are estimated with Ordinary Least Squares. Standard errors are clustered by 6-digit industry and reported in brackets. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

In Table 3, we perform robustness checks using the richest specification in first differences. First, we add import penetration in column (1). The results show that sale dispersion correlates significantly with export opportunities but not with import competition. In columns (2) and (3), we re-estimate our specification computing the dependent variable on two subsets of establishments, which exclude the smallest plants (bottom 25 per cent of the sample) and largest plants (top 21 per cent of the sample), respectively. The main results hold in both cases, suggesting that they are not driven by large or small establishments. As a further robustness check, in column (4) we re-compute the standard deviation of log plant sales using all sales-size bins available in each period, so as to fully exploit the information contained in the data set. The coefficients are similar to those in column (7) of Table 2, suggesting that the number of bins does not affect our results. Finally, we estimate specifications using the standard deviation of log sales per

tion.

worker as the dependent variable. The coefficients in columns (5)-(7) show that labor productivity is also more dispersed in industries with larger plant size and higher export intensity, and is positively correlated with the business cycle.

Our next step is to dig deeper into the empirical results found so far. The positive correlation between the mean and the dispersion of sales is an important property of the empirical distribution. This property holds both for Pareto and log-normal distributions, which are known to describe the data well. Hence, the positive coefficient for average establishment sales is reassuring, but probably not too surprising.²⁰ The strong correlation between export intensity and the standard deviation of log plant sales is instead in our view more interesting and deserves further investigation.

2.3 IDENTIFYING THE EFFECT OF TRADE

The positive correlation between export intensity and sale dispersion may be consistent with causality running in both directions. On the one hand, industries with higher dispersion of sales may have a higher export intensity, because they are more likely to host large firms which tend to participate more in export markets (see Bernard et al., 2012). On the other hand, better export opportunities may amplify differences between firms. Hence, understanding this correlation requires sorting out the direction of causality. We try to do so using two alternative identification strategies. First, in the spirit of Autor, Dorn and Hanson (2013), we follow an instrumental variables approach to identify the effect of exogenous export shocks on the standard deviation of log plant sales. Second, following Hummels (2007) and Hummels et al. (2014), we adopt a difference-in-differences strategy based on changes in transportation costs with heterogeneous effects across industries.

In Table 4, we estimate the main specifications instrumenting export intensity with the exports of all non-U.S. countries to the destination markets of the United States in each industry and year. Our aim is to identify variation in U.S. exports generated by foreign demand shocks, which would raise both U.S. and third countries' exports to the same destination market, while cleaning out variation due to U.S. industry-specific technological shocks, which would only raise U.S. exports and could induce reverse causality. The exclusion restriction is that, conditional on our large set of control variables, technological shocks originated in U.S. industries are uncorrelated with those originated in other countries' industries. Because the vast majority of countries used to construct the instrument are low- or middle-income economies, their technological

²⁰Explaining why firm heterogeneity is Pareto or log-normal is a very interesting question on its own, but it is not the goal of this paper.

Table 4: Dispersion of Sales and Industry Characteristics: IV Regressions

	Dep. Var.: S.D. of Log Establishment Sales			Dep. Var.: S.D. of Log Labor Productivity		
	Pooled	Industry Fixed Eff.	Industry FE + First Diff.	Pooled	Industry Fixed Eff.	Industry FE + First Diff.
	(1)	(2)	(3)	(4)	(5)	(6)
Log establishment sales	0.227*** [0.058]	0.700*** [0.103]	0.812*** [0.229]	0.172*** [0.030]	0.267*** [0.064]	0.301*** [0.102]
Log export intensity	0.141*** [0.030]	0.297** [0.137]	0.707*** [0.167]	0.050*** [0.017]	0.181* [0.094]	0.316*** [0.095]
Log number of establishments	0.089 [0.072]	0.599*** [0.170]	1.129*** [0.424]	0.259*** [0.036]	0.407*** [0.107]	0.743*** [0.240]
Log employment	0.092 [0.071]	-0.221* [0.133]	-0.486 [0.322]	-0.205*** [0.033]	-0.249*** [0.084]	-0.440*** [0.152]
Log capital intensity	0.054 [0.111]	-0.154 [0.227]	0.039 [0.216]	0.003 [0.066]	0.007 [0.147]	0.096 [0.171]
Log skill intensity	-0.106* [0.058]	0.368*** [0.121]	0.007 [0.205]	-0.099*** [0.032]	0.025 [0.086]	0.021 [0.091]
Log material intensity	-0.283* [0.164]	-0.837** [0.347]	-0.314 [0.464]	-0.109 [0.098]	-0.266 [0.272]	0.043 [0.266]
Log average education	-0.035 [0.455]	-0.907 [0.875]	-0.880 [1.308]	-0.018 [0.237]	0.073 [0.572]	0.176 [0.630]
S.D. log education	0.863** [0.378]	0.009 [0.505]	-0.279 [0.710]	0.317 [0.221]	0.003 [0.311]	0.120 [0.327]
GDP growth	1.925*** [0.470]	2.940*** [0.557]		0.377* [0.208]	0.551 [0.335]	
Obs.	1,015	1,001	630	1,008	992	620
First-stage results						
Log world exports	0.447*** [0.068]	0.397*** [0.044]	0.546*** [0.068]	0.447*** [0.068]	0.398*** [0.045]	0.568*** [0.064]
Kleibergen-Paap <i>F</i> -statistic	43.8	81.1	65.3	43.0	79.9	77.7

Notes: The dependent variables, indicated in columns' headings, are computed for each 6-digit NAICS industry on six sales-size bins homogeneous for all years. All variables except for GDP growth are observed at the 6-digit NAICS industry level in the years 1997, 2002 and 2007. All industry-level controls are contemporaneous to the dependent variable. GDP growth is computed over the two years prior to each observation. All specifications are estimated with Two-Stage Least Squares. Export intensity is instrumented with non-U.S. exports to the destination markets of the U.S. in each industry and year. Standard errors are clustered by 6-digit industry and reported in brackets. *F*-statistics are reported for the Kleibergen-Paap test for weak instruments. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

shocks are unlikely to be correlated with those hitting the United States.

The first-stage results reported in the bottom of the table show that the instrument has strong power for predicting U.S. export intensity. The first-stage coefficient is always precisely estimated and large, ranging between 0.4 and 0.57 across specifications, and the Kleibergen-Paap *F*-statistics for excluded instruments are remarkably high, ranging from 43 to 81.1. The second-stage results show that the coefficient on export intensity is positive and statistically significant in all specifications, consistently with a causal effect of exports on the dispersion of sales and productivity. If anything, the IV coefficients are slightly larger than their OLS counterparts, suggesting that the potential upward bias due to reverse causality is more than compensated in our data by attenuation bias due to measurement error.

To have a sense of the size of the effect of trade, note that in our sample export intensity has a standard deviation of 1.34, while the standard deviation of sale dispersion is equal to 0.58. Then, our coefficients imply that a 1 s.d. increase in export intensity would raise the dispersion of sales by 0.17-0.69 s.d.. The observed increase in export intensity over the sample period (21 per cent) explains between 13 and 52 per cent of the increase in sale dispersion over 1997-2007.

Finally, we show that our evidence is qualitatively unchanged when using an alternative strategy, which does not rely on instrumental variables but identifies the effect of exports exploiting heterogeneous changes in transportation costs across industries. In particular, we regress our dispersion measures on the time series for oil price (Brent), the bulk weight (in Kg per US\$) of U.S. shipments in each industry, and the interaction between the two. The interaction coefficient is identified by the differential response to a common oil price shock across industries that produce goods of different weight and are thus characterized by a different importance of transportation costs. Hence, we refer to this approach as a difference-in-differences strategy. A negative estimate for the interaction coefficient would imply that, when hit by a reduction in oil price, industries shipping heavier goods - which see a larger drop in trade costs (increase in export opportunities) - experience a larger increase in dispersion.

To construct the bulk weights we use product-level export data. We define the bulk weight of a given industry as the export-weighted average of the bulk weights of its constituent products, which in turn are computed as averages between air and vessel transportation in 1995, to ensure that the choice of transport mode does not react to changes in oil price. Table 5 reports coefficient estimates for various specifications. The interaction coefficient is always negative and precisely estimated, which also points in the direction of a causal effect of export opportunities on the dispersion of sales and labor productivity.

2.4 DISCUSSION

The stylized facts documented in this section raise a number of important questions. Why is higher export intensity associated to more heterogeneity across establishments? More generally, what drives changes in the distribution of sales and productivity? Before proposing our explanation, we pause to discuss why we think that existing models do not provide complete and fully satisfactory answers.

Natural candidates for explaining changes in the distribution of firm characteristics could be “granularity” in the data and misallocation. As documented by a recent literature, the law of large numbers may fail at the industry level, especially if the distribution

Table 5: Dispersion of Sales and Industry Characteristics: Diff-in-Diff Specifications

	Dep. Var.: S.D. of Log Establishment Sales			Dep. Var.: S.D. of Log Labor Productivity		
	Pooled	Industry Fixed Eff.	Industry FE + Time Dummies	Pooled	Industry Fixed Eff.	Industry FE + Time Dummies
	(1)	(2)	(3)	(4)	(5)	(6)
Log establishment sales	0.170*** [0.062]	0.677*** [0.158]	0.677*** [0.158]	0.150*** [0.036]	0.119 [0.074]	0.119 [0.074]
Bulk weight * Log oil price	-0.071** [0.030]	-0.097** [0.040]	-0.097** [0.040]	-0.027** [0.012]	-0.032** [0.014]	-0.032** [0.014]
Bulk weight	0.081 [0.109]			0.062 [0.040]		
Log oil price	0.142*** [0.034]	0.053 [0.056]		0.058** [0.026]	0.093** [0.040]	
Log number of establishments	-0.020 [0.074]	0.680*** [0.226]	0.680*** [0.226]	0.228*** [0.041]	0.324** [0.127]	0.324** [0.127]
Log employment	0.144* [0.076]	-0.401** [0.162]	-0.401** [0.162]	-0.200*** [0.041]	-0.201** [0.099]	-0.201** [0.099]
Log capital intensity	-0.014 [0.121]	-0.331 [0.216]	-0.331 [0.216]	-0.008 [0.079]	-0.016 [0.172]	-0.016 [0.172]
Log skill intensity	-0.111* [0.060]	0.280* [0.155]	0.280* [0.155]	-0.104*** [0.035]	0.079 [0.108]	0.079 [0.108]
Log material intensity	-0.199 [0.171]	-0.874** [0.376]	-0.874** [0.376]	-0.051 [0.111]	-0.211 [0.317]	-0.211 [0.317]
Log average education	0.498 [0.457]	-1.090 [1.107]	-1.090 [1.107]	0.252 [0.246]	-0.115 [0.731]	-0.115 [0.731]
S.D. log education	1.056*** [0.392]	-0.444 [0.627]	-0.444 [0.627]	0.495** [0.236]	-0.116 [0.389]	-0.116 [0.389]
GDP growth	1.634*** [0.520]	3.307*** [0.556]		0.278 [0.232]	0.506 [0.366]	
Obs.	776	776	776	771	771	771
R ²	0.32	0.66	0.66	0.18	0.48	0.48

Notes: The dependent variables, indicated in columns' headings, are computed for each 6-digit NAICS industry on six sales-size bins homogeneous for all years. All variables except for GDP growth and oil price are observed at the 6-digit NAICS industry level in the years 1997, 2002 and 2007. All industry-level controls, except for the bulk weight, are contemporaneous to the dependent variable. The bulk weight is expressed in Kg per US\$ shipped by air and/or vessel, and refers to the year 1995. GDP growth is computed over the two years prior to each observation. All specifications are estimated with Ordinary Least Squares. Standard errors are clustered by 6-digit industry and reported in brackets. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

of sales is very fat tailed. Hence, sales differences can simply be due to granularity. There is also a growing literature suggesting that firm heterogeneity partly reflects misallocation, i.e., the presence of frictions allowing inefficient firms to survive. However, these explanations do not seem fully consistent with the observation that changes in heterogeneity are driven neither by small firms (which may reflect more inefficient entry) nor by large firms (which should matter more under the granular hypothesis).

Turning to the correlation between the standard deviation of log sales and export intensity, candidate explanations could be reverse causality, selection effects and reallocations towards exporters. Once again, although these mechanisms can play an important role, they do not seem to provide a complete account of the patterns in the data. As already argued, more dispersion in productivity can increase the fraction of exporters

when the latter are in the tail of the distribution, but the IV and difference-in-difference results suggest that reverse causality is not the whole story.

