

Essays on Firms in Developing Countries

David Alfaro Serrano

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2020

© 2020

David Alfaro Serrano

All Rights Reserved

Abstract

Essays on Firms in Developing Countries

David Alfaro Serrano

Understanding firm behavior is key to understand the process of economic development. Firm choices affect labor market outcomes and the economy's ability to increase productivity and living standards.

In this dissertation, I study two important aspects of firm behavior: technological upgrading and exporting. In the first chapter, I analyze the role of adoption costs and technological complementarities in the process of managerial upgrading, and propose a feasible way to promote the adoption of better management practices by firms. Using a regression discontinuity strategy, I show that a subsidy to certify process standards, such as ISO 9001, increases certification probability and, additionally, induces the adoption of modern management practices that are beyond the standards' scope. The managerial improvement is concentrated in monitoring and target-setting practices, while no change is detected in practices related to incentives for employees. These findings are consistent with a model in which process documentation, which is required by the standards, and modern management practices are complementary and suggest that subsidizing the certification of process standards is a feasible way to improve management.

While the first chapter focuses on the adoption of an already known technology, the second chapter is concerned with the capacity of R&D subsidies to induce the adoption of new technologies in companies. Despite their popularity, there is little evidence of the effect R&D subsidies on the adoption of new technologies by companies. Using a regression discontinuity strategy, I

show that an R&D subsidy program in Peru was not able to induce the adoption of new products and processes by beneficiary firms. Qualitative evidence suggests that the main obstacles were not the technical challenges of developing the new technologies, but their implementation. Together with the results presented in the first chapter, these findings suggest that firms' lack of capacity to handle complex projects might be an important barrier for the success of policy interventions to promote technological upgrading.

In the third chapter, co-authored with Judith A. Frías, David S. Kaplan, and Eric Verhoogen, we explore the impact of exports on wage premia. There is evidence showing that exporting firms pay higher average wages. However, it is still unclear whether these results are due to changes in the wage premia or changes in workforce composition. In our study, we use employer-employee and longitudinal plant data from Mexico to address this question. We do so by decomposing plant-level average wages into a component reflecting wage premia and a component reflecting workers' skill composition. Using the late-1994 peso devaluation interacted with initial plant size as a source of exogenous variation in exports, we find that exports have a significant positive effect on wage premia, and that the effect on wage premia accounts for essentially all of the medium-term effect of exporting on plant-average wages.

Table of Contents

List of Tables	iv
List of Figures	vii
Acknowledgments	ix
Dedication	x
Chapter 1: Process Standards and Management Practices: Evidence from Peru	1
1.1 Introduction	1
1.2 Literature review	5
1.2.1 Management and productivity	5
1.2.2 Process standards	9
1.3 Model	11
1.3.1 Setup	11
1.3.2 Exploiting complementarity to promote upgrading	13
1.4 Empirical strategy	15
1.4.1 The Certification Program and the identification strategy	15
1.4.2 Data sources	18
1.4.3 Validity of the identification strategy	20

1.5	Effects of the certification program	21
1.5.1	Effects on certification and management practices	21
1.5.2	Other changes within the firm	28
1.5.3	Effects on productivity and performance	29
1.6	Conclusion	36
Chapter 2: R&D Subsidies and Technology Adoption: Evidence from Peru		39
2.1	Introduction	39
2.2	Empirical strategy	41
2.2.1	The R&D subsidy program and the identification strategy	41
2.2.2	Data sources	44
2.2.3	Validity of the identification strategy	44
2.3	Effects of the R&D subsidy program	46
2.3.1	Effects on technology adoption	46
2.3.2	Effects on productivity and performance	50
2.4	Conclusion	53
Chapter 3: Exports and Wage Premia: Evidence from Mexican Employer-Employee Data		56
3.1	Introduction	56
3.2	Econometric Strategy	62
3.3	Data and Background	69
3.4	Estimating Wage Premia	76
3.5	Estimating the Effect of Exporting on Wage Premia	81
3.6	Conclusion	89

References	106
Appendix A: Additional information for chapter 1	107
A.1 Certification Program sample	107
A.2 Management index and complementary data sources	108
A.3 Balance tests using honest confidence intervals	112
A.4 Effect on employment and sales	113
A.5 Robustness of results	115
A.5.1 Results using the restricted sample	115
A.5.2 Results using honest confidence intervals	122
A.6 Effect on management index and subindices controlling for productivity	123
A.7 Changes in recent years	126
A.7.1 Questions per category of change	126
A.7.2 Results for other changes	128
A.8 Exploiting complementarity under financial constraints	133
A.9 Proof of supermodularity	136
Appendix B: Additional information for chapter 2	137
B.1 R&D program sample	137
Appendix C: Additional information for chapter 3	138
C.1 IMSS individual-level data	138
C.2 EIA plant-level data	139

List of Tables

1.1	Baseline descriptive statistics	16
1.2	Continuity of baseline covariates	22
1.3	Continuity of baseline covariates	23
1.4	Effect on certification probability	26
1.5	Effect on management index	27
1.6	Effect on monitoring & target-setting	30
1.7	Effect on incentives	31
1.8	Effect on likelihood of improving machinery and infrastructure	32
1.9	Effect on log-sales per worker	35
2.1	Baseline descriptive statistics	42
2.2	Continuity of baseline covariates	47
2.3	Effect on technology adoption	49
2.4	Reasons for not adopting new technology	50
2.5	Effect on log-sales per worker	51
2.6	Effect on log-employment	53
2.7	Effect on log-sales	54
3.1	Summary statistics, IMSS individual-level data, largest connected sets	75

3.2	Summary statistics, EIA-IMSS plant panel, 1993	76
3.3	Estimation results for AKM-type model, per period	81
3.4	Baseline estimates, exports and plant-level outcomes	93
3.5	Cross-sectional correlations, alternative proxies, period 1 (1992-1994)	94
3.6	First stage, reduced-form, and IV estimates, alternative proxies	95
3.7	Differential effect of devaluation, IMSS 1988-2002 panel	96
3.8	Differential effect of devaluation, IV estimates, stayers only	97
A.1	Balance of number of previous applications	108
A.2	Correlation between performance and mgt. index	110
A.3	Continuity of baseline covariates	112
A.4	Effect on log-employment	113
A.5	Effect on log-sales	114
A.6	Effect on certification probability	115
A.7	Effect on management index	116
A.8	Effect on monitoring & target-setting	117
A.9	Effect on incentives	118
A.10	Effect on log-sales per worker	119
A.11	Effect on log-employment	120
A.12	Effect on log-sales	121
A.13	Results using honest confidence intervals	122
A.14	Effect on management index controlling for productivity	123
A.15	Effect on monitoring & target-setting controlling for productivity	124

A.16 Effect on incentives controlling for productivity	125
A.17 Effect on changes in customers	128
A.18 Effect on changes in suppliers	129
A.19 Effect on changes in the workforce	130
A.20 Effect on cost reduction	131
A.21 Effect on changes in organizational structure	132
B.1 Balance of number of previous applications	137
C.2.1 Aggregate labor force statistics	145
C.2.2 Summary statistics, IMSS individual-level data, before limiting to largest connected sets	146
C.2.3 Summary statistics, EIA panel, 1993	147
C.2.4 Number of EIA plants linked to IMSS data, by connected set status	148
C.2.5 OLS estimates, exports and plant-level outcomes	149
C.2.6 Construction of alternative proxies, TFP and predicted export share	150
C.2.7 Correlations between proxies, EIA-IMSS panel, pooling periods	151