The effect of exports may also be due to the fact that trade opening forces the least productive firms to exit and reallocates market shares towards larger firms. Yet, the impact of this reallocation is theoretically ambiguous, in that it depends on the characteristics of the firms it affects. Instead, we have found that trade is associated to more dispersion even when removing smaller firms, which are more likely to exit, or bigger firms, which are more likely to export. In particular, only about 18 per cent of U.S. firms export and yet all our findings are unchanged when we remove the top 21 per cent of our sample. Moreover, trade raises also the dispersions of labor productivity, which should be less directly affected by reallocations. Finally, the effect of trade seems to go beyond a mere scale effect, because all regressions control for average firm size.

In the remainder of the paper, we will propose a novel explanation based on the idea that firm heterogeneity stems from endogenous technological choices made at the time when new products are introduced. Besides its intuitive appeal and its ability to fit the empirical findings of this section, our modelling strategy is inspired by several observations. First, Dunne et al. (2004) and Faggio, Salvanes and Van Reenen (2010) show evidence suggesting that changes in productivity dispersion appear to be related to new technologies. Second, by focusing on technological choices made before heterogeneity is realized, our theory will be able to explain changes in dispersion across the whole size distribution. This is important in light of recent findings that inequality is “fractal”. Moreover, the emphasis on product innovation is empirically relevant, given that every year about 25 per cent of consumer goods sold in U.S. markets are new (Broda and Weinstein, 2010). Third, in our model export opportunities will induce all firms to choose technologies with more uncertain outcomes, a prediction that seems consistent with the finding by di Giovanni and Levchenko (2012) that volatility is higher in sectors that are more open to trade.

3 CLOSED-ECONOMY MODEL

We now build a multi-sector, one factor, model of monopolistic competition between heterogeneous firms along the lines of Melitz and Redding (2014). After paying an entry (innovation) cost, firms draw their productivity from some distribution and exit if they cannot profitably cover a fixed cost of production. Differently from Melitz (2003), we allow firms to affect the variance of their productivity draws. In this section, we characterize the resulting endogenous distribution of firm-level variables in a closed economy. We defer to the next section the case in which firms can engage in costly trade. For

simplicity, we consider a static model in which entry and production decisions are all simultaneous.

3.1 PREFERENCES

Consider an economy populated by a unit measure of identical households of size L with quasi-linear preferences over consumption of a homogenous good q_0 and differentiated goods produced in I industries:

$$U = q_0 + \sum_{i=1}^I \frac{\alpha_i X_i^{\zeta_i}}{\zeta_i}, \quad \zeta_i \in (0, 1) \quad \alpha_i > 0$$

Each industry $i \in \{1, \dots, I\}$ produces differentiated varieties and preferences over these varieties take the constant elasticity of substitution form:

$$X_i = \left[\int_{\omega \in \Omega_i} x_i(\omega)^{\frac{\sigma_i-1}{\sigma_i}} d\omega \right]^{\frac{\sigma_i}{\sigma_i-1}}, \quad \sigma_i > 1$$

where $x_i(\omega)$ is consumption of variety ω , Ω_i denotes the set of varieties produced in sector i and σ_i is the elasticity of substitution between varieties within an industry. We denote by $p_i(\omega)$ the price of variety ω in industry i and by P_i the ideal price of the consumption basket X_i :

$$P_i = \left[\int_{\omega \in \Omega_i} p_i(\omega)^{1-\sigma_i} d\omega \right]^{1/(1-\sigma_i)}.$$

The demand for the differentiated basket X_i is $X_i = (\alpha_i / P_i)^{1/(1-\zeta_i)}$ and the demand for each individual variety is

$$x_i(\omega) = X_i \left(\frac{P_i}{p_i(\omega)} \right)^{\sigma_i}. \quad (1)$$

The demand for the homogenous good q_0 is residual. We assume that income of each household is sufficiently high to always guarantee a positive consumption of the homogenous good, which is chosen as the numeraire. In the remainder of the paper, we focus on a single sector and derive results that do not depend on general equilibrium effects. For this reason, and to save notation, from now on we remove the index i with the understanding that all parameters can potentially vary across sectors.

3.2 PROBLEM OF THE FIRM

Within each sector, every variety ω is produced by monopolistically competitive firms that are heterogeneous in their labor productivity, φ . Since all firms with the same productivity behave symmetrically, we index firms by φ . There are fixed costs of production and of entry, all in units of labor. At the entry stage, a firm can choose how much to invest in innovation, a choice that affects the variance of the possible realizations of productivity. Next, the firm faces standard production and pricing decisions. We solve the problem backwards: first, we describe the strategy of a firm with a given productivity and then solve for investment at the entry stage given rational expectations on the industry equilibrium. Note also that, as it is customary, we follow the convention of identifying firms with varieties. We also assume that labor productivity in the homogeneous sector is one so that the wage is one.

A firm with productivity φ will choose its price and whether to exit so as to maximize profit, $\pi(\varphi)$, subject to a downward-sloping demand curve with elasticity σ . The first-order conditions for this problem imply that firms set prices equal to a constant markup over the marginal cost,

$$p(\varphi) = \frac{\sigma}{\sigma - 1} \frac{1}{\varphi}, \quad (2)$$

and exit if $\pi(\varphi) < 0$. Using (1) and (2), we can express profit as a function of productivity:

$$\pi(\varphi) = A\varphi^{\sigma-1} - f, \quad (3)$$

where $A = \left(\frac{\sigma}{\sigma-1}\right)^{1-\sigma} \frac{XP^\sigma}{\sigma}$. Since profits are increasing in φ , the firm will exit whenever its productivity is below the cutoff $\varphi^* = (f/A)^{1/(\sigma-1)}$.

We now consider the entry stage. As in Melitz (2003), firms pay a sunk innovation cost to be able to manufacture a new variety with productivity drawn from some distribution with c.d.f. $G(\varphi)$. Hence, combining the pricing and exit decision, we can write *ex-ante* expected profit as:

$$\mathbb{E}[\pi] = \int_0^\infty \pi(\varphi) dG(\varphi) = \int_{\varphi^*}^\infty (A\varphi^{\sigma-1} - f) dG(\varphi). \quad (4)$$

We depart from the canonical approach by making the distribution $G(\varphi)$ endogenous. More precisely, we now develop a simple model of investment in innovation projects generating a Pareto distribution for φ with a mean and variance that depend on firms' decisions. The model will formalize the intuitive idea that firms can choose between small projects with relatively safe returns and larger projects with more unpredictable outcomes.

Before continuing, we pause to discuss briefly why we focus on Pareto distributions. Our choice is based both on empirical and theoretical considerations. First, the Pareto distribution is widely used in the literature and has been shown to approximate well some observed firm-level characteristics.²¹ Thus, it is an empirically plausible assumption. The second reason is analytical tractability. The convenient properties of Pareto distributions allow us to derive closed-form solutions for various measures of firm heterogeneity, which helps in mapping the model to the data. In particular, recall that our motivation is to explain some stylized facts about the variance of the logarithm of firm characteristics, as documented in Section 2. Since the standard deviation of the log of a Pareto-distributed variable is just equal to the inverse of the shape parameter, it is clear why we are interested in endogenizing this parameter. We now show a simple way of doing this.

Suppose that, in order to enter, a firm must invest in an innovation project. The outcome of the project is a technology allowing the firm to manufacture a new variety with productivity φ . The realization of this productivity depends both on the quality of the project, which is uncertain, and size of it, which is a choice variable. More precisely, assume that quality, q , of new projects is random and exponentially distributed:

$$\Pr [q > x] = \exp(-\kappa x),$$

with support $x \in [0, \infty)$ and rate $\kappa > 0$, capturing how “compressed” the distribution is. Notice that quality is inherently uncertain and its realization is beyond control.²² The firm can instead choose the size of the project, s , with minimum $\underline{s} > 0$.²³ To obtain a Pareto distribution for φ , we assume that productivity depends both on the quality and the size of the project as follows:

$$\ln \varphi = sq + \ln \varphi_{\min}.$$

with $\varphi_{\min} > 0$. This equation embeds a complementarity between quality and size: resources invested in a bad project ($q = 0$) are wasted, in that they do not increase φ , while even a great idea is useless without some investment to implement it. More

²¹Although Head, Mayer and Thoenig (2014) argue that the log-normal distribution provides a better description of the empirical distribution of firm sales, they also find a considerable overlap with the Pareto distribution, a result echoed in Mrazova, Neary and Parenti (2015).

²²In reality, firms might be able to distinguish between “riskier” and “safer” ideas. Yet, our purpose here is to show that exposure to this kind of innovation risk can be controlled even when the deep source of uncertainty is exogenous.

²³A positive minimum, even if arbitrarily small, simplifies the analysis by ruling out the case of a degenerate distribution.

importantly, these assumptions imply that φ is Pareto distributed with minimum φ_{\min} and shape κ/s , as can be seen from:

$$1 - G(\varphi) = \Pr \left[q > \frac{\ln(\varphi/\varphi_{\min})}{s} \right] = \left(\frac{\varphi}{\varphi_{\min}} \right)^{-\frac{\kappa}{s}}.$$

Hence, by choosing the size of the project, the firm is choosing to draw φ from different Pareto distributions, identified by the new parameter $v \equiv s/\kappa$. Note that the standard deviation of the log of φ is equal to v , which can therefore be taken as an index of the dispersion of the distribution. As such, v can be interpreted as a measure of “risk”, in the sense that the outcome of the investment project becomes less predictable and more dispersed when v is higher. As shown below, v is one of the key determinants of the equilibrium distributions of the log of firm characteristics, such as sales. Moreover, v also affects the expected value of φ , which is equal to $\varphi_{\min} (1 - v)^{-1}$, so that mean and variance are linked. Although we consider this a realistic property, we show in the Appendix that our main results hold in an alternative model in which firms can choose between distributions that are a mean-preserving spread.

The next step is to study the value of drawing productivity from more or less spread-out distributions. To simplify the notation, from now on we rewrite the entry problem of the firm as one of choosing directly v , rather than s . Substituting $A(\varphi^*)^{\sigma-1} = f$ into (4), assuming $\varphi^* > \varphi_{\min}$ (so that there is selection), $v < 1/\zeta$ (for $\mathbb{E}[\pi]$ to be finite) and using $G(\varphi)$, we can solve for expected profits as a function of v :²⁴

$$\mathbb{E}[\pi] = f \int_{\varphi^*}^{\infty} \left[\left(\frac{\varphi}{\varphi^*} \right)^{\sigma-1} - 1 \right] dG(\varphi) = \frac{f\zeta}{1/v - \zeta} \left(\frac{\varphi_{\min}}{\varphi^*} \right)^{1/v}, \quad (5)$$

where it proves convenient to define $\zeta \equiv \sigma - 1$. It is easy to see that expected *ex-ante* profits are increasing in v with elasticity equal to:

$$\frac{\partial \ln \mathbb{E}[\pi]}{\partial \ln v} = \frac{1}{1 - v\zeta} + \ln \left(\frac{\varphi^*}{\varphi_{\min}} \right)^{1/v} > 0. \quad (6)$$

There are three reasons why a higher v , and hence more dispersion in the distribution of productivity draws, implies higher expected profits. First, a higher v increases average productivity directly, by raising the mean of φ . Second, given the shape of the Pareto distribution, a higher v increases the probability of drawing a productivity above the exit cutoff φ^* .²⁵ Third, even in the absence of the previous effects, more dispersion

²⁴More precise parametric restrictions are stated later on.

²⁵An increase in v raises the density at any $\varphi > \varphi^*$. As shown below, this is true even holding constant

increases expected profits whenever the profit function is convex in prices and hence in φ . As equation (3) shows, this is the case when $\sigma > 2$ (i.e., for $\varsigma > 1$). To see this, suppose now that $\varphi_{\min} = \bar{\varphi} (1 - v)$ so that the mean of the distribution is constant at $\bar{\varphi}$ and an increase in v corresponds to a mean-preserving spread. Then:

$$\frac{\partial \ln \mathbb{E} [\pi]}{\partial \ln v} = \frac{1}{1 - v\varsigma} - \frac{1}{1 - v} + \ln \left(\frac{\varphi^*}{\varphi_{\min}} \right)^{1/v}$$

which is necessarily positive when $\varsigma > 1$ ($\sigma > 2$), even in the absence of selection effects (i.e., when $\varphi^* \rightarrow \varphi_{\min}$). The intuition is that firms can expand to take advantage of good realizations of productivity and contract to insure against bad realizations, making them potentially risk loving. This is a well-known result, sometimes referred to as the Hartman (1972) and Abel (1983) effect.

Having characterized the value of investing in larger and riskier projects, we now turn to its cost. In order to have a well-defined trade off, we assume that there are diminishing returns to investing in project size so that the total entry cost is a convex function of v . Formally, we denote the entry cost $\lambda F(v)$ with $F : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $F'(v) > 0$, $F''(v) > 0$ and $\lambda > 0$. We interpret the factor λ as a positive shifter parametrizing all the costs of financing the entry investment $F(v)$. Furthermore, to make sure that expected profits are finite, we also assume $\lim_{v \rightarrow \bar{v}} F'(v) = \infty$ with $\bar{v} < 1/\varsigma$.

We are now in the position to solve the entry stage. The problem is simplified by the fact that all firms entering a given sector are *ex-ante* identical and therefore face the same problem of choosing v so as to maximize expected profits minus the entry cost:

$$\max_{v \in [\underline{v}, \bar{v}]} \{ \mathbb{E} [\pi] - \lambda F(v) \},$$

where $\underline{v} \equiv \underline{s}/\kappa$. To ensure that the maximand is concave, we assume $\partial^2 \mathbb{E} [\pi] / \partial v^2 < F''(v)$. Then, the first-order condition for v is

$$\frac{\mathbb{E} [\pi]}{v} \left[\frac{1}{1 - v\varsigma} + \ln \left(\frac{\varphi^*}{\varphi_{\min}} \right)^{1/v} \right] = \lambda F'(v). \quad (7)$$

Concavity and implicit differentiation allow us to sign the comparative statics for v . The equilibrium choice of v is increasing in the elasticity of substitution, ς , average profit, $\mathbb{E} [\pi]$, and the exit cutoff, $\varphi^* / \varphi_{\min}$. However, both $\mathbb{E} [\pi]$ and $\varphi^* / \varphi_{\min}$ are endogenous and to solve for them we now turn to the industry equilibrium.

the mean of the Pareto distribution.

3.3 INDUSTRY EQUILIBRIUM

Free entry implies that *ex-ante* expected profits must be equal to the entry cost: $\mathbb{E}[\pi] = \lambda F(v)$. Substituting (5) into this condition, we can solve for the exit cutoff:

$$\left(\frac{\varphi^*}{\varphi_{\min}}\right)^{1/v} = \frac{f}{\lambda F(v)} \frac{\varsigma}{1/v - \varsigma}. \quad (8)$$

To make sure that $\varphi^*/\varphi_{\min} > 1$, we assume $f > \max_{v \in [\underline{v}, \bar{v}]} \{\lambda F(v) (1/\varsigma v - 1)\}$. Next, using $\mathbb{E}[\pi] = \lambda F(v)$ and (8), we can rewrite the first-order condition for v (7) as:

$$\ln\left(\frac{f}{\lambda F(v)} \frac{\varsigma}{1/v - \varsigma}\right) + \frac{1}{1 - v\varsigma} = \frac{vF'(v)}{F(v)} \quad (9)$$

Under regularity conditions that we take for granted, equation (9) has an interior solution over the relevant range $v \in [\underline{v}, \bar{v}]$.²⁶ Moreover, provided that $F(v)$ is sufficiently convex, the solution will be unique. Although the possibility of multiple equilibria is interesting, exploring it goes outside the scope of this paper and we therefore disregard this possibility.²⁷ As an illustration, we show in the Appendix functional-form assumptions yielding simple closed-form solutions and continue here with the more general case.

We can now study the equilibrium determinants of v . A higher fixed cost of production, f or a lower entry cost λ , increases the exit cutoff and hence raises the benefit of choosing a more dispersed distribution. A higher elasticity of substitution raises the value of v by making profits more convex in productivity and by increasing the exit cutoff. Interestingly, the choice of innovation size and hence v does *not* depend on the size of the market, captured by the parameter A . The reason is that a higher demand increases entry so as to keep expected profit per firm constant.

The choice of v affects the equilibrium distribution of firm characteristics. Consider the distribution of revenues, which matches closely the variable documented in Section 2. It is easy to show that revenues are a power function of productivity: $r(\varphi) = r(\varphi^*) (\varphi/\varphi^*)^\varsigma$. Then, from the properties of the Pareto distribution, $r(\varphi)$ is also Pareto distributed with c.d.f. $G_r(r) = 1 - (r_{\min}/r)^{1/v\varsigma}$, for $r > r_{\min} = \sigma f$.²⁸ Hence, the log of

²⁶Sufficient conditions are $F'(\underline{v}) = 0$ with $F(\underline{v}) > 0$ and $vF'(\bar{v})/F(\bar{v}) \rightarrow \infty$ with $\bar{v} < 1/\varsigma$.

²⁷One can imagine a situation in which the expectation of high profits in equilibrium induces firms to choose a high initial investment, which in turn confirms the initial expectation that firms be large. Similarly to Bonfiglioli and Gancia (2014), this multiplicity may help explain cross-country differences in the prevalence of small and large firms and other outcomes.

²⁸If φ follows a $\text{Pareto}(\varphi^*, z)$, then $x \equiv \ln(\varphi/\varphi^*)$ is distributed as an exponential with parameter z . Then, any power function of φ of the type $A\varphi^B$, with A and B constant, is distributed as a $\text{Pareto}(A(\varphi^*)^B, z/B)$, since $A\varphi^B = A(\varphi^*)^B e^{Bx}$ with $Bx \sim \text{Exp}(z/B)$, by the properties of the exponential

revenue is exponential with a standard deviation equal to $v\zeta$. This immediately implies that differences in the choice of entry risk across sectors will translate into differences in the equilibrium distributions of firm characteristics as summarized in the following Proposition.

Proposition 1 *Assume that the solution to (9) is unique and interior. Then, the equilibrium dispersion of firm productivity and revenue, as measured by the variance of the log of φ and $r(\varphi)$, is larger in sectors with a higher fixed cost, f , higher elasticity of substitution between varieties, ζ , and a lower entry cost as parametrized by λ .*

These results are broadly consistent with the empirical evidence documented in Section 2. As long as economic expansions are associated to lower entry costs, for instance by lowering financing costs, the effect of λ is compatible with the finding that faster economic growth correlates to a rise in dispersion. This prediction seems also consistent with the casual observation that in industries with very low entry barriers there are many startups, but only a few giants survives (e.g., Amazon, Google and Facebook, in the universe of “dot-com” companies). The model also reproduces the positive correlation between the variance and the mean of revenues observed in the data. To see this, note that average revenue, \bar{r} , is equal to $\sigma f / (1 - v\zeta)$, which is increasing in σ , f and v .

Although the model implies a positive correlation between firm heterogeneity and the elasticity of substitution, σ , this prediction is not straightforward to take to the data because it is difficult to find time-varying proxies of the latter parameter.²⁹ Note, however, that σ is an inverse measure of how fast marginal revenue falls with size, thereby also capturing the notion of “scalability”. Then, the model predicts firm heterogeneity to be higher in more scalable industries. Since labor-intensive firms are typically considered less scalable, this may explain why in Section 2 the coefficient for employment is negative and often significant.

Furthermore, it is possible to show that revenue-based labor productivity is also an increasing function of φ , and hence shares the same properties.³⁰ Finally, since the model is static, the results in Proposition (1) should be interpreted as capturing the long-run distribution. Nevertheless, given the high rate of product turnover observed in the data, we expect realistic transitional dynamics to be reasonably fast.

distribution.

²⁹Also, a high σ makes the restriction $\bar{v} < 1/(\sigma - 1)$ more binding.

³⁰In this model revenue per worker is an increasing function of φ because the fixed cost of production is in units of labor. The model in Section 5 generates variation in revenue per worker across firms through another channel.

4 TRADE AND EQUILIBRIUM FIRM HETEROGENEITY

We now extend the model by adding the possibility for firms to export their varieties subject to fixed and variable costs. This will lead to the familiar results that only the most productive firms export and that trade forces the least productive firms out. This *ex-post* reallocation of revenues will have new implications for the *ex-ante* entry stage: by increasing the payoffs in the tail, trade will induce firms to draw their productivity from more dispersed distributions.

Consider a world economy composed, for simplicity, of two symmetric countries. To serve the foreign market, firms must incur a fixed cost f_x in units of labor and an iceberg variable cost such that $\tau > 1$ units must be shipped for one unit to arrive at destination. The presence of a fixed trade cost implies that only the most productive firms choose to serve the foreign market. Formally, notice that, in analogy to (3), profits from exporting are $\pi_x(\varphi) = A(\varphi/\tau)^{\sigma-1} - f_x$. These profits would be negative for firms with productivity $\varphi < \varphi_x^* = \tau(f_x/A)^{1/\varsigma}$. As usual, we restrict attention to the space of parameters such that $\varphi_x^*/\varphi^* = \tau(f_x/f)^{1/\varsigma} > 1$, so that there is a range of firms with $\varphi \in [\varphi^*, \varphi_x^*]$ operating in the domestic market only, while the most productive firms also export.

Under these assumptions, *ex-ante* expected profits are:

$$\mathbb{E}[\pi] = f \int_{\varphi^*}^{\infty} \left[\left(\frac{\varphi}{\varphi^*} \right)^{\varsigma} - 1 \right] dG(\varphi) + f_x \int_{\varphi_x^*}^{\infty} \left[\left(\frac{\varphi}{\varphi_x^*} \right)^{\varsigma} - 1 \right] dG(\varphi), \quad (10)$$

where the two terms represent expected profits from the domestic and the foreign market. Solving the integrals yields:

$$\mathbb{E}[\pi] = \frac{\varsigma}{1/v - \varsigma} \left[f \left(\frac{\varphi_{\min}}{\varphi^*} \right)^{1/v} + f_x \left(\frac{\varphi_{\min}}{\varphi_x^*} \right)^{1/v} \right].$$

To study how export opportunities affect the value of drawing productivity from a riskier distribution, we compute again the elasticity of expected profits to v :

$$\frac{\partial \ln \mathbb{E}[\pi]}{\partial \ln v} = \frac{1}{1 - v\varsigma} + \ln \left(\frac{\varphi^*}{\varphi_{\min}} \right)^{1/v} + \frac{\ln(\varphi_x^*/\varphi^*)^{1/v}}{(\varphi_x^*/\varphi^*)^{1/v} f/f_x + 1} \quad (11)$$

Comparing this derivative to (6), we see that choosing a riskier distribution yields now a new advantage: conditional on surviving, it increases the probability of reaching the export cutoff, φ_x^* . Moreover, as it is well known and we show next, φ^*/φ_{\min} is higher with trade.

As in autarky, we solve for the equilibrium v by imposing the free-entry condition, $\mathbb{E}[\pi] = \lambda F(v)$. This condition allows us to find the exit cutoff:

$$\left(\frac{\varphi^*}{\varphi_{\min}}\right)^{1/v} = \frac{\varsigma}{1/v - \varsigma} \frac{f + f_x (\varphi_x^* / \varphi^*)^{-1/v}}{\lambda F(v)}. \quad (12)$$

As expected, the exit cutoff is higher than in autarky and is increasing in the barriers to export. For convenience, we now define $\rho \equiv \varphi^* / \varphi_x^* = (f/f_x)^{1/\varsigma} / \tau$ and use it as a synthetic measure of trade openness. This index, which varies between zero and one, only depends on exogenous parameters and determines the fraction of exporting firms, which is equal to $\rho^{1/v}$. Using this notation and (12) into (11), we can show how trade affects the elasticity of expected profits to v , and hence the incentive to draw productivity from riskier distribution:

$$\frac{\partial^2 \ln \mathbb{E}[\pi]}{\partial \ln v \partial \rho} = \frac{f/f_x}{\rho^{1+1/v}} \frac{\ln \rho^{-1/v}}{v (\rho^{-1/v} f/f_x + 1)^2} > 0. \quad (13)$$

In words, more openness raises *unambiguously* the return from riskier productivity draws. This result is intuitive: trade offers new profitable opportunities, but only to the most productive firms and hence reallocates profits to the right tail of the distribution. In turn, a higher v increases the probability mass in that tail. This is one of the main results of the paper: the chance of winning the extra prize of exporting induces firms to take a riskier bet at the entry stage.