List of Figures

1.1	Possible actions and their payoff	13
1.2	Treatment probability and score distribution	17
1.3	Continuity of baseline covariates	24
1.4	Effect on certification and management index	25
1.5	Effect on monitoring & target-setting, and incentives	29
1.6	Effect on likelihood of improving machinery and infrastructure	33
1.7	Effect on productivity and performance	34
1.8	Comparing effects on productivity	34
2.1	Treatment probability and score distribution	44
2.2	Continuity of baseline covariates	45
2.3	Effect on technology adoption	48
2.4	Effect on productivity and performance	50
3.1	Wage percentiles, top- and bottom-codes, IMSS data, 1985-2005	70
3.2	Real exchange rate	77
3.3	Shift to export market, EIA panel	78
3.4	Movers' mean real wages, IMSS data, 1992-1995	79
3.5	Movers' mean real wages, IMSS data, 2000-2003	80

3.6	Non-parametric regressions, variables from EIA plant data	91
3.7	Non-parametric regressions, wage components from IMSS employer-employee data	92
A.1	Possible parameter configurations	134
C.2.1	Comparing upward and downward moves, IMSS data, 1992-1995	141
C.2.2	Comparing upward and downward moves, IMSS data, 2000-2003	142
C.2.3	Movers' mean nominal wages, IMSS data, 1992-1995	143
C.2.4	Movers' mean nominal wages, IMSS data, 2000-2003	144

Acknowledgements

I would like express my gratitude to the people and institutions that, in different ways, contributed to this dissertation.

I thank Eric Verhoogen, Jonas Hjort, and Réka Juhász for their advice, support, and guidance. I would also like to thank Michael Best, Laura Boudreau, Donald Davis, Christian Pop-Eleches, Amit Khandelwal, Suresh Naidu, David Weinstein, and Jack Willis for their valuable comments and suggestions.

The research presented in this dissertation could not have been completed without the help of Alejandro Afuso, Celia Cornejo, Rafael Castillo, Luis Mesías, Juan Manuel García, Santiago Cueto, and Eduardo Zegarra. I am also indebted to the *Grupo de Análisis para el Desarrollo* (GRADE) for providing a productive and stimulating work environment during my stay in Peru; and to *Innovate Perú* for providing information needed for this project. Moreover, I also thank the financial support provided by Columbia University's Department of Economics Development Colloquium and Program of Economic Research, the Center for Development Economics and Policy, and the National Science Foundation.

My years in graduate school have been an exciting and challenging period of my life. I thank the intellectual and emotional support provided by my friends Lorenzo Lagos, Divya Singh, and Pablo Warnes.

Finally, I want to express my utmost gratitude to my family and partner. I thank my mother Ana, my sister Nuria, and my father David for their love, inspiration, and support. They put the base of what I am and have achieved today. I also thank my partner, Alejandra, whose patience, strength, and love were key for my happiness and professional accomplishment.

A Ana, Nuria, y David.
A Alejandra. ¡Lo logramos, Lunita!

Chapter 1: Process Standards and Management Practices: Evidence from Peru¹

1.1 Introduction

Several studies suggest that the use of modern management practices, such as monitoring internal processes, setting explicit targets, and incentivizing employees, is important for firm growth and therefore the economic progress of developing countries (Bloom and Van Reenen (2010), Syverson (2011), Bloom et al. (2013), Bloom et al. (2019)). However, many firms, especially those in developing countries, do not adopt such practices (Bloom et al. (2013), McKenzie and Woodruff (2017), Giorcelli (2019)). Studies seeking to understand this lack of managerial upgrading have mainly focused on the role of informational limitations (Bloom et al. (2013), McKenzie and Woodruff (2014), Cai and Szeidl (2017), Bloom et al. (2018), Bloom et al. (2019)). At the same time, studies have paid limited attention to the role of adoption costs, even though such costs play an important role in the analysis of the adoption of other technologies.

Understanding the impediments to managerial upgrading is important because such knowledge sheds light on the barriers to firm growth and it may help us explain the large variation in management practices and productivity that has been observed across firms (Syverson (2004), Hsieh and Klenow (2009), Bloom and Van Reenen (2007), Bloom and Van Reenen (2010), McKenzie and

¹I would like to thank professors Eric Verhoogen, Jonas Hjort, and Réka Juhász for their advice and support; professors Michael Best, Donald Davis, Christian Pop-Eleches, Amit Khandelwal, Suresh Naidu, David Weinstein, and Jack Willis for their valuable comments; Alejandro Afuso, Celia Cornejo, Rafael Castillo, Luis Mesías, Juan Manuel García, Santiago Cueto, and Eduardo Zegarra for their help during this project; Andrea Quijano and Johnathon F. Cherry for their research assistance; and Nuria Alfaro Serrano for her help with the model figures. I am also indebted to the *Grupo de Análisis para el Desarrollo* (GRADE) for providing a productive and stimulating work environment during my stay in Peru; and to *Innovate Perú* for providing information needed for this project. This research would not have been possible without the financial support provided by Columbia University's Department of Economics Development Colloquium and Program of Economic Research, the Center for Development Economics and Policy (CDEP), and the National Science Foundation (NSF, grant number 1829326). All views, opinions, and conclusions expressed are my own.

Woodruff (2017)). Additionally, a good understanding of the reasons behind the lack of managerial upgrading is important for policy design. If informational limitations are the main obstacle to upgrading, actions to promote it should focus on facilitating training and the flow of information about management techniques. If adoption costs also play a relevant role, these initiatives could be complemented with subsidies, a type of intervention that governments have experience dealing with. In fact, papers assessing the effectiveness of policy-feasible interventions that address informational constraints have found mixed results (Bruhn and Zia (2013), Karlan and Valdivia (2011), Giné and Mansuri (2014), Drexler et al. (2014), Valdivia (2015), Karlan et al. (2015), Cai and Szeidl (2017), Higuchi et al. (2019)) while a recent paper by Bruhn et al. (2018) suggests that subsidies might be a helpful policy tool.

An often underappreciated aspect of the cost of adoption of modern management practices is that they rely on the availability of data about processes within the firm, and gathering that data is in itself a difficult task. One of the determinants of this cost is whether firms have documented processes. Documenting firm processes entails explicitly deciding how they will be carried out, designing them in a way that allows the recovery of records about their execution, and writing them down. Once processes have been documented, collecting data about them becomes easier. In this sense, there is complementarity between process documentation and the modern management practices.

This paper tests this idea of complementarity and suggests using it to design a policy-feasible way to promote the adoption of modern management practices. I do this by studying whether a subsidy for the certification of widely known process standards, such as ISO 9001, induces the adoption of modern management practices beyond their scope. In what I call the *certification program*, Peruvian firms received a government subsidy to adopt and certify process standards. This required standardizing and documenting the internal processes of the firm but did not require the adoption of specific practices or the achievement of specific goals. However, as mentioned above, it may be that these standards facilitate the adoption of better management practices by reducing their adoption cost. In other words, in the presence of complementarity between the

standards and sound management practices, subsidizing the adoption of the former should also induce the adoption of the latter.

The study of the complementarity between process documentation and management practices is difficult due to identification and data challenges. The identification challenge is that, in general, firms that acquire process standards are different from those that do not in aspects that might also be relevant for the adoption of modern management practices. I address this issue with the use of a regression discontinuity design in a real policy setting, exploiting the subsidy assignment mechanism. To apply for the subsidies, firms had to submit a project, which was reviewed and given a score. Only those with a score above a certain threshold were funded. In the neighborhood of the threshold, the subsidy can arguable be considered as good as randomly assigned. In addition to the identification challenge, the fact that data about management practices are rarely available also complicates the analysis. To solve this issue, I ran a survey to collect information about management practices in the last quarter of 2018, three years after treatment assignment, and built a management index. Like the index in Bloom and Van Reenen (2010), this index measures the adoption of practices related to monitoring, target-setting, and provision of incentives to employees.