Following the same steps as in autarky, the equilibrium v is implicitly determined by:

$$\frac{1}{1 - v\varsigma} + \ln \left(\frac{\varsigma}{1/v - \varsigma} \frac{f + f_x \rho^{1/v}}{\lambda F(v)} \right) + \frac{\ln \rho^{-1/v}}{\rho^{-1/v} f/f_x + 1} = \frac{vF'(v)}{F(v)} \quad (14)$$

Since the left-hand side is increasing in openness (this follows from 13), and assuming again the solution to be unique and interior, more openness leads to a higher equilibrium v and hence more productivity dispersion.

These results are summarized in the following Proposition.

Proposition 2 *An increase in openness triggered by a fall in the variable cost of trade, τ , induces firms to choose riskier productivity draws (higher v) and raises the equilibrium dispersion of firm productivity, as measured by the variance of the log of φ .*

An additional interesting implication of this model is that trade has a new effect on productivity. Since a higher v also raises the unconditional mean of φ , export opportunities induce firms to choose riskier technologies with higher expected returns. As a

result, in an equilibrium with trade firms will be more productive, not just because of the usual selection effect, but also because firms choose more costly, but on average more efficient, technologies. This prediction is also of intuitive appeal: the higher premium for success in the global economy makes firms more “ambitious” by choosing a bigger (and riskier) investment in the entry stage.

Of course, the analytical results derived in this section partly hinge on functional form assumptions and on the convenient properties of Pareto distributions. Yet, we expect the main mechanism to hold more in general. In particular, as long as trade reallocates profits in favor of exporters and exporting firms are in the upper tail of the distribution, trade will raise the return from taking a riskier bet on productivity.³¹

5 FROM FIRM HETEROGENEITY TO INCOME INEQUALITY

We now explore the implications of our theory for income and wage inequality. This is a natural step: the distribution of productivity is likely to be a major determinant of the distribution of wages because in the data more productive firms pay higher wages. We therefore extend the model to allow for differences in wages across firms. This will yield two main results: first, it will highlight a new channel through which trade can increase wage inequality and, second, it will identify some additional variables affecting the choice of dispersion at the entry stage.

In principle, our theory can be used to study top-income inequality. An immediate way of doing this is to draw a link between profits and entrepreneurial income. For example, one could assume that there is a class of agents, entrepreneurs, who are the only ones who can enter and start new firms. These agents may be able to finance part of the entry cost externally and will be the residual claimants on a share of profits. Recent models along these lines include Jones and Kim (2014) or Grossman and Helpman (2014). Since trade increases the dispersion of profits, it will also make entrepreneurial income more unequal. Several contributions in corporate finance, such as Gabaix and Landier (2008), have indeed shown that CEO compensations are proportional to firm size and that this can explain why they have increased so much in recent decades. Our theory can then help rationalize some of the changes in the firm size distribution behind this phenomenon.

Another possibility, that we consider more in detail, is to extend the model to study implications for wage dispersion. In the literature, there are several ways of linking firm productivity to wages. With competitive labor markets, wages can vary because

³¹Numerical exercises performed using realistically-calibrated log-normal distributions seem to confirm this intuition.

of differences in workforce composition across firms (e.g., Sampson, 2014, Monte, 2011, Yeaple, 2005). Alternatively, workers could be paid different wages due to labor market frictions (e.g., Helpman, Itskhoki and Redding, 2010, Amiti and Davis, 2012, Egger and Kreickemeier, 2009, Felbermayr, Impullitti and Prat, 2014). For example, in Helpman, Itskhoki and Redding (2010, HIR henceforth) workers matched randomly with heterogeneous firms draw a match-specific ability which is not observed and firms can invest in costly screening. In equilibrium, more productive firms screen workers more intensively to exclude those with lower ability. As a result, they have workforce of higher average ability and pay higher wages. These models yield an exporter wage premium and have been found to have considerable empirical support (e.g., Helpman et al. 2015). We therefore now borrow the framework of HIR to study the implications of our theory for wage dispersion. One key advantage of HIR is that it preserves the main equations of the basic Melitz model, thereby allowing us to apply our previous results in a relatively straightforward manner.

We briefly derive the equations of HIR that are relevant for our purpose and refer the reader to the original article for more details. For ease of comparison, we try to follow the original notation whenever possible. Production depends on the productivity of the firm, φ , the measure of hired workers, h , and the average ability of these workers, \bar{a} :

$$y = \varphi h^\gamma \bar{a},$$

where $\gamma \in (0, 1)$ implies diminishing returns to hired workers. Two important properties of this production function are the complementarity between firm productivity and average worker ability and a trade-off between the quantity and quality of hired workers. Workers' ability is assumed to be independently distributed and drawn from a Pareto distribution with shape parameter $k > 1$ and c.d.f. $G_a(a) = 1 - (a_{\min}/a)^{-k}$. Search frictions in the labor market imply that a firm has to pay bn units of the numeraire to be matched randomly with a measure n of workers. Ability is unknown. However, once the match is formed, the firm can use a screening technology to identify workers with ability below a_c at the cost of ca_c^δ/δ units of the numeraire, with $c > 0$ and $\delta > k$. Given the distribution of ability, a firm matched with n workers and screening at the cutoff a_c will hire a measure $h = n(a_{\min}/a_c)^k$ of workers with an average ability of $\bar{a} = a_c k / (k - 1)$. Following the notation in HIR, we define $\beta \equiv 1 - 1/\sigma$. Then, total revenue of a firm with productivity φ can be written as

$$r(\varphi) = (1 + \mathbb{I}^{1-\sigma})^{1-\beta} P X^{1-\beta} (\varphi \bar{a})^\beta h^{\beta\gamma}$$

where \mathbb{I} is an indicator function taking value 1 if the firm decides to export and zero otherwise.

Wages are determined through strategic bargaining between the firm and workers, after the firm has paid all the costs, which are now defined in units of the numeraire. HIR show that the outcome is that the firm retains a fraction of revenues equal to the Shapley value, $1/(1 + \beta\gamma)$, and pays the rest to the workers. Thus, the profit maximization problem of the firm is:

$$\pi(\varphi) = \max_{n, a_c, \mathbb{I}} \left\{ \frac{r(\varphi)}{1 + \beta\gamma} - bn - \frac{ca_c^\delta}{\delta} - f - \mathbb{I}f_x \right\},$$

and the first-order conditions for n and a_c are

$$\begin{aligned} \frac{\beta\gamma}{1 + \beta\gamma} r(\varphi) &= bn(\varphi) \\ \frac{\beta(1 - \gamma k)}{1 + \beta\gamma} r(\varphi) &= ba_c(\varphi)^\delta. \end{aligned}$$

Inspection reveals immediately that firms with higher revenue sample more workers (higher n) and screen more intensively (higher a_c). Assuming $\delta > k$ also ensures that firms with higher revenue hire more workers.

Substituting the first-order conditions for n and a_c into the profit function yields $\pi(\varphi) = \frac{\Gamma r(\varphi)}{1 + \beta\gamma} - f - \mathbb{I}f_x$, with $\Gamma \equiv 1 - \beta\gamma - (1 - \gamma k)\beta/\delta$. Since revenues are increasing in productivity, the fixed costs imply that firms with $\varphi < \varphi^*$ exit (where $\pi_{\mathbb{I}=0}(\varphi^*) = 0$) and firms with $\varphi > \varphi_x^*$ export (where $\pi_{\mathbb{I}=0}(\varphi_x^*) = \pi_{\mathbb{I}=1}(\varphi_x^*)$). Moreover, the relative revenue of any two firms only depends on their relative productivity and export status:

$$\frac{r(\varphi)}{r(\varphi^*)} = \left(1 + \mathbb{I}\tau^{1-\sigma}\right)^{(1-\beta)/\Gamma} \left(\frac{\varphi}{\varphi^*}\right)^{\beta/\Gamma}.$$

Combining these results, we find an expression for *ex-ante* expected profits, $\mathbb{E}[\pi]$, which turns out to be identical to the one in the previous section (equation 10) after the redefinition of the parameter $\varsigma \equiv \beta/\Gamma$ (instead of $\sigma - 1$). Openness is now

$$\rho = \frac{\varphi^*}{\varphi_x^*} = (f/f_x)^{1/\varsigma} \left[\left(1 + \tau^{1-\sigma}\right)^{\varsigma(1-\beta)/\beta} - 1 \right]^{1/\varsigma}.$$

The equilibrium v depends on ς , f and ρ as implied by equation (14) and, in particular, it is still increasing in ς . The difference, however, is that ς corresponds now to a combination of more parameters, $\varsigma = [\beta^{-1} - \gamma - (1 - \gamma k)/\delta]^{-1}$, so that in this ex-

tended version of the model there are more determinants of v . In particular, through their impact on ς , an increase in γ or a fall in k and δ leads firms to draw from more dispersed distributions. These results are intuitive. As already discussed, more risk taking is optimal for the firm when profits are more convex in productivity. In the simpler version of the model, convexity only depends on σ . Now, instead, the profit function is more convex also when there are weaker diminishing returns (high γ) and when screening - which is disproportionately beneficial to more productive firms - is more effective, i.e., when worker ability is more dispersed (low k) and the screening cost not too elastic (low δ).

Proposition 3 *The dispersion of firm productivity, as measured by the variance of the log of φ , is larger in sectors with more ability dispersion and weaker decreasing returns to scale.*

These results may also contribute at explaining why the dispersion of firm productivity varies across countries and over time. For example, it suggests that firms in countries with a more heterogeneous labor force will choose riskier investments and hence be more unequal in equilibrium. Likewise, the growing evidence on the “flattening of the firm” may indicate a rise in the span-of-control parameter and this may help explain the generalized increase in productivity dispersion documented in Section 2.

What are the implications for wages? Using the definition of wages as a share of revenue per hired worker yields:

$$w(\varphi) \equiv \frac{\beta\gamma}{1 + \beta\gamma} \frac{r(\varphi)}{h(\varphi)} = b \left[\frac{a_c(\varphi)}{a_{\min}} \right]^k.$$

Since $a_c(\varphi)$ is increasing in productivity, more productive firms pay higher wages. Due to the complementarity in production between average worker ability and productivity, more productive firms have a stronger incentive to be more selective, hire workers with higher ability and pay them higher wages. Moreover, since wages are proportional to revenue, which jumps at the export cutoff $\varphi = \varphi_x^*$, the model implies an exporter wage premium. More precisely, the wage paid by firms with productivity φ can be written as

$$w(\varphi) = \left(1 + \mathbb{I}\tau^{1-\sigma}\right)^{\frac{k(1-\beta)}{\delta\Gamma}} \varphi^{\frac{\beta k}{\delta\Gamma}} w(\varphi^*).$$

Finally, since employment, $h(\varphi)$, is also a power function of productivity, the wages of workers employed by domestic firms and exporters follow Pareto distributions with shape parameter:

$$1 + \delta[(v\varsigma)^{-1} - 1]/k,$$

which is decreasing in v . Thus, heterogeneity in productivity maps into wage dispersion. This allows us to state the following proposition on the impact of trade on wage inequality.

Proposition 4 *More openness raises unambiguously sectoral wage dispersion among workers employed by domestic firms and among workers employed by exporters. Conditional on not changing export status, more openness increases wage inequality between workers employed by any pair of firms with different productivity.*

Before concluding, it is important to highlight the qualitative and quantitative differences between our result and HIR. In HIR and some other existing models, trade affects wage dispersion through the exporter wage premium. The sign of the effect then depends on the fraction of exporters. As long as exporters are a minority, trade increases wage dispersion by raising the share of firms paying high wages. Once exporters are a majority, instead, trade decreases wage dispersion by pushing low-wage domestic firms to exit and making the surviving firms more equal. Thus, the overall effect of trade on inequality is inverted-U shaped. This effect is present also in our model. But there is now another, potentially more powerful, force: by making all firms more unequal, trade is changing the slope of the entire wage schedule. This second effect, which is absent in HIR, implies that trade now increases wage inequality within exporters, within non-exporters, and also between the two groups of firms.

6 A FURTHER LOOK AT THE EVIDENCE

Having derived the theoretical model, we now turn again to the data to test some of its predictions. We start by focusing on the implications for wage inequality. We show that, consistently with the theory, export opportunities make the distribution of wages significantly more spread out in a panel of U.S. industries. Next, we probe deeper into the mechanism linking trade, innovation and firm heterogeneity.

6.1 WAGE DISPERSION AND EXPORT OPPORTUNITIES

Our measure of wage dispersion is the standard deviation of log hourly wages. As detailed in the Appendix, we construct this variable using individual-level wage data from the CPS Merged Outgoing Rotation Groups for the years 1997, 2002 and 2007. For each year, we compute the standard deviation of log wages in 74 industries, defined according to the Census industry classification. We also use a second measure of wage dispersion, which captures *residual* rather than total wage inequality and thereby controls

for differences in observable worker characteristics. To construct this measure, we follow Helpman et al. (2015) and run Mincer wage regressions of log wages on a large set of covariates, separately for each year to allow for changes in the effects of these characteristics over time. Then, we use the residuals from these regressions to construct the standard deviation of log residual wages in each industry and year. The residual component explains the majority (60 per cent) of the variance of log wages in 2007 and essentially all (97 per cent) of its growth between 1997 and 2007.³² Besides its relevance, residual inequality is also more closely related to our model, in which wages do not depend on observable characteristics of workers.

The upper part of Table 6 reports descriptive statistics on the dispersion of wages and residual wages in the U.S. manufacturing sector. In the first four columns, we show the mean of both standard deviations across the 74 industries in 2007, the minimum and maximum value, and the average percentage change between 1997 and 2007. The standard deviation of log wages equals 0.54 and that of residual wages 0.43, implying that wage dispersion is smaller than sale dispersion. Similarly to sale dispersion, wage dispersion exhibits large variation across industries, with the maximum exceeding the minimum by more than 3 times for both standard deviations. Table 6 also confirms the substantial increase in U.S. wage inequality documented in previous studies. In particular, the standard deviation of log wages and residual wages rose by 12 and 16 per cent, respectively, between 1997 and 2007.

In the next two columns, we perform variance decomposition exercises analogous to those in Helpman et al. (2015). We decompose the total variance of log wages and residual wages into a within-industry and a between-industry component, and show the percentage contribution of the former to the level (in 2007) and growth (between 1997 and 2007) of both variances. These exercises confirm the well-known fact that wage inequality and its growth are not explained by dispersion of wages between industries but rather by wage dispersion within industries. The within-industry component accounts for 89 per cent of the variance of log wages and for 98 per cent of the variance of residual wages. Moreover, changes in wage inequality within industries explain essentially all of the increase in both variances over the sample period. While we cannot assign workers to firms using our data, recent studies show that most of the within-industry increase in wage inequality is due to rising inequality between firms operating in the same industry, as predicted by the model we use (Barth et al., 2014; Song et al., 2015).

In the bottom part of Table 6, we study how wage inequality changed between 1997

³²See Barth et al. (2014) for related evidence on residual wage inequality based on matched employer-employee data for the U.S. over 1977-2002.

Table 6: Descriptive Statistics on the Dispersion of Wages in U.S. Manufacturing

	Standard Deviation				Variance Decomposition: Within-Industry Contribution	
	Mean	Minimum	Maximum	% Change	Level	Change
Log wages	0.54	0.38	1.10	0.12	0.89	1.06
Log residual wages	0.43	0.26	1.10	0.16	0.98	1.01

	Standard Deviation (% Change)			
	Top 20	Top 40	Top 60	Top 80
Log wages	0.07	0.08	0.08	0.08
Log residual wages	0.13	0.08	0.09	0.09

Notes: Standard deviations are computed using worker-level data for 74 Census industries and are weighted by the number of full-time equivalent hours of labor supply. The sample consists of workers aged 18-64. Residual wages are obtained from yearly Mincer regressions of log hourly wages on log age, log age squared and dummies for race, gender, type of job, country of birth, educational level, union membership, full-time / part-time status, 3-digit occupations, Census industries and states. Statistics on the standard deviations are computed across the 74 industries. Mean, minimum and maximum refer to the year 2007; percentage changes refer to the period 1997-2007. The unreported between component of the variance decomposition is equal to 1 minus the within component. The bottom part of the table reports percentage changes in the standard deviations computed on the top 20, 40, 60 and 80 per cent of the wage distributions.

and 2007 across different parts of the wage distribution. To this purpose, we compute the average percentage change in each standard deviation after excluding observations in the left tail of the corresponding wage distribution. Moving from left to right in the table, we restrict to the top 20, 40, 60 and 80 per cent of the distribution.³³ These exercises show that the increase in wage inequality remains fairly similar when excluding increasingly larger shares of low-wage individuals. This suggests that the recent rise in U.S. wage inequality corresponds to a widening of the entire distribution, a result also found in Song et al. (2015).

Next, we study how the dispersion of wages relates to industry characteristics. To this purpose, we regress the standard deviation of log wages and residual wages on a large number of covariates, using the same sample of 6-digit NAICS industries employed in Section 2. We link the Census industries to 6-digit NAICS codes using correspondence tables from the U.S. Census Bureau. Because we attribute the wage dispersion of each Census industry to all 6-digit NAICS codes corresponding to it, we weight the regressions so as to give less weight to smaller industries; as weights, we use the 1997 shares of the 6-digit industries in total manufacturing employment.

The results are reported in Tables 7 and 8. The only difference between the two tables is that the former uses the standard deviation of log wages as the dependent vari-

³³An advantage of this approach is that if wages are Pareto distributed then we should obtain the same measure of dispersion, equal to the inverse of the shape parameter, on the truncated sub-samples.

Table 7: Dispersion of Wages and Industry Characteristics

	OLS						2SLS		
	Ind. FE	Ind. FE	Ind. FE	Ind. FE + FD	Ind. FE + FD	Pooled	Pooled	Ind. FE	Ind. FE + FD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
S.D. log establ. sales	0.044*** [0.012]								
Log establ. sales		0.070*** [0.014]	0.086*** [0.016]	0.052* [0.029]	0.051* [0.029]	0.004 [0.009]	0.007 [0.008]	0.080*** [0.016]	0.075** [0.036]
Log export intensity		0.037*** [0.011]	0.016** [0.007]	0.020** [0.009]	0.020** [0.009]	0.008*** [0.002]	0.014*** [0.004]	0.075** [0.032]	0.105*** [0.036]
Log exp. int. squared		0.002* [0.001]							
Log number of establ.			0.054* [0.029]	0.036 [0.054]	0.033 [0.052]	0.018* [0.009]	0.024** [0.010]	0.062* [0.036]	0.051 [0.063]
Log employment			-0.065*** [0.018]	0.002 [0.043]	0.004 [0.043]	-0.022** [0.009]	-0.026*** [0.009]	-0.035 [0.026]	0.014 [0.062]
Log capital intensity			-0.044 [0.045]	-0.040 [0.043]	-0.043 [0.043]	0.017 [0.018]	0.019 [0.018]	-0.026 [0.052]	-0.103 [0.065]
Log skill intensity			-0.006 [0.023]	-0.047** [0.023]	-0.050** [0.023]	-0.005 [0.009]	-0.005 [0.009]	-0.014 [0.025]	-0.044 [0.031]
Log material intensity			-0.081 [0.056]	-0.044 [0.087]	-0.051 [0.087]	0.029 [0.031]	0.030 [0.030]	-0.093 [0.067]	-0.193* [0.110]
Log average education			0.198* [0.111]	0.274 [0.190]	0.304 [0.195]	0.580*** [0.088]	0.524*** [0.086]	0.168 [0.122]	0.310 [0.254]
S.D. log education			-0.034 [0.107]	-0.048 [0.141]	-0.032 [0.139]	0.567*** [0.092]	0.524*** [0.093]	-0.024 [0.115]	-0.028 [0.165]
GDP growth			0.197 [0.120]			0.121 [0.102]	0.115 [0.102]	0.121 [0.135]	
Log import penetration					0.010 [0.010]				
Obs.	1,036	1,036	1,036	676	676	1,036	1,015	1,001	636
R ²	0.57	0.58	0.60	0.50	0.50	0.26	-	-	-
First-stage results									
Log world exports	-	-	-	-	-	-	0.535*** [0.113]	0.440*** [0.069]	0.520*** [0.084]
Kleibergen-Paap <i>F</i> -statistic	-	-	-	-	-	-	22.3	40.2	37.9

Notes: The dependent variable is the standard deviation of log hourly wages. It is computed using worker-level data for 74 Census industries. The sample consists of workers aged 18-64. The standard deviation is weighted by the number of full-time equivalent hours of labor supply. The 74 Census industries are mapped into the corresponding 6-digit NAICS industries using correspondence tables from the U.S. Census Bureau. All variables except for GDP growth are observed at the 6-digit NAICS industry level in the years 1997, 2002 and 2007. All industry-level controls are contemporaneous to the dependent variable. GDP growth is computed over the two years prior to each observation. The specifications in columns (1)-(6) are estimated with Ordinary Least Squares, those in columns (7)-(9) with Two-Stage Least Squares. Export intensity is instrumented with non-U.S. exports to the destination markets of the U.S. for each industry and year. All regressions are weighted with the shares of the 6-digit NAICS industries in total manufacturing employment in 1997. Standard errors are clustered by 6-digit industry and reported in brackets. *F*-statistics are reported for the Kleibergen-Paap test for weak instruments. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

able while the latter uses the standard deviation of log residual wages. Standard errors are clustered by industry. In column (1), we regress wage dispersion on the standard deviation of log establishment sales, controlling for industry fixed effects. The coefficients are positive, very precisely estimated and similar in size across the two tables. Consistently with our model, higher dispersion of sales is associated with greater wage inequality. This result also echoes recent firm-level evidence, according to which the dispersions of productivity and wages are positively correlated across U.S. firms (Barth et al., 2014). To have a sense of the magnitude of this correlation, note that our mea-

Table 8: Dispersion of Residual Wages and Industry Characteristics

	OLS						2SLS		
	Ind. FE	Ind. FE	Ind. FE	Ind. FE + FD	Ind. FE + FD	Pooled	Pooled	Ind. FE	Ind. FE + FD
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
S.D. log establ. sales	0.051*** [0.015]								
Log establ. sales		0.075*** [0.016]	0.106*** [0.020]	0.083* [0.043]	0.083* [0.043]	0.018* [0.010]	0.021** [0.009]	0.100*** [0.021]	0.102** [0.044]
Log export intensity		0.038*** [0.012]	0.022*** [0.008]	0.027** [0.011]	0.027** [0.011]	0.005*** [0.002]	0.012*** [0.004]	0.076** [0.035]	0.097** [0.041]
Log exp. int. squared		0.001 [0.001]							
Log number of establ.			0.078** [0.034]	0.099 [0.068]	0.099 [0.068]	0.029*** [0.011]	0.035*** [0.011]	0.085** [0.040]	0.112 [0.071]
Log employment			-0.075*** [0.022]	0.010 [0.057]	0.011 [0.057]	-0.031*** [0.011]	-0.036*** [0.011]	-0.048 [0.031]	0.021 [0.068]
Log capital intensity			-0.054 [0.053]	-0.063 [0.050]	-0.063 [0.051]	-0.010 [0.020]	-0.008 [0.020]	-0.038 [0.059]	-0.115 [0.071]
Log skill intensity			0.006 [0.029]	-0.032 [0.033]	-0.032 [0.034]	0.004 [0.010]	0.004 [0.010]	-0.001 [0.030]	-0.030 [0.037]
Log material intensity			-0.096 [0.065]	-0.052 [0.099]	-0.052 [0.100]	0.012 [0.027]	0.014 [0.026]	-0.106 [0.075]	-0.176 [0.123]
Log average education			-0.024 [0.143]	0.032 [0.238]	0.036 [0.247]	0.218** [0.085]	0.161** [0.082]	-0.050 [0.151]	0.062 [0.266]
S.D. log education			-0.201 [0.130]	-0.235 [0.164]	-0.233 [0.164]	0.224** [0.091]	0.183* [0.094]	-0.190 [0.138]	-0.217 [0.175]
GDP growth			0.224* [0.135]			0.139 [0.117]	0.135 [0.116]	0.157 [0.144]	
Log import penetration					0.001 [0.013]				
Obs.	1,036	1,036	1,036	676	676	1,036	1,015	1,001	636
R ²	0.41	0.42	0.44	0.52	0.52	0.07	-	-	-
First-stage results									
Log world exports	-	-	-	-	-	-	0.535*** [0.113]	0.440*** [0.069]	0.520*** [0.084]
Kleibergen-Paap <i>F</i> -statistic	-	-	-	-	-	-	22.3	40.2	37.9

Notes: The dependent variable is the standard deviation of log residual hourly wages. It is computed using worker-level data for 74 Census industries. The sample consists of workers aged 18-64. The standard deviation is weighted by the number of full-time equivalent hours of labor supply. Residual wages are obtained from yearly Mincer regressions of log hourly wages on log age, log age squared and dummies for race, gender, type of job, country of birth, educational level, union membership, full-time / part-time status, 3-digit occupations, Census industries and states. The 74 Census industries are mapped into the corresponding 6-digit NAICS industries using correspondence tables from the U.S. Census. All variables except for GDP growth are observed at the 6-digit NAICS industry level in the years 1997, 2002 and 2007. All industry-level controls are contemporaneous to the dependent variable. GDP growth is computed over the two years prior to each observation. The specifications in columns (1)-(6) are estimated with Ordinary Least Squares, those in columns (7)-(9) with Two-Stage Least Squares. Export intensity is instrumented with non-U.S. exports to the destination markets of the U.S. for each industry and year. All regressions are weighted with the shares of the 6-digit NAICS industries in total manufacturing employment in 1997. Standard errors are clustered by 6-digit industry and reported in brackets. *F*-statistics are reported for the Kleibergen-Paap test for weak instruments. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

sure of sale dispersion has a standard deviation of 0.58 while the dispersion of wages and residual wages have standard deviations of 0.09 and 0.10, respectively. Then, the estimated coefficients imply that a 1 s.d. increase in the dispersion of sales is associated with a 0.29 s.d. increase in both measures of wage dispersion.