This paper has five main findings. First, the certification program accomplished its stated goal of promoting the certification of process standards. The intervention triggered an increase of 65% in the certification probability. Second, more importantly, the program induced the adoption of modern management practices beyond the scope of the certification. The management index, which runs between zero and one, increased by 0.13. This value is equivalent to 33% of the index mean, 0.56 times its standard deviation, and 35% of the difference between the 25th and 75th percentiles. This finding is consistent with a model in which the process documentation required by the standards and the modern management practices measured by the index exhibit complementarity. Third, managerial upgrading is explained by improvements in monitoring and target-setting practices, while no change is detected in the use of incentives for employees. Fourth, firms that receive the subsidy also show evidence of an increase in the likelihood of improvements

in machinery and infrastructure, consistent with the previous evidence showing that the adoption of modern management practices is associated with firm growth. Fifth, the study also finds a large positive, although statistically nonsignificant, increase in productivity measured as sales per worker. The magnitude of this change is similar to other estimates in the literature.

This paper relates to several strands of literature. The most direct link is with the literature on management practices. In this regard, the contribution is twofold. First, this paper presents evidence consistent with the idea that adoption costs are a relevant barrier to the adoption of better management practices. This complements previous studies that have focused on the role of informational limitations in an effort to understand the lack of managerial upgrading. Second, this paper also suggests a policy-feasible intervention to help reduce those adoption costs: subsidizing the certification of widely popular process standards, which in turn promotes the adoption of better management practices due to their complementarity. From an implementation perspective, this type of intervention has several advantages compared with others that have been studied. First, given that the standards are already codified, the implementing agency does not have to create curricula of knowledge to be transmitted nor choose a set of practices to be promoted. Second, the existence of a market of support services for the adoption of popular process standards means that the implementing agency can limit itself to providing funding and letting the beneficiary firms hire help if they need it. Third, a subsidy for the certification of process standards is simple to monitor because the certification is an easily verifiable signal of project completion. Fourth, the cost is within what a government or development agency can afford. On average, firms treated by the certification program received USD 11,000, which is similar to the cost of other government interventions promoting managerial improvement² and lower than the cost of the intensive interventions studied, for instance, by Bloom et al. (2013) and Giorcelli (2019).³

This paper also relates to the literature on process standards. While this field has seen few

²For example, the Mexican program studied by Bruhn et al. (2018) had a cost of USD 10,670 per firm treated with largest subsidy option.

³The consulting intervention in Bloom et al. (2013) had an approximate market value of USD 250,000 per treated firm, while the intervention studied by Giorcelli (2019) had a cost of USD 38,723 (without considering the fact that its viability relied on the backing provided by the U.S. government).

papers in economics (Volpe Martincus et al. (2010), Masakure et al. (2009), Sun and Ouyang (2014), Bernini et al. (2017), Calza et al. (2019)), there is a larger body of work among management scholars (Sampaio et al. (2009), Tarí et al. (2012), Heras-Saizarbitoria and Boiral (2013), Boiral et al. (2018), Heras-Saizarbitoria et al. (2019), Riaz et al. (2019)). Understanding the changes that these standards trigger within firms is one of the main questions in the field. In this paper, I present quasi-experimental evidence of the impact of adopting process standards on the internal practices of companies.

More broadly, this paper also contributes to the literature on technology adoption. The lack of technology upgrading is a pervasive development problem, and low levels of firm capacity have been advanced as one of the possible causes (Bustos (2011b), Goñi and Maloney (2017), Cirera and Maloney (2017)). This paper presents evidence consistent with this view, as the presence or absence of complementary technologies can be thought of as a component of firm capacity.

The rest of the paper proceeds as follows. Section 1.2 reviews the literature on process standards, management practices, and firm performance. To guide analysis, section 1.3 presents a model that illustrates the idea of complementarity between process standards and modern management practices, and describes how policy can exploit this relationship to foster managerial improvement. Section 1.4 describes the empirical strategy, and section 1.5 presents the results. Finally, section 3.6 concludes.

1.2 Literature review

1.2.1 Management and productivity

The management practices used by a company are one of the determinants of the way in which inputs are transformed into outputs; hence, they are part of technology of the firm⁴. Although it is conceivable that some management techniques are convenient or not dependent on context, some studies have shown that certain management practices are associated with higher firm growth in

⁴Foster and Rosenzweig (2010) define technology as “the relationship between inputs and outputs” (p.396). Bloom et al. (2016) explicitly argue that management can be treated as a technology.

different sectors and countries. Bloom and Van Reenen (2007) and Bloom and Van Reenen (2010) have shown that monitoring what occurs within the firm, setting explicit targets, and providing incentives to employees are practices associated with higher performance. Similar results have been found by McKenzie and Woodruff (2017) in a study focusing on the monitoring and target-setting practices of small firms in developing countries.

Some experimental and quasi-experimental studies have suggested that at least part of the relationship between management practices and performance is causal. Bloom et al. (2013) and Bloom et al. (2018) analyze an experiment in which they assigned Indian textile firms to receive management consultancy services from a top U.S. consultancy company. They find that the intervention led to an increase in the adoption of recommended management practices a few months into the treatment, and produced an increase in productivity (TFP) of 16.6% after one year of treatment, although this figure is noisily estimated. Based on these results, the authors impute an annual increase of approximately USD 325,000 in annual profits, which they estimate to represent, on average, a doubling of profitability. Giorcelli (2019) performed a quasi-experimental analysis of the effect of a U.S. assistance program to improve management on a set of Italian firms. She finds that the program induced the adoption of sound practices, such as performing regular maintenance on machinery or recruiting managers who were not kin. Importantly, this study shows that the effects of the intervention unfolded over a long period of time. Beneficiary firms increased their productivity (TFP) by 15.1% one year after the beginning of the 3-year intervention, and this number reached 49% after 15 years. These findings raise an obvious question: why did the firms not adopt these practices on their own? Bloom et al. (2013) collected data suggesting that informational barriers, in the form of ignoring the existence of the practices or their profitability, were the main obstacles. Giorcelli (2019) also suggests that this type of barrier plays an important role in explaining why the nontreated firms in her setting did not adopt better management practices even when it was apparent that treated firms had improved their performance. In a different study, Bloom et al. (2019) also present evidence suggesting that lack of information plays an important role in explaining why U.S. firms do not upgrade their management. Moreover, Bloom

et al. (2018) and Bloom et al. (2019), also provide evidence suggesting that the use of modern management practices is associated with the mobility of managerial personnel who are knowledgeable about such practices, reinforcing the idea that information availability is an important determinant of managerial upgrading.

While the studies by Bloom et al. (2013) and Giorcelli (2019) suggest that management is important and that firms do not improve it on their own, they do not provide a clear path to creating policies to promote upgrading. In both cases, the interventions cannot just be replicated given their cost and complexity. The consultancy intervention in Bloom et al. (2013) was provided by a top U.S. consultancy firm, had a market cost of USD 250,000 per treated company, and included a one-month diagnostic phase followed by four months of support to implement the recommendations. This support phase was quite intense, as it involved approximately 15 days of consultant time per month-plant and had the explicit goal of helping to implement and stabilize the new practices. The quasi-experimental study by Giorcelli (2019) considered an actually implemented policy, but it was also an abnormally intense one. Managers of treated firms traveled to visit U.S. plants for eight to twelve weeks to learn about their management practices. These trips were followed by a 3-year monitoring period in which U.S. experts visited the treated plants, observed the state of the implementation and provided advice. The explicit cost of the program was USD 38,723 per treated firm. Although this is substantial, it underestimates the complexity of the intervention. The U.S. government had to not only cover the program's explicit cost but also convince U.S. companies to spend time and effort to gratuitously transfer part of their know-how to Italian firms. Clearly, this coordination effort was possible only due to the geopolitical interests at stake.