In the next columns, we replace the dispersion of sales with its main determinants, to uncover new facts on the drivers of wage dispersion across U.S. industries. We are especially interested in the role of trade, since export opportunities should make the

distribution of wages more spread out according to our model. In column (2), we start by including log establishment sales and log export intensity. In this specification we also control for the square of log export intensity, to allow for possible non linearities in the relation between exports and wage inequality. Note that the dispersion of wages is increasing in average firm size. More importantly, wage inequality is strongly positively correlated with export intensity, consistently with our model. The quadratic term of export intensity is always small, positive and estimated with little precision, suggesting that in our data wage dispersion and export intensity do not display an inverted-U shaped pattern.

In column (3) we include the full set of control variables used in Section 2. Strikingly, wage dispersion covaries with industry characteristics in a way that mirrors remarkably well the pattern found for the dispersion of sales. In particular, the coefficient on average sales is always positive and significant. Moreover, wage inequality tends to correlate negatively with employment and positively with the number of establishments and GDP growth. Most importantly, the coefficients on export intensity remain positive and highly significant. In the following columns we show that this is still the case when we estimate the first-difference specification with industry fixed effects to account for industry-specific trends (column 4) and when we control for import penetration to account for the effect of foreign competition (column 5). The export intensity coefficients are also positive, albeit slightly smaller, when we estimate our preferred specification with pooled-OLS to fully exploit cross-industry variation (column 6).

Finally, in columns (7)-(9) we use Two-Stage Least Squares. As before, we instrument export intensity with non-U.S. countries' exports to the destination markets of the U.S..³⁴ The coefficient on export intensity is positive and precisely estimated across the board. In terms of magnitude, the specifications in levels with industry fixed effects (column 8 of Tables 7 and 8) imply that a 1 s.d. increase in log export intensity (1.34) would raise the dispersion of wages by 1.12 s.d. and the dispersion of residual wages by 1 s.d.. The observed increase in export intensity during the sample period (21 per cent) then accounts for 13 per cent of the growth in wage inequality and for 10 per cent of the increase in residual wage inequality. To conclude, the evidence presented so far paints a remarkably consistent picture, suggesting that the distribution of sales and wages behave very similarly in a panel of U.S. industries and that export opportunities have made both distributions more spread out.

³⁴Our second identification strategy - based on the interaction of bulk weight and oil price - is less promising in this case, because much of the variation in the bulk weight takes place within the 74 Census industries and does not contribute to identification.

6.2 TRADE, INNOVATION INTENSITY AND FIRM HETEROGENEITY

In this section we provide evidence on the mechanism that links export opportunities to the dispersion of sales. In our model, export opportunities raise expected profits in the tail of the productivity distribution and this induces firms to invest in technologies with higher costs and more dispersed outcomes. We now show that specific predictions of this mechanisms are that (i) more export opportunities lead to an increases in the share of revenue invested in innovation (total entry cost) and (ii) a higher investment in innovation is associated to more heterogeneity across firms at the sector level. We then take these predictions to the data.

Since aggregate profits must cover exactly the aggregate entry cost, we can express the share of revenue invested in innovation as:

$$i \equiv \frac{\bar{\pi}}{\bar{r}},$$

where $\bar{\pi}$ is average profit and \bar{r} average revenue made by active firms. We will use i as our measure of “innovation intensity”. To solve for it, note that free entry requires expected profits to be equal to the entry cost, which can be written as $\bar{\pi} = (\varphi^* / \varphi_{\min})^{1/v} \lambda F$, where $(\varphi_{\min} / \varphi^*)^{1/v}$ is the probability of successful entry. Next, recall that average profit is a fraction $1/\sigma$ of average revenue minus average fixed costs, which implies $\bar{r} = \sigma[\bar{\pi} + f + f_x (\varphi^* / \varphi_x)^{1/v}]$. Finally, combining these expressions for $\bar{\pi}$ and \bar{r} and using (12) yields

$$i = \frac{\varsigma v}{\varsigma + 1}. \quad (15)$$

Innovation intensity is a positive function of ς and v only. It follows that a fall in trade costs, τ , affects the share of revenue invested in innovation only through its effect on the dispersion of productivity, v . Since a fall in trade costs induces firms to choose technologies with a higher variance, the model predicts export opportunities to have a positive correlation with innovation intensity. But this correlation should be zero in a model in which v is exogenous. The same reasoning applies to a fall in entry costs, as parametrized by λ . Thus, investigating empirically the determinants of innovation intensity allows us to test a specific prediction of the model with endogenous heterogeneity in productivity.

To make progress, we follow Aghion et al. (2015) and switch from industrial to geographic data. In publicly available databases, innovation measures are typically aggregated into a small number of manufacturing industries, leaving us with insufficient degrees of freedom.³⁵ Using geographic data we can instead obtain a measure of inno-

³⁵For instance, the Bureau of Economic Analysis releases data on R&D investment for 13 broad manu-

vation for a wide and long panel of U.S. states. In particular, we proxy for innovation intensity, i , using the log number of granted utility patents, by year of application, per 1,000 workers in each U.S. state, using patent data sourced from Aghion et al. (2015).³⁶ To construct a state-level measure of export intensity, we instead follow Autor, Dorn and Hanson (2013) and compute the weighted average of export intensities across 6-digit NAICS industries, using as weights the industries' shares in the states' manufacturing employment by year. This measure of export intensity varies across states due to their different patterns of industrial specialization. A higher value of export intensity is thus associated with U.S. states specializing in more export-oriented industries.

We start our analysis by showing that, consistently with the model, innovation is strongly correlated with export opportunities. The results are reported in Table 9. To maximize sample size, we use data for all years between 1989 and 2007, the period for which we observe both innovation and export intensity. Since yearly changes in patent counts are typically noisy, we estimate all specifications in levels. In column (1) we start with a parsimonious regression of innovation on exports. We control for state fixed effects, so that the coefficient on export intensity is identified through within-state variation. Standard errors are clustered at the state level. The correlation is positive and very precisely estimated, with a t -statistic of 7.3. In column (2) we add a set of controls to account for other factors that may affect innovation and exports. In particular, we include the growth of U.S. GDP over the two years prior to each observation to proxy again for cyclical factors, including changes in entry costs. Moreover, to account for heterogeneous trends in other observable characteristics across states, we include two-year changes in the following variables: the number of bank deposits per capita; the shares of manufacturing and financial sector in each state's GDP; and the share of working-age population with at least a college degree in the state. The idea is that innovation intensity might increase more in states where access to credit and skilled labor is growing at a faster rate. Reassuringly, the inclusion of these controls leaves our main coefficient essentially unchanged.

Next, we re-estimate the last specification replacing the current value of export intensity with its first or fifth lag as, for some patents, the innovation process could have started in previous years triggered by past export opportunities. The results, reported in columns (3) and (4), are similar to our baseline specification and continue to highlight a strong positive association between innovation and exports. These regressions

facturing industries.

³⁶Utility patents are meant to protect new and useful innovations, and differ from design and plant patents which protect new product designs or new plants, respectively. Thus, utility patents are probably the best proxy for innovation. They cover 90 per cent of all patents registered at the US patent office (USPTO). See Aghion et al. (2015) for more discussion on this point.

Table 9: Determinants of Innovation

	OLS				2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log export intensity (t)	0.706*** [0.096]	0.693*** [0.098]			1.634*** [0.294]	1.769*** [0.303]		
Log export intensity (t-1)			0.639*** [0.092]				1.850*** [0.409]	
Log export intensity (t-5)				0.380*** [0.062]				1.335** [0.572]
GDP growth		-0.197 [0.268]	0.815*** [0.210]	0.351 [0.247]		-0.813** [0.398]	0.812* [0.436]	0.663 [0.413]
Growth in bank deposits per capita		0.111 [0.075]	0.097 [0.071]	-0.009 [0.069]		-0.071 [0.059]	-0.046 [0.058]	-0.123 [0.079]
Growth in manufacturing share of GDP		0.036 [0.098]	-0.005 [0.083]	-0.109* [0.057]		0.310* [0.185]	0.176 [0.172]	0.033 [0.130]
Growth in financial sector share of GDP		0.159 [0.160]	0.081 [0.155]	0.039 [0.142]		0.214 [0.179]	0.043 [0.166]	0.514 [0.335]
Growth in share of college-educ. persons		-0.007 [0.104]	-0.001 [0.123]	0.007 [0.102]		0.374** [0.146]	0.194 [0.158]	-0.176 [0.179]
Obs.	969	969	918	714	663	663	612	408
R ²	0.90	0.90	0.91	0.95	-	-	-	-
First-stage results								
Log world exports	-	-	-	-	0.097*** [0.018]	0.102*** [0.018]	0.105*** [0.024]	0.166*** [0.058]
Kleibergen-Paap F-statistic	-	-	-	-	29.3	30.8	19.7	8.2

Notes: The dependent variable is the log number of granted utility patents, by year of application, per 1000 workers in each U.S. state. All variables except for GDP growth are observed at the state level. Export intensity is obtained as the weighted average of export intensities across 6-digit NAICS industries; the weights are the industries' shares in each state's manufacturing employment in each year. All other controls are computed over the two years prior to each observation. The sample period is 1989-2007. In columns (5)-(8), export intensity is instrumented with non-U.S. exports to the destination markets of the U.S.. The state-level instrument, available since 1995, is constructed similarly to state-level export intensity, but the weights are given by the industries' shares in each state's manufacturing employment five years prior to each observation. All specifications include state fixed effects. Standard errors are clustered by state and reported in brackets. *F*-statistics are reported for the Kleibergen-Paap test for weak instruments. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

also yield a positive correlation between innovation and GDP growth, which might be consistent with the model if GDP growth captures lower entry costs.

In columns (5)-(8) we re-estimate the above specifications with Two-Stage Least Squares. To construct an instrument that varies across states and over time, we follow similar steps as in the construction of export intensity. In particular, we compute the weighted average of our industry-level instrument, using industries' employment shares as weights. Because innovation may affect the industrial composition of state employment, we use five-year lags rather than current values of the employment shares to mitigate endogeneity concerns. The resulting instrument attributes stronger demand shocks to states that are specialized in industries where non-U.S. exports to the U.S. destination markets grow faster.³⁷ The number of observations drops because the bilateral trade data used to construct the instrument are available since 1995. The first-stage results show

³⁷We have also computed state-specific bulk weights using a similar approach. This variable varies little across states, because industries with high bulk weights are active in most states with similar employment shares. We therefore lack sufficient cross-sectional variation to identify the effect of oil price shocks.

Table 10: Trade, Innovation and Sale Dispersion

	(1)	(2)	(3)	(4)	(5)	(6)
Log export intensity	1.017*** [0.249]		0.637* [0.349]	0.895*** [0.256]		0.501 [0.316]
Log number of patents per 1000 workers		0.644*** [0.128]	0.423** [0.192]		0.631*** [0.122]	0.462*** [0.170]
GDP growth				1.224* [0.726]	1.218* [0.670]	1.128 [0.676]
Growth in bank deposits per capita				-0.155 [0.268]	-0.092 [0.212]	-0.077 [0.205]
Growth in manufacturing share of GDP				0.258 [0.225]	0.414* [0.232]	0.407* [0.209]
Growth in financial sector share of GDP				-0.731 [0.503]	-0.711* [0.354]	-0.598 [0.417]
Growth in share of college-educated persons				-0.093 [0.322]	0.075 [0.349]	0.005 [0.328]
Obs.	153	153	153	153	153	153
R^2	0.32	0.33	0.39	0.42	0.46	0.49

Notes: The dependent variable is the standard deviation of log establishment sales in each state. This variable and export intensity are obtained as the weighted average of their counterparts across 6-digit NAICS industries; the weights are the industries' shares in each state's manufacturing employment in each year. All other controls are computed over the two years prior to each observation and, except for GDP growth, are observed at the state level. The sample period consists of the years 1997, 2002 and 2007. All specifications are estimated with Ordinary Least Squares and include state fixed effects. Standard errors are clustered by state and reported in brackets. ***, **, and * denote significance at 1, 5, and 10 per cent, respectively.

that the instrument is strongly correlated with export intensity also across states. At the same time, the second-stage estimates confirm that export intensity has a positive effect on innovation. The effect is not only statistically significant but also quantitatively large. Noting that export and innovation intensity have standard deviations of 0.3 and 0.68, respectively, the coefficient from our richest specification (column 6) implies that a 1 s.d. increase in export intensity leads to a substantial 0.78 s.d. increase in innovation intensity.

Finally, we take a further step to test the main mechanism highlighted in the paper, namely, that export opportunities make the distribution of sales more spread out by stimulating innovation. The results are reported in Table 10. We start by regressing the dispersion of sales at the state-level - our proxy for v , obtained by aggregating the industry-level dispersions as explained before - on export intensity. We report results from a parsimonious specification with only state fixed effects (column 1) and from a richer specification with additional controls (column 4). Both specifications are estimated with 153 observations, corresponding to the 51 states in 1997, 2002 and 2007. The results show that the dispersion of sales is strongly correlated with export intensity also across states. In columns (2) and (5) we regress the dispersion of sales on innovation intensity. Independent of the specification, we find a positive correlation as in

(15). Finally, in columns (3) and (6) we simultaneously control for innovation and export intensity. While the positive coefficient for innovation remains highly significant, the estimate for export intensity drops in both magnitude and precision, suggesting that export opportunities increase firm heterogeneity mostly through innovation.

7 CONCLUSIONS

In this paper, we made several contributions to the literature. First, we started documenting some little-known facts regarding how the distribution of firms varies across U.S. sectors and over time. We have found that the extent of heterogeneity, measured by the standard deviation of log sales, changes systematically with industry characteristics, especially export opportunities, and has increased significantly over time. Second, we have proposed one possible explanation, based on the idea that firms can affect the variance of their productivity at the entry stage. The model formalizes the hypothesis that firms can choose between larger and riskier innovation projects with high expected payoffs, and smaller but safer projects with lower expected returns. It shows that export opportunities, by reallocating profits to the most productive firms, increase the return to technological risk and induce firms to bet on more ambitious project with higher variance. Third, we have explored the implications for wage inequality and found a new channel through which trade liberalization can affect the entire wage distribution and increase its dispersion. Fourth, we have found evidence that the distribution of wages varies across U.S. industries and time in a way that mirrors remarkably well the distribution of firms' sales. Finally, we have used patent data for a panel of U.S. states to provide a first attempt at testing a key mechanism of our model, namely, that export opportunities increase firm heterogeneity by fostering innovation.

The analysis in this paper could be extended in several directions. To focus on one mechanism shaping the equilibrium distribution of firms and preserve tractability, we have abstracted from firm dynamics and *ex-post* innovation. Yet, our approach to the choice of technological risk could be applied to all innovation strategies of existing firms. Making the model dynamic, for instance along the lines of Arkolakis (2015) or Gabler and Poschke (2013), would also allow to study quantitative implications and transitional adjustments. Within our theory, we also restricted the attention to positive implications. Yet, the model suggests interesting normative questions: do firms take too much risk, especially if workers are risk averse and insurance markets are imperfect? Does international trade introduce new externalities in the technology choice at the entry stage? The answers to these questions may be important for quantifying the welfare implications of trade liberalization (Antràs, de Gortari and Itskhoki, 2015). Finally, our look at the

data has uncovered a number of new findings, but much remains to be done to deepen our understanding of how and why productivity and wages vary across firms, sectors and time.

REFERENCES

- [1] Abel, Andrew (1983). "Optimal Investment under Uncertainty." *American Economic Review*, 73: 228–233.
- [2] Acemoglu, Daron and Dan Cao (2015). "Innovation by Entrants and Incumbents." *Journal of Economic Theory*, 157: 255–294.
- [3] Aghion, Philippe, Ufuk Akcigit, Antonin Bergeaud, Richard Blundell, David Hémous (2015). "Innovation and Top Income Inequality." NBER Working Paper No. 21247.
- [4] Antràs, Pol, Alonso de Gortari, and Oleg Itskhoki (2015). "Inequality, Costly Redistribution and Welfare in an Open Economy." Working Paper.
- [5] Amiti, Mary and Donald Davis (2012). "Firms, Trade, and Wages: Theory and Evidence." *Review of Economic Studies*, 79: 1–36.
- [6] Arkolakis, Costas (2015). "A Unified Theory of Firm Selection and Growth." NBER Working Paper No. 17553.
- [7] Atkeson, Andrew and Ariel Burstein (2010). "Innovation, Firm Dynamics and International Trade." *Journal of Political Economy*, 118: 433–484.
- [8] Autor, David H., David Dorn and Gordon H. Hanson (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review*, 103: 2121–2168.
- [9] Autor, David H., Frank Levy and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*, 118: 1279–1333.
- [10] Axtell, Robert L. (2001). "Zipf Distribution of U. S. Firm Sizes." *Science*, 293: 1818–1820.
- [11] Bartelsman, Eric J., Pieter Gautier and Joris de Wind (2015). "Employment Protection, Technology Choice, and Worker Allocation." *International Economic Review*, forthcoming.
- [12] Bartelsman, Eric, John Haltiwanger and Stefano Scarpetta (2009). "Measuring and Analyzing Cross Country Differences in Firm Dynamics." in *Producer Dynamics: New Evidence from Micro Data* (Dunne, Jensen and Roberts, eds.), University of Chicago Press.
- [13] Barth, Erling, Alex Bryson, James C. Davis and Richard Freeman (2014). "It's Where

You Work: Increases in Earnings Dispersion across Establishments and Individuals in the U.S.,” NBER Working Paper No. 20447.

- [14] Bernard, Andrew, Bradford Jensen, Stephen J. Redding and Peter K. Schott (2012). “The Empirics of Firm Heterogeneity and International Trade.” *Annual Review of Economics*, 4: 283-313.
- [15] Bernard, Andrew, Stephen J. Redding and Peter K. Schott (2011). “Multi-Product Firms and Trade Liberalization.” *Quarterly Journal of Economics*, 126: 1271–1318.
- [16] Bonfiglioli, Alessandra and Gino Gancia (2014). “Heterogeneity, Selection and Labor Market Disparities.” CEPR Discussion Paper No. 9981.
- [17] Broda, Christian and David E. Weinstein (2010). “Product Creation and Destruction: Evidence and Price Implications.” *American Economic Review*, 100: 691–723.
- [18] Bustos, Paula (2011). “Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms.” *American Economic Review*, 101: 304-340.
- [19] di Giovanni, Julian and Andrei Levchenko (2012). “Country Size, International Trade, and Aggregate Fluctuations in Granular Economies.” *Journal of Political Economy*, 120: 1083-1132.
- [20] Dhingra, Swati (2013). “Trading Away Wide Brands for Cheap Brands.” *American Economic Review*, 103: 2554–2584.
- [21] Doraszelski, Ulrich and Jordi Jaumandreu (2013). “R&D and Productivity: Estimating Endogenous Productivity.” *Review of Economic Studies*, 80: 1338-1383.
- [22] Dunne, Timothy, Lucia Foster, John Haltiwanger and Kenneth Troske (2004). “Wage and Productivity Dispersion in US Manufacturing: The Role of Computer Investments.” *Journal of Labor Economics*, 22: 397-429.
- [23] Egger, Hartmut and Udo Kreickemeier (2009). “Firm Heterogeneity and the Labor Market Effects of Trade Liberalization.” *International Economic Review*, 50: 187–216.
- [24] Faggio, Giulia, Kjell Salvanes and John Van Reenen (2010). “The Evolution of Inequality in Productivity and Wages: Panel Data Evidence.” *Industrial and Corporate Change*, 19: 1919-1951.
- [25] Felbermayr, Gabriel, Giammario Impullitti and Julien Prat (2014). “Firm Dynamics and Residual Inequality in Open Economies.” Working Paper, University of Nottingham.
- [26] Fillat, José L. and Stefania Garetto (2015). “Risk, Returns, and Multinational Production.” *Quarterly Journal of Economics*, forthcoming.
- [27] Gabaix, Xavier and Augustin Landier (2008). “Why Has CEO Pay Increased So Much?” *Quarterly Journal of Economics*, 123: 49-100.

- [28] Gabler, Alain and Markus Poschke (2013). "Experimentation by Firms, Distortions, and Aggregate Productivity." *Review of Economic Dynamics*, 16: 26-38.
- [29] Grossman, Gene and Elhanan Helpman (2014). "Growth, Trade, and Inequality." Working Paper, Harvard University.
- [30] Hartman, Richard (1972). "The Effects of Price and Cost Uncertainty on Investment." *Journal of Economic Theory* 5: 258-266.
- [31] Head, Keith, Thierry Mayer and Mathias Thoenig (2014). "Welfare and Trade without Pareto." *American Economic Review*, 104: 310-316.
- [32] Helpman, Elhanan, Oleg Itskhoki, Marc-Andreas Muendler and Stephen J. Redding (2015). "Trade and Inequality: From Theory to Estimation." Working Paper, Harvard University.
- [33] Helpman, Elhanan, Oleg Itskhoki and Stephen J. Redding (2010). "Inequality and Unemployment in a Global Economy." *Econometrica*, 78: 1239-1283.
- [34] Helpman, Elhanan, Marc J. Melitz and Stephen R. Yeaple (2004). "Export Versus FDI with Heterogeneous Firms." *American Economic Review*, 94: 300-316.
- [35] Hopenhayn, Hugo A. (2014). "Firms, Misallocation, and Aggregate Productivity: A Review." *Annual Review of Economics*, 6: 735-770.
- [36] Hummels, David (2007). "Transportation Costs and International Trade in the Second Era of Globalization." *Journal of Economic Literature*, 21: 131-154.
- [37] Hummels, David, Rasmus Jørgensen, Jakob Munch and Chong Xiang (2014). "The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data." *American Economic Review*, 104: 1597-1629.
- [38] Jones, Charles I. and Jihee Kim (2014). "A Schumpeterian Model of Top Income Inequality." Working Paper, Stanford University.
- [39] König, Michael D., Jan Lorenz and Fabrizio Zilibotti (2015). "Innovation vs. Imitation and the Evolution of Productivity Distributions." *Theoretical Economics* forthcoming.
- [40] Lileeva, Alla and Daniel Trefler (2010). "Improved Access to Foreign Markets Raises Plant-Level Productivity... for Some Plants." *Quarterly Journal of Economics*, 125: 1051-1099.
- [41] Luttmer, Erzo G.J. (2010). "Models of Growth and Firm Heterogeneity." *Annual Reviews of Economics*, 2: 547-576.
- [42] Mayer, Thierry, Marc J. Melitz and Gianmarco Ottaviano (2015). "Product Mix and Firm Productivity Responses to Trade Competition." Working Paper, Harvard University.
- [43] Melitz, Marc J. (2003). "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica*, 71: 1695-1725.