Given the relevance of management and the idea that lack of information is one of the main obstacles to its improvement, different papers have analyzed policy-feasible ways to help overcome this barrier. These studies have found mixed results, as summarized by McKenzie and Woodruff (2014). Among others, Karlan and Valdivia (2011), Bruhn and Zia (2013), and Drexler et al. (2014) study the impact of in-class business training programs. Karlan and Valdivia (2011) find that adding business training to a standard microcredit program improves knowledge about sound

business practices and induces the adoption of some of them, such as keeping records of bank withdrawal, but not of other important ones, such as keeping records of payments to workers. The authors also find no effects on performance. Bruhn and Zia (2013) find that the beneficiaries of a training program increased their business knowledge and declared with higher probability that they introduced new production processes. Although the authors did not codify the content of the new processes, they indicate that adopting systematic stock management was part of them. Increases in profits were only detected among female-run firms. Drexler et al. (2014) explore in more detail how different teaching techniques alter the effects of training interventions. They compare the traditional approach, which is based on transmitting comprehensive business knowledge in class, with a novel approach based on rules-of-thumbs. Their finding is that the rules-of-thumbs approach induces changes in management practices, while the traditional approach does not. Consistently, the rules-of-thumbs treatment increased an index of revenue measures while the traditional treatment failed to achieve this.

Instead of pure in-class training, Karlan et al. (2015), Valdivia (2015), and Higuchi et al. (2019) explored the effectiveness of on-site consultancy components. Karlan et al. (2015) find that management consultancy administered to tailoring microenterprises in Ghana led to a temporary improvement in management practices, but they reverted back to normal after the first year. Additionally, there were no effects on profitability. Valdivia (2015) randomly allocated Peruvian businesswomen to receive in-class training, or a combination of training and on-site technical assistance. He finds that only a few of the suggested practices were actually adopted and that they depended on the type of treatment received. Regarding performance, the study also finds that both treatment arms led to a 15% increase in sales. However, the link between this improvement and the treatment is hard to establish because the level of take-up and the completion rate were low. Higuchi et al. (2019) randomly allocated a sample of Tanzanian firms to receive in-class training, on-site consultancy, or a combination of both. The content of the classes and consultancy included the standard modules on marketing and accounting and basic modules of Kaizen. The authors find that the three treatment arms induced changes in management practices one year after treatment,

that the combined treatment induced increases in sales and value-added two years after treatment and that the on-site consultancy treatment induced noisily estimated increases in these variables three years after treatment.

In a variation of the previously described consultancy interventions, Bruhn et al. (2018) study the effect of a subsidy that covered between 70% and 90% of the cost of hiring local consultants in Puebla (Mexico). The content of the service was jointly decided by them and the business-owners, and the consultants worked with the firms for one year. The authors find a 26% increase in productivity (TFP) immediately after treatment, although this figure is imprecisely estimated and was not concurrent with an increase in performance. However, within five years after treatment, the authors find that average employment and total wage bill were higher among treated firms.

1.2.2 Process standards

Using process standards to promote the adoption of better management practices is a promising but still unexplored possibility, as the nature and popularity of these standards makes information about them accessible and allows the existence of experienced consultancy firms that provide support.

Process standards⁵ are guidelines to systematize and document the internal processes of the firm. Given that their aim is to be universally applicable, they do not prescribe specific practices or goals. Instead, they require the processes to be standardized, written down, and designed to leave a record (Heras-Saizarbitoria and Boiral (2013), Braun (2005)). Adopting a process standard is challenging and usually involves hiring experienced consultants to help with the task. Firms can certify compliance with these standards. If they decide to follow that path, they need to be audited by a “registrar”, which is a company authorized to certify compliance with a given process standard. This audit includes a review of the required documentation, a review of the process records, and on-site visits and interviews with employees. Different process standards refer to different aspects of firm operations, such as quality management (ISO 9001) or workplace safety

⁵Also called meta-standards (Uzumeri (1997)) in the management literature.

(OHSAS 18001). However, as noted by Heras-Saizarbitoria and Boiral (2013), they are similar with regard to the methodology used for their creation, implementation, and monitoring.

To be more concrete, let us consider the case of the standard ISO 9001, which is the oldest and most popular process standard. To comply with it, firms must produce documentation describing the processes related to the creation of their product. Those documents must indicate the responsibilities of the senior management regarding quality, how resources will be procured and allocated to those processes, and how compliance with the documented procedures will be verified. Additionally, firms have to establish a procedure to create new processes if needed; and must create a “quality manual”, which is a document listing the other documents (similar to an index of processes).

Despite their importance in the business world, process standards have received little attention from economics scholars. Using a combination of matching and diff-in-diff, Volpe Martincus et al. (2010), Masakure et al. (2009), and Sun and Ouyang (2014) found that certifying a management standard has a positive effect on exports. Using a similar method, Bernini et al. (2017) also found a positive effect of certification exports, but no effect on productivity. Employing instrumental variables for identification and a sample of Vietnamese small and medium enterprises, Calza et al. (2019) find that certifying a management standard has a positive effect on productivity, particularly among firms with previous innovation experience and from the most developed part of the country.

Management scholars have produced a notable amount of work on process standards.⁶ One of the central questions in this field is whether their adoption actually induces changes in how firms operate and whether it increases performance. Studies on this topic have found mixed results and, as noted by Heras-Saizarbitoria and Boiral (2013), they have faced important challenges when assessing whether the relationships they find are of a causal nature.

Given that process standards do not prescribe specific modern management practices, why is it reasonable to expect that they will help to foster their adoption? The hypothesis behind this idea is that there is complementarity between the standards’ requirements and the modern managerial

⁶Sampaio et al. (2009), Tarí et al. (2012), Heras-Saizarbitoria and Boiral (2013), and Boiral et al. (2018) are important reviews of this literature. Heras-Saizarbitoria et al. (2019) and Riaz et al. (2019) are recent contributions.

practices related with monitoring, target-setting, and incentives described in the previous subsection. Once processes have been properly described and designed in such a way that their execution leaves a record behind, other management practices become easier to implement. For example, if the execution of the different processes of the firm does not leave information about what was done, who did it, and the results of those actions; then, it is not possible to monitor that processes, or to set verifiable performance goals. The following section presents a model that illustrates this intuition.

From a policy perspective, the key advantage of a subsidy for standards certification over other types of interventions that have been used to try to improve management is that it can be handled like any other subsidy. Implementing agencies do not need to create curricula of knowledge to teach because the content of the standards has already been created, tested, and codified. Additionally, the wide popularity of these standards allows the existence of a market of consultancy firms that can be hired by the beneficiaries to provide support. Their popularity has also incentivized the publication of easily accessible information about the adoption process, and its benefits and challenges.

1.3 Model

In this section, I present a simple model that illustrates the idea that in the presence of complementarity between two technologies, a subsidy to promote one of them (for example, process documentation required by process standards), can also help to promote the adoption of the other (for example, modern management practices). In the model, a firm adopts technologies in two periods, and the technologies adopted in those periods exhibit complementarity, meaning that technology upgrading in one moment increases the benefit of upgrading in the next one.

1.3.1 Setup

In this model, a *firm* uses a single input l and the production function $f(l) = Al^\alpha$ ($0 < \alpha < 1$). It maximizes profit taking the prices of the product (p) and input (w) as given. This way, the

problem of the firm is:

$$\max_l pAl^\alpha - wl \quad (1.1)$$

The optimum level of input is given by $l = \left(\frac{\alpha p A}{w}\right)^{\frac{1}{1-\alpha}}$ and the corresponding operating profit (meaning, the profit without considering technological investments) is $A^{\frac{1}{1-\alpha}} b$ with $b := p^{\frac{1}{1-\alpha}} w^{\frac{-\alpha}{\alpha-1}} (\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}})$. A is the productivity provided by the technology used by the firms. For algebraic convenience I represent this productivity with the monotonically increasing transformation $\phi := A^{\frac{1}{1-\alpha}}$.

7

The set of existing *technologies* is $\{t_0, t_1, t_2, \dots\}$, and ϕ_{t_i} is the productivity provided by technology t_i ($\phi_{t_0} < \phi_{t_1} < \phi_{t_2} < \dots$). The cost of adopting a new technology depends on the current technology of the firm. More specifically, for a firm that has already adopted technology t_i , which provides productivity ϕ_{t_i} , the cost of adopting technology t_j , which provides productivity ϕ_{t_j} , is $(\phi_{t_j} - \phi_{t_i})^2$.

The *timing* of actions is as follows: The firm is born with technology $T_0 \in \{t_0, t_1, t_2, \dots\}$, which provides productivity ϕ_{T_0} , and has two periods to adopt new technologies before producing. The variables T_1 and T_2 represent the technologies adopted by the firms in periods 1 and 2. Naturally, their support is also the set of existing technologies $\{t_0, t_1, t_2, \dots\}$ and they provide productivity ϕ_{T_1} and ϕ_{T_2} respectively. Given that adoption costs are sunk, it never makes sense for the company to downgrade its technology, hence that option is ruled out. After these two periods, the firm produces using the last technology adopted, and obtains the corresponding profit $\pi(\phi_{T_1}, \phi_{T_2})$, which is equal to the operating profit minus the cost of technological adoption.

$$\pi(\phi_{T_1}, \phi_{T_2}) = \phi_{T_2} b - (\phi_{T_2} - \phi_{T_1})^2 - (\phi_{T_1} - \phi_{T_0})^2 \quad (1.2)$$


Note that in this setup, there is complementarity between the technologies adopted in different periods: The better the technology adopted in the first period, the higher the marginal benefit of

⁷This is similar to the presentation of the model in Bustos (2011b), in which the productivity parameter is graphically represented by a monotonically increasing transformation.

technology upgrading in the second one. Formally, this means that the profit function is supermodular in (ϕ_{T_1}, ϕ_{T_2}) . The proof of this property is in appendix A.9. At the core of this complementarity is the quadratic adoption cost, which represents the idea that big technological jumps are more costly than a series of incremental improvements.

For simplicity, let's consider the case in which there are three possible technologies t_0, t_1, t_2 with productivities $\phi_{t_0}, \phi_{t_1}, \phi_{t_2}$ ($\phi_{t_0} < \phi_{t_1} < \phi_{t_2}$), the firm is born with technology $T_0 = t_0$, and the productivity of this initial technology is normalized to 1 ($\phi_{T_0} = \phi_{t_0} = 1$). The possible actions and payoffs are represented in figure 1.1.

Figure 1.1: Possible actions and their payoff



Different actions could be optimal for the firm depending on the specific values of the parameters ϕ_{t_1} and ϕ_{t_2} .

1.3.2 Exploiting complementarity to promote upgrading

In this subsection I show that, given an initial situation with no technological upgrading, a subsidy for the adoption of a simple technology can be used to foster the adoption of a more

advanced one.

Consider a situation in which the following two conditions hold:

$$\phi_{t_2}b - (\phi_{t_2} - \phi_{t_1})^2 - (\phi_{t_1} - 1)^2 > \phi_{t_1}b - (\phi_{t_1} - 1)^2 \quad (1.3)$$

$$b > \phi_{t_2}b - (\phi_{t_2} - \phi_{t_1})^2 - (\phi_{t_1} - 1)^2 \quad (1.4)$$

In this case, the firm would keep technology t_0 until production time; despite the fact that, if it upgraded to t_1 in period 1, it would also update to t_2 in period 2.⁸

Now consider a subsidy S for the adoption of technology t_1 in period 1. Naturally, the value of this subsidy can be at most the cost of adopting this technology in the first period, $S < (\phi_{t_1} - 1)^2$. If S were at least $b - \phi_{t_2}b + (\phi_{t_2} - \phi_{t_1})^2 + (\phi_{t_1} - 1)^2$, the direction of the inequality in condition (1.4) would change, and the firm would adopt t_1 and t_2 .

Note that for this to be possible, it has to be the case that $b - \phi_{t_2}b + (\phi_{t_2} - \phi_{t_1})^2 + (\phi_{t_1} - 1)^2 < (\phi_{t_1} - 1)^2$ which is equivalent to $\phi_{t_2} < \frac{1}{2} \left(\sqrt{b(b + 4\phi_{t_1} - 4)} + b + 2\phi_{t_1} \right)$. This inequality puts a limit on how far can ϕ_{t_2} be from ϕ_{t_1} . Intuitively, in order to exploit the complementarity, the final technology t_2 cannot be substantially more advance than t_1 .

Wouldn't be more convenient to simply subsidize the adoption of t_2 ? No. To induce the direct adoption of t_2 , the magnitude of the subsidy S would need to be at least $b - \phi_{t_2}b + (\phi_{t_2} - 1)^2$, which is larger than the required subsidy to induce the adoption of t_1 . Additionally, if t_1 represented process documentation and t_2 represented one of the modern management practices, subsidizing t_2 would be more difficult than subsidizing t_1 .

In this subsection, I have shown how a small subsidy for the adoption of a simple technology can be used to foster the adoption of a more advanced one. In appendix A.8, I show that this type of intervention can be welfare improving if the initial situation exhibits market imperfections.

⁸Note that $b > \phi_{t_1}b - (\phi_{t_1} - 1)^2$ and $b > \phi_{t_2}b - (\phi_{t_2} - 1)^2$ are implied conditions (1.3) and (1.4).

1.4 Empirical strategy

1.4.1 The Certification Program and the identification strategy

The certification program, which was implemented for the first time in 2015, is managed by *Innovate Perú*, a Peruvian government agency in charge of programs to promote productive development. This program supports projects to certify process standards. Beneficiary firms receive a subsidy to cover up to 50% of the cost of the project, with a maximum of USD 14,000. Formal firms of any sector can apply provided that they had sales of less than USD 2.8 million the year before. In this analysis, I use data from the first three rounds (03/2015, 05/2015, and 10/2015). They include 250 applications, of which 127 were funded⁹. Table 1.1 reports descriptive statistics at the moment of application and compares them with country-level figures. Applicants were, on average, 10 years old and were concentrated in the secondary and tertiary sectors. Micro- and small-sized firms accounted for 81% of the participants, and 89% of them came from the coastal region of the country, which is the most economically developed region. Successful applicants were slightly, larger and older than unsuccessful ones. Compared with national figures, participants were larger, more focused on the secondary sector, and slightly more concentrated in coastal areas of the country.

To apply, firms submitted a project detailing their plan. Those projects were evaluated by external reviewers, who assigned them a score between 0 and 100 on four criteria: impact and relevance (35%), viability (25%), cost-benefit relationship (25%), and complementary factors (15%). The final score of the project was the average of those values. The subsidy was offered to projects with 70 points or more. I exploit this fact to implement a regression discontinuity analysis, using the score as the running variable. Figure 1.2a shows the treatment probability as a function of the running variable. It increases by approximately 80% at the cutoff. Figure 1.2b shows the distribution of the running variable. As can be seen, the score of the applications is a discrete variable. This is because reviewers used only a few salient values to evaluate the projects, and because only four

⁹Appendix A.1 provides details of sample construction.

Table 1.1: Baseline descriptive statistics

	Total (1)	Nontreated (2)	Treated (3)	National (4)
A. Sector				
primary	6.0%	5.7%	6.3%	2.3%
secondary	32.0%	29.3%	34.6%	11.0%
tertiary	62.0%	65.0%	59.1%	86.7%
B. Size				
mse	81.2%	81.1%	80.5%	99.4%
medium	18.8%	18.1%	19.5%	0.6%
C. Trade				
exported	13.2%	13.8%	12.6%	
imported	26.4%	26.0%	26.8%	
D. Other characteristics				
coast	88.8%	89.4%	88.2%	72.7%
age (yr)	10.2	9.5	10.9	

Column (1) shows percentages corresponding to the total of received applications, column (2) refers only to nonsubsidized applications, and column (3) refers to subsidized applications. Column (4) refers to national statistics recovered from the report *Perú: Estructura Empresarial 2015* prepared by the Peruvian national statistical institute (*Instituto Nacional de Estadística e Informática (INEI)*). Micro and small firms (MSE) are those with sales of less than USD 2.1 million. Medium firms are those with sales between 2.1 and 2.9 million.

criteria were assessed.