- [44] Melitz, Marc J. and Stephen J. Redding (2014). "Heterogeneous Firms and Trade." *Handbook of International Economics* (Gopinath, Helpman and Rogoff, eds.), North Holland.
- [45] Monte, Ferdinando (2011). "Skill Bias, Trade, and Wage Dispersion." *Journal of International Economics*, 83: 202–218.
- [46] Mrazova, Monika, J. Peter Neary and Mathieu Parenti (2015). "Technology, Demand and the Size Distribution of Firms." Working Paper, University of Oxford.
- [47] Poschke, Markus (2014). "The Firm Size Distribution across Countries and Skill-Biased Change in Entrepreneurial Technology." Working Paper, McGill University.
- [48] Romalis, John (2004). "Factor Proportions and the Structure of Commodity Trade." *American Economic Review*, 94: 67-97.
- [49] Rossi-Hansberg, Esteban and Mark L. J. Wright (2007). "Establishment Size Dynamics in the Aggregate Economy." *American Economic Review*, 97: 1639-1666.
- [50] Sampson, Thomas (2014). "Selection into Trade and Wage Inequality." *American Economic Journal: Microeconomics*, 6: 157-202.
- [51] Schott, Peter (2008). "The Relative Sophistication of Chinese Exports." *Economic Policy*, 53: 5-49.
- [52] Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom and Till von Wachter (2015). "Firming Up Inequality." NBER Working Paper No. 21199.
- [53] Syverson, Chad (2004a). "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112: 1181-1222.
- [54] Syverson, Chad (2004b). "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics*, 86: 534-550.
- [55] Syverson, Chad (2011). "What Determines Productivity." *Journal of Economic Literature*, 49: 326–365.
- [56] Vannoorenberghe, Gonzague (2014). "International Trade, Risk Taking and Welfare." *Journal of International Economics*, 92: 363-374.
- [57] Yeaple, Stephen R. (2005). "A Simple Model of Firm Heterogeneity, International Trade, and Wages." *Journal of International Economics*, 65: 1–20.

8 APPENDIX

8.1 DATA APPENDIX

Here we provide details on data sources and variables definitions.

Dispersion of sales and labor productivity The dispersion of sales and labor productivity are computed with Census data drawn from the ‘Statistics of U.S. Businesses’. For Census years 1997, 2002 and 2007, the database contains receipts of sales and the number of firms, establishments and employees, disaggregated by industry and sales-size category. The publicly available data aggregate confidential establishment-level information from the Business Registry, which covers the universe of establishments with paid employees in the U.S.. The number of sales-size bins is equal to 10 for 1997, 8 for 2002 and 18 for 2007. The lowest bin contains firms with revenue smaller than 50 thousand US\$ while the highest bin contains firms with revenue larger than 100 million US\$. In our main analysis, we aggregate the data into six bins consistently defined over the sample period. They refer to firms with US\$ revenue of: (1) less than 100,000; (2) 100,000-499,999; (3) 500,000-999,999; (4) 1,000,000-4,999,999; (5) 5,000,000-99,999,999; (6) 100 million or more. As for the industrial breakdown, the 1997 data are reported at the 4-digit level of the SIC-87 classification, while the 2002 and 2007 data are reported at the 6-digit level of the 1997 and 2002 NAICS classifications, respectively. We map the original data into 453 6-digit NAICS industries consistently defined over the three years. To this purpose, we use crosswalks provided by the NBER (between SIC-87 and NAICS-97) and the U.S. Census Bureau (between different versions of the NAICS classification). Finally, we compute the dispersion of sales in each industry and year as the standard deviation of log average sales per establishment across the six bins, weighting the observations by the number of establishments in each bin. The standard deviation of log labor productivity (sales per worker) is computed analogously, using the number of firms in each bin as weights.

Dispersion of wages and educational attainment Our measures of wage dispersion are the standard deviations of log hourly wages and log hourly residual wages. We construct these measures with individual-level data from the CPS Merged Outgoing Rotation Groups (CPS-MORG). We focus on working-age individuals (18-64 years old) and compute hourly wages as weekly earnings divided by the usual number of hours worked per week. Residual wages are the residuals from yearly Mincer regressions of log wages on log age, log age squared and dummies for race, gender, type of job, country of birth, educational level, union membership, full-time / part-time status, 3-digit occupations,

Census industries and states. With this data in hand, we compute the yearly standard deviations of log wages and log residual wages for 74 Census industries, consistently defined over the sample period; both standard deviations are weighted with full-time equivalent hours of labor supply as in Autor, Levy and Murnane (2003). We map the 74 Census industries into the 6-digit NAICS industries using correspondence tables from the U.S. Census Bureau. We proceed similarly to construct the mean and standard deviation of log educational attainment. Education is a discrete variable ranging from 1 to 16, with higher values corresponding to more advanced degrees.

Export intensity and import penetration Export intensity is the ratio of exports to total shipments in a given industry and year. Shipment data come from the NBER-CES Manufacturing Industry Productivity Database. Export data are sourced from Schott (2008). For the years 1989-2005, the data are provided by 6-digit NAICS industry and destination country. For later years, they are instead disaggregated by destination and product (10-digit Harmonized System classification), so we compute industry-level bilateral exports by summing exports across all products belonging to a given industry and exported to a given destination. Import penetration is the ratio of c.i.f. imports to apparent consumption, defined as production plus imports minus exports. Import data come from Schott (2008) and are constructed analogously to the export data. After excluding 144 observations with exports greater than total shipments, our final data set contains export intensity and import penetration for 375 6-digit NAICS industries over 1989-2007.

Instrument We instrument export intensity using non-U.S. countries' exports to the destination markets of the U.S. in each industry and year. To construct the instrument we use data on bilateral exports from BACI, available since 1995 at the HS 6-digit product level. We convert these data into 6-digit NAICS industries using correspondence tables from the World Integrated Trade Solution and the U.S. Census Bureau.

Bulk weight and oil price We construct industry-level bulk weights (expressed in Kg per US\$) as the export-weighted average of individual products' bulk weights across air and vessel shipments in 1995. We source data from Schott (2008). We combine the bulk weights with data on oil price (Brent) from FRED (Federal Reserve of St. Louis).

Patents To construct our measure of innovation intensity we use the number of granted utility patents, by year of application, in each U.S. state and year. The original data are sourced from Aghion et al. (2015) and are available for the period 1970-2009. The

data contain information on all patents registered at the US Patent Office (USPTO); each patent is located to the state where its inventor works or live. We normalize patent counts by total employment in each state and year, sourced from the County Business Patterns (CBS).

Other variables The number of establishments and workers in each industry and year come from the ‘Statistics of U.S. Businesses’. Factor intensities are computed as in Romalis (2004) using data from the NBER-CES Manufacturing Industry Productivity Database. Material intensity is equal to material costs divided by material costs plus value added. Capital intensity is computed as 1 minus material intensity, times the non-labor share of value added. Skill intensity is computed as 1 minus material intensity, times the product between the labor share of value added and the employment share of non-production workers.

We construct export intensity, sale dispersion and the instrument for each U.S. state and year as the weighted average of these variables across 6-digit NAICS industries. The weights are given by the industries’ shares in each state’s manufacturing employment by year. We obtain yearly information on the industrial composition of state employment from the CBS. For some state-industry-year cells, employment is not disclosed for confidentiality reasons. To estimate employment in these cells, we use information on the number of establishments in nine employment-size bins, available for each industry and year from the CBS. We first estimate employment in each bin as the product between the number of establishments in the bin and the mid-point of the bin itself. Then, we sum the results across bins to obtain an estimate of total employment in the state-industry-year cell.

Data on state population and GDP by sector come from the regional accounts of the BEA. The number of bank deposits in each state and year is instead sourced from the ‘USA Counties Database’. Finally, the share of working-age population with at least a college degree in each state and year is constructed with educational attainment data from the CPS-MORG, weighting the individual observations with full-time equivalent hours of labor supply.

8.2 A SIMPLE CASE WITH CLOSED-FORM SOLUTIONS

To guarantee $v < 1/\zeta$, assume that $\kappa = \alpha\zeta$ with $\alpha > 1$ and $s \in (0, 1]$. The first assumption implies that quality of potential ideas is more dispersed in industry producing more differentiated varieties. This is intuitive: there is little scope for technological differences between very homogeneous products. The second assumption normalizes

the size of the largest project to one. Then, $v = s/(\alpha\zeta)$. The parameter $\alpha > 1$ pins down the upper bound to v : $\bar{v} = 1/(\alpha\zeta)$. Next, assume that the entry cost is proportional to $s/(1-s)$, which is increasing and tends to infinity as s approaches the maximum size. Thus, $F(v) = v\alpha\zeta/(1-v\alpha\zeta)$. Substituting $F(v)$ and $v = s/(\alpha\zeta)$ into (9) yields:

$$(1-s) \ln \left(\frac{f}{\lambda} \frac{1-s}{\alpha-s} \right) + \alpha \frac{1-s}{\alpha-s} = \frac{1}{\alpha\zeta}. \quad (16)$$

Note that $f/\lambda > \alpha$ is required to have $\varphi^* > \varphi_{\min}$. Assuming

$$\frac{\alpha\lambda}{f} < \frac{1}{\alpha\zeta} < \ln \left(\frac{f}{\lambda\alpha} \right) + 1,$$

guarantees that (16) has a unique interior solution for s . Since $v = s/(\alpha\zeta)$, the log of revenue has a standard deviation equal to s/α . In the limit case $\alpha \rightarrow 1$, which however requires the restrictive condition $\sigma < 2$, (16) collapses to $s = 1 - \frac{1-\zeta}{\zeta \ln(f/\lambda)}$.

8.3 MEAN-PRESERVING SPREADS

We now solve the model under the assumption that $\varphi_{\min} = \bar{\varphi}(1-v)$ so that the mean of the distribution is constant at $\bar{\varphi}$ and an increase in v corresponds to a mean-preserving spread. Although the model in the main text suggests that the mean and the variance of productivity are likely to be linked, our goal here is to show that the main results in the paper do not change if firms can only choose the variance, and hence the riskiness, of their initial draw.

Using the expression for $\frac{\partial \ln \mathbb{E}[\pi]}{\partial \ln v}$ derived in the text, the first-order condition for the problem $\max_{v \in [\underline{v}, \bar{v}]} \{ \mathbb{E}[\pi] - \lambda F(v) \}$ becomes:

$$\frac{\partial \mathbb{E}[\pi]}{\partial v} = \frac{\mathbb{E}[\pi]}{v} \left[\frac{1}{1-v\zeta} - \frac{1}{1-v} + \ln \left(\frac{\varphi^*}{\varphi_{\min}} \right)^{1/v} \right] = \lambda F'(v).$$

Imposing free-entry, $\mathbb{E}[\pi] = \lambda F(v)$, yields the same expression for the exit cutoff, φ^*/φ_{\min} . Using these results in the new first-order condition yields:

$$\frac{1}{1-v\zeta} - \frac{1}{1-v} + \ln \left(\frac{f}{\lambda F(v)} \frac{\zeta}{1/v - \zeta} \right) = \frac{vF'(v)}{F(v)}.$$

This equation, which pins down implicitly the equilibrium v , is identical to (9), except for the new term $-1/(1-v)$ on the left-hand side. Intuitively, the fact that higher risk is not associated to higher expected productivity draw lowers the value of v . This imme-

diately implies that firms will choose a lower equilibrium level of entry risk. However, the comparative static for all the parameters are unchanged. Moreover, revenue, $r(\varphi)$, is still Pareto distributed with c.d.f. $G_r(r) = 1 - (r_{\min}/r)^{1/v\varsigma}$, for $r > r_{\min} = \sigma f$ as in the benchmark case. Hence, all the results in Proposition 1 still hold.

Consider now the model with trade. Deriving expected profit (10) with respect to v when $\varphi_{\min} = \bar{\varphi}(1 - v)$ yields:

$$\frac{\partial \ln \mathbb{E}[\pi]}{\partial \ln v} = \frac{1}{1 - v\varsigma} - \frac{1}{1 - v} + \ln \left(\frac{\varphi^*}{\varphi_{\min}} \right)^{1/v} + \frac{\ln(\varphi_x^*/\varphi^*)^{1/v}}{(\varphi_x^*/\varphi^*)^{1/v} f/f_x + 1}.$$

This expression for the return to risk is again, is identical to (11), except for the new term $-1/(1 - v)$ on the left-hand side.

Imposing free-entry, $\mathbb{E}[\pi] = \lambda F(v)$, yields the same expression (12) for the exit cut-off, φ^*/φ_{\min} . Then, the effect of openness, ρ , on the value of risk, as captured by $\frac{\partial^2 \ln \mathbb{E}[\pi]}{\partial \ln v \partial \rho}$ is identical to (13). Following the same steps as in autarky, the equilibrium v is implicitly determined by:

$$\frac{1}{1 - v\varsigma} - \frac{1}{1 - v} + \ln \left(\frac{\varsigma}{1/v - \varsigma} \frac{f + f_x \rho^{1/v}}{\lambda F(v)} \right) + \frac{\ln \rho^{-1/v}}{\rho^{-1/v} f/f_x + 1} = \frac{vF'(v)}{F(v)}.$$

Once again, the lower value of risk (the new term $-1/(1 - v)$ on the left-hand side) implies the firms will draw from less dispersed distributions, but the effect of openness and other parameters on the choice of v is qualitatively unchanged. However, the result that trade induces firms to choose technologies with higher expected productivity is now absent.