I estimate both intent-to-treat effects (ITT) and local average treatment effects (LATE) using local linear regressions.¹⁰ To estimate ITT, I compute the OLS estimate of β_1 in the following specification:

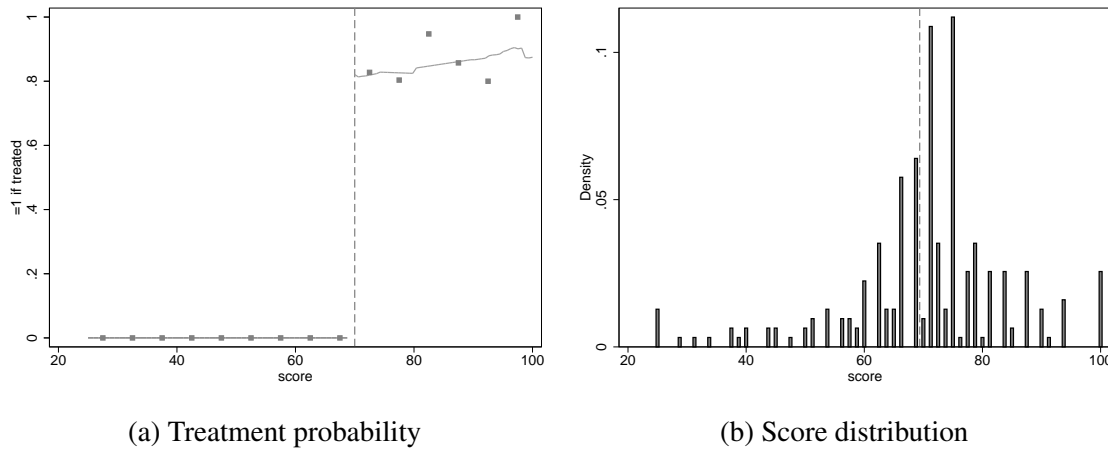
$$y = \beta_0 + \beta_1 \mathbf{1}(s \geq 70) + f(s) + u \quad (1.5)$$

where $\mathbf{1}(s \geq 70)$ is a dummy variable taking value one if the running variable s is to the right of the cutoff, $f(s)$ is polynomial of the running variable that allows for different slopes on different sides of the cutoff. To estimate LATE, I compute the IV estimate of β_1 in the following specification:

$$y = \beta_0 + \beta_1 \textit{treated} + f(s) + u \quad (1.6)$$

where *treated* is a dummy variable that takes value one for treated observation. I instrument this variable with $\mathbf{1}(s \geq 70)$. In both cases, I use observations restricted to different bandwidths around the cutoff and uniform kernel.

Figure 1.2: Treatment probability and score distribution



The graph on the left shows a local linear regression of a dummy variable that takes value one if a subsidy for the project was granted and zero otherwise, on the running variable. The local linear regression uses uniform kernel and a bandwidth of 20 points. The graph on the right shows the probability distribution of the running variable.

Although the discreteness of the running variable does not affect the identification arguments,

¹⁰In this case, both ITT and LATE are local in the sense that that apply only to observations at the cutoff. LATE is additionally restricted to compliers.

it might be a challenge for inference, as noted by Lee and Card (2008) and Kolesár and Rothe (2018). The challenge is that in some cases, standard confidence intervals for treatment effects might have incorrect coverage when the running variable is discrete because the uncertainty about the behavior of the conditional expectation function between the running variable mass points is not eliminated even as the sample size grows. In practice, this is a problem only when the number of support points is small; hence, in general, there is no need to distinguish sharply between the discrete and continuous cases. However, to be cautious, I check that the conclusions obtained with the standard methods are robust to the use of methods that are not affected by the discreteness of the running variable. More specifically, I check the robustness of my results using the local randomization approach proposed by Cattaneo et al. (2015) and Cattaneo et al. (2018), and the “honest confidence intervals” proposed by Kolesár and Rothe (2018). The first approach consists of treating the observations in a narrow window around the cutoff as coming from an actual experiment and using finite sample randomized inference in that window to test the sharp null hypothesis of no effect. This renders the discreteness of the running variable irrelevant because the actual value of the running variable is not used after the window around the cutoff has been selected. In my estimations, I use different windows to show robustness. These windows are narrower than the bandwidths used for local linear regression because the local randomization assumption is stronger than the continuity one that underlies the usual procedure. The local randomization approach is my preferred alternative method because the use of finite sample inference methods makes it well suited to address my reduced sample size. The second approach consists of an alternative way to build confidence intervals. These honest confidence intervals have good coverage when the running variable is discrete provided that the underlying conditional expectation function is bounded, although they tend to be very broad.

1.4.2 Data sources

The main sources of data I use to study the certification program are the administrative records provided by *Innovate Perú*, the implementing agency, and a survey I ran in the last quarter of

2018. The administrative data contain some baseline information from the applications and their scores. The survey was administered to managers of the firms and was used to collect data about the certification of standards, the use of different management practices, and standard information about sales and employment.

To learn about the management practices of the firm, I asked a series of questions to assess management regarding monitoring, target-setting, and incentives. As mentioned by Bloom and Van Reenen (2010) these practices capture what is commonly understood as good management and are associated with high productivity and performance. Monitoring refers to the extent to which the activities that take place within the firms are supervised. A firm with good monitoring practices is one that systematically tracks and reviews the performance of different processes. Target-setting refers to the extent to which firms set explicit and verifiable goals. A company with sound target-setting practices is one that sets goals that are comprehensive and expressed in terms of variables that can be measured (as opposed to vague statements). The use of incentives refers to the extent in which an employee's effort is rewarded. Better incentive-related practices are those that allow the employee to benefit from the effort exerted.

To be more specific, monitoring was assessed with questions about whether the firm uses performance indicators, how many are used, and how often they are reviewed by managers. Target-setting was measured with questions about whether the firm uses explicit targets and the time-horizon considered. The information about incentives was collected by asking about the use of bonuses, worker reassignment and firing procedures, and promotion procedures. As in Bloom et al. (2019), I scored the answers to these questions with a variable taking value 0 if the practices are not used and 1 if it is used to a large extent. Then, I averaged those variables per type of practice. My management index is equal to the average of the scores for monitoring, target-setting, and incentives practices. Appendix A.2 describes the construction of the index in more detail.

In addition to my main data sources, I also use trade data provided by the Chamber of Peruvian Exporting Firms¹¹ and data manually collected from the tax authority's website. The trade data

¹¹*Asociación de Exportadores (ADEX).*

allow me to observe the export and import behavior of the companies and are provided through the commercial service ADEX Data Trade, which collects data from the customs administration and makes them accessible through a web-based interface. From the tax authority's website, I collected information regarding firm's sector and age through the service called *Consulta RUC*. Appendix A.2 also provides more details about how these additional data sources were used.

1.4.3 Validity of the identification strategy

The validity of a regression discontinuity design as an identification strategy relies on the assumption that the value of the running variable has not been manipulated around the cutoff. In my case, firms cannot directly manipulate their score, so altering it would require some form of collusion between the firm and the reviewer. I do not believe that is a relevant concern in the case of the certification program, as the identity of the reviewer is decided after submitting the project and the firm never learns who that person was. Additionally, the implementing agency monitors that the project is actually executed and the money spent on it through on-site visits and requiring proof of payment and execution. This reduces the amount of resources that could be split between reviewer and firm-owner.

In addition to these considerations, the lack of manipulation has observable consequences that can be tested. In the absence of manipulation, there should not be discontinuities in the baseline covariates at the cutoff. Indeed, there is no evidence of such discontinuity. Figure 1.3 displays local linear regressions of different baseline covariates against the running variable, and shows that there is no evidence of an abrupt change at the cutoff value of the running variable. This is confirmed in table 1.2, which reports OLS estimates of the change in baseline characteristics at the cutoff using different bandwidths and polynomial degrees. None of the changes are statistically significant. Similarly, table 1.3 reports the mean difference of the baseline covariates between both sides of the cutoff and the randomized inference p-value of tests of the sharp null hypothesis of no difference. No significant difference is detected with this procedure and this conclusion is robust to the use of different windows around the threshold value. Appendix A.3 that shows the conclusions

are the same when using honest confidence intervals. Additionally, if manipulation did not occur, the running variable should be smoothly distributed around the cutoff. The usual practice is testing this using the McCrary (2008) test; however, this procedure is not valid when the running variable is discrete. An alternative test for the discrete case has been proposed by Frandsen (2016). This test compares the second differences around the cutoff with those occurring in other parts of the distribution. The null hypothesis is that the second differences in the cutoff area are not different from those occurring in other parts of the running variable's support. One of the disadvantages of this procedure compared with the McCrary (2008) test is that its ability to detect deviations from the null hypothesis decreases with the jaggedness of the running variable distribution. However, it is still informative that, in my case, this procedure fails to reject the null hypothesis of smoothness at the cutoff with a p-value of 0.648.¹²

1.5 Effects of the certification program

In this section, I show that the certification program accomplished its stated goal of promoting the certification of process standards and, additionally, induced the improvement of management practices beyond the scope of the certification. Moreover, I also provide evidence suggesting that the reason for the additional managerial improvement is the complementarity between the certification requirements and the management practices adopted. Regarding productivity, I find large positive effects, although they are not statistically significant at the conventional levels.

1.5.1 Effects on certification and management practices

I find that assignment to treatment increased the probability of certifying a process standard by approximately 65 percentage points, compared with a certification probability of 17% among non-treated firms. Figure 1.4a shows the result of a linear local regression of a dummy variable

¹²To implement this test, the researcher has to provide the value of a parameter k measuring the degree of jaggedness of the running variable's probability function. Following Frandsen (2016), I based my choice of k on the behavior of the probability function away from the cutoff. More specifically, I computed k for all the support points, except the cutoff and its adjacent mass points, and used the average of those values in the test reported above. Additionally, as a robustness check, I also implemented the test using the average k in the full support, and the average k excluding only the cutoff value. In all these cases, the test fails to reject the null hypothesis of continuity.

Table 1.2: Continuity of baseline covariates

	(1) age (mo)	(2) mse	(3) manuf.	(4) services	(5) coast	(6) export	(7) import
A. bw=10							
right	-13.874 (18.583)	0.107 (0.111)	0.036 (0.130)	0.078 (0.139)	0.067 (0.083)	0.009 (0.097)	0.050 (0.125)
obs	172	172	172	172	172	172	172
B. bw=20							
right	12.783 (17.247)	-0.064 (0.100)	0.033 (0.106)	0.042 (0.113)	0.019 (0.060)	0.049 (0.077)	0.091 (0.103)
obs	219	219	219	219	219	219	219
C. bw=30							
right	10.945 (15.056)	-0.073 (0.091)	0.100 (0.097)	-0.037 (0.104)	0.026 (0.055)	0.033 (0.069)	0.024 (0.093)
obs	240	240	240	240	240	240	240
D. Full sample							
right	10.747 (14.778)	-0.064 (0.081)	0.093 (0.089)	-0.043 (0.095)	0.021 (0.052)	0.020 (0.064)	0.024 (0.087)
obs	250	250	250	250	250	250	250
E. Full sample, cubic polynomial							
right	-2.690 (23.382)	0.135 (0.127)	0.124 (0.148)	-0.013 (0.156)	0.047 (0.096)	0.094 (0.116)	0.191 (0.141)
obs	250	250	250	250	250	250	250

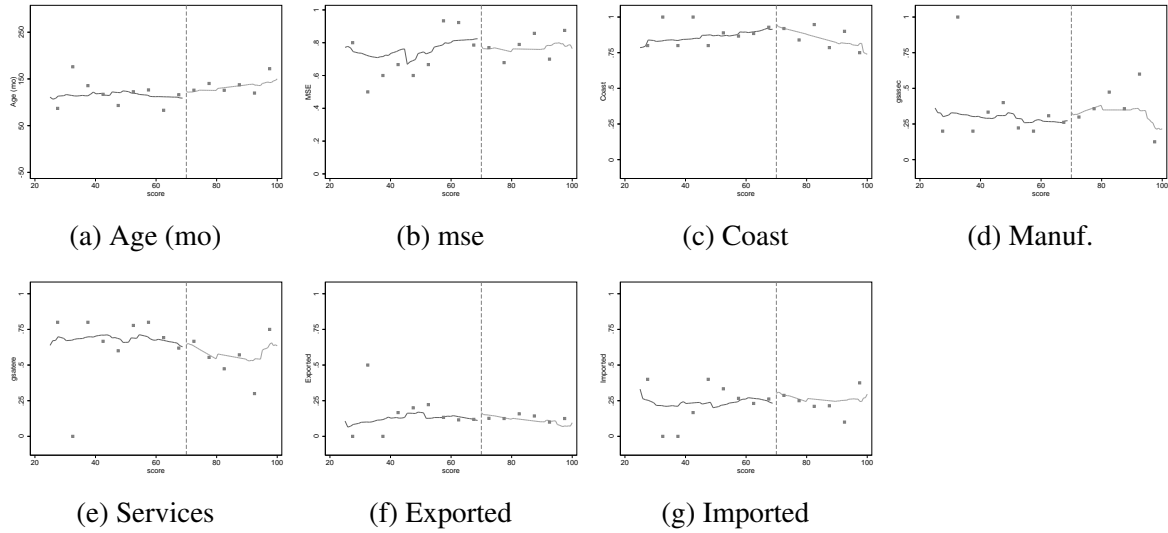
***, **, * indicate significance at 1%, 5%, and 10%. The values reported in this table are the OLS estimates of β_1 in equation (1.5) taking the variable indicated at the top as the dependent variable. Values in parenthesis are standard errors clustered at the level of the firm. Panels A - C report estimates using a bandwidth of 10, 20, and 30 points around the cutoff, uniform kernel, and a linear polynomial of the running variable. Panel D reports similar estimates using the full sample. Panel E reports similar estimates using the full sample and a cubic polynomial of the running variable.

Table 1.3: Continuity of baseline covariates

	(1) age (mo)	(2) mse	(3) manuf.	(4) services	(5) coast	(6) export	(7) import
A. Two mass points							
diff. in means	8.658	0.102	0.142	-0.010	0.052	0.111	0.087
rand. inf. p-value	0.573	0.375	0.202	1.000	0.600	0.218	0.411
obs	75	75	75	75	75	75	75
B. 67.5 - 72.5							
diff. in means	8.296	-0.038	0.162	-0.183	0.058	0.067	0.162
rand. inf. p-value	0.641	0.777	0.226	0.146	0.584	0.687	0.216
obs	68	68	68	68	68	68	68
C. 65 - 75							
diff. in means	9.752	-0.016	0.037	0.048	-0.009	0.007	0.025
rand. inf. p-value	0.543	1.000	0.688	0.686	1.000	1.000	0.849
obs	129	129	129	129	129	129	129
D. 62.5 - 72.5							
diff. in means	18.055	-0.057	0.021	0.014	0.015	0.013	0.008
rand. inf. p-value	0.187	0.460	0.841	0.854	0.754	1.000	1.000
obs	153	153	153	153	153	153	153

In each panel of this table, each column reports three values. The first value is the mean difference of the variable indicated at the top between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the difference is zero. The third value is the number of observations. Estimates in panel A use only observations within two mass points of the cutoff, the narrowest window in which the test can be implemented. Estimates in panels B, C, and D use windows with radius of 2.5, 5, and 7.5 points.

Figure 1.3: Continuity of baseline covariates



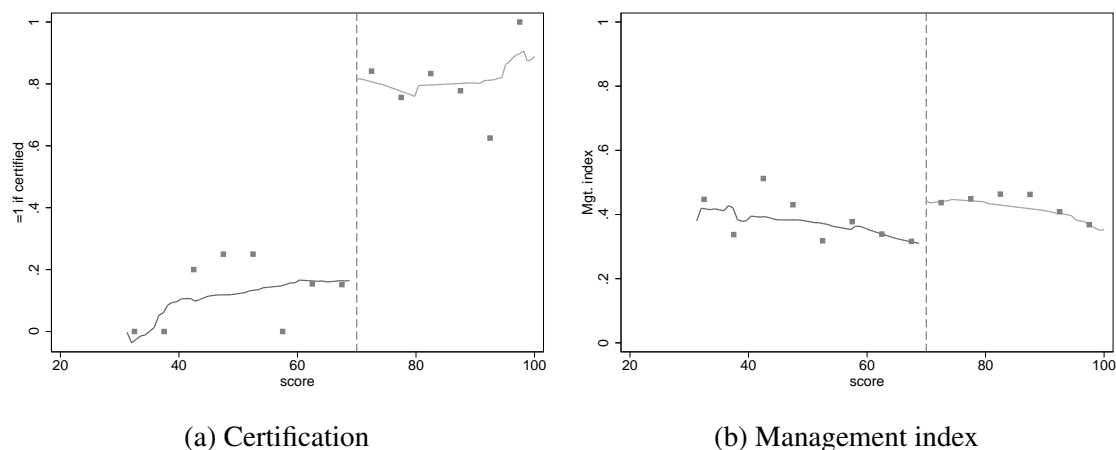
Each of these graphs shows the local linear regression of a baseline covariate on the running variable, using a uniform kernel and a bandwidth of 20 points. In graph (a), the outcome variable is the age of the firm in months. In graph (b), it is a dummy variable that indicates the firm falls in the category of micro and small firms. In graph (c), it is a dummy variable that indicates the firm is located in the coast region. In graph (d), it is a dummy variable that indicates the firm is classified in the manufacturing sector. In graph (e), it is a dummy variable that indicates the firm is classified in the service sector. In graph (f), it is a dummy variable that indicates that firm had exported. In graph (g), it is a dummy variable that indicates that the firm had imported.

indicating that the standard certification was obtained, on the running variable. The vertical axis indicates the probability of being certified, the horizontal axis displays the running variable, and the dashed vertical line indicates the cutoff value. The increase in the treatment probability is apparent, and the estimates in table 1.4 confirm this finding. Columns 1 - 3 in table 1.4 report OLS and 2SLS estimates of the ITT and LATE of interest using local linear regression for bandwidths 10, 20, and 30. Columns 4 and 5 report similar estimates using the full sample with linear and cubic polynomials. Finally, columns 6 - 8 report the p-values of randomized inference tests using observations in different narrow windows around the cutoff. In column 6, only observations located within two running variable mass points of the cutoff are used. In columns 7 and 8, observations within 5 and 7.5 points, respectively are considered.

Consistent with the hypothesis of complementarity between process standards and modern management practices, I also find that assignment to treatment increases the management index, which varies between 0 and 1, by approximately 0.13 points. This is a sizable increase. Its value

is equal to 33% of the mean value of the management index in 2018 and equal to 35% of the difference between the 25th and 75th percentiles of its distribution. The increase in the management index is graphically reported in figure 1.4b, and formally corroborated in table 1.5, which follows a format similar to the previous table. A potential concern with interpreting the positive estimate of the effect of the program on the management index as supporting the complementarity hypothesis is that this effect could also be explained by the fact that the subsidy represented an increase in income for firms that would have sought the certification regardless of their treatment status. Although it is not possible to rule out the existence of this income effect, I believe it does not drive the effect on management as the proportion of always-takers is just around 17% according to the proportion of non-treated firms that certified process standards.¹³

Figure 1.4: Effect on certification and management index



Each of these graphs shows a local linear regression of the dependent variable indicated below, on the running variable using a uniform kernel and a bandwidth of 20 points. The dots represent the mean of the outcome variable in 5-point bins. In the graph on the left, the dependent variable is a dummy variable that indicates whether the proposed certification was obtained. In the graph on the right, the outcome variable is the management index.

If the complementarity between process standards and modern management practices is one of the reasons why the certification program improved management quality, it would be reasonable to expect that an important part of the improvement is explained by changes in practices closely related to process standardization and documentation. To check if this is the case, I divide the

¹³A similar concern is discussed by Duflo et al. (2019) in a study about the impact of high school scholarships in Ghana.

Table 1.4: Effect on certification probability

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.689*** (0.135)	0.654*** (0.100)	0.620*** (0.089)	0.607*** (0.084)	0.590*** (0.157)	0.655	0.690	0.642
rand. inf. p-val.						0.000	0.000	0.000
B. LATE estimates								
treated	0.824*** (0.148)	0.798*** (0.115)	0.734*** (0.100)	0.720*** (0.094)	0.777*** (0.190)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a dummy variable that indicates if the proposed certification was obtained. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table 1.5: Effect on management index

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.134* (0.081)	0.129** (0.064)	0.164*** (0.058)	0.148*** (0.057)	0.181** (0.090)	0.140	0.121	0.115
rand. inf. p-val.						0.033	0.020	0.014
B. LATE estimates								
treated	0.160* (0.097)	0.158** (0.078)	0.195*** (0.069)	0.175** (0.068)	0.239** (0.122)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the management index. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.


management index into two subindices: one containing information about monitoring and target-setting and the other one containing information about incentives for employees. The practices included in the first subindex rely more heavily on the availability of process-level information than do the practices included in the second subindex, and are expected to react more intensely if complementarity is playing a role. Consider, for example, the case of performance indicators. Using performance indicators is part of having a good monitoring system. To build such indicators, managers must be able to gather information about the processes they want to monitor and to summarize that information in one or more numbers. Process standardization and documentation make it easier to construct those indicators because they help to delimit the different tasks and require leaving records that can later be used to collect the information required. This way, even though adopting a process standard is not the same as having a good monitoring system in place, the former eases the challenge of adopting the latter. Unlike monitoring and target-setting practices, incentive practices require not only information but also other changes such as modifying contracts with employees, which are not provided by a process standard.

As expected, the improvement in the overall management index occurred due to improvements in monitoring and target-setting, while no change is detected in incentives for employees. This difference is apparent in figures 1.5a and 1.5b, and is corroborated in tables 1.6 and 1.7.

1.5.2 Other changes within the firm

In addition to management upgrading, I also find evidence of improvements in the firms machinery and infrastructure. As part of the survey, I asked managers whether some changes had taken place between 2015 and 2018 in the following categories: i) changes in organizational structure, ii) changes in the workforce, iii) improvements in machinery and infrastructure, iv) reductions in production costs, v) changes in customers, vi) changes in suppliers. The details of the questions asked per category are presented in appendix A.7.1. I summarized their answers using the Kling et al. (2007) index, which is also used by Bruhn et al. (2018) for a similar purpose. This procedure consists in recentering the answer to each question at the mean of the control group, rescaling it

Figure 1.5: Effect on monitoring & target-setting, and incentives



Each of these graphs shows the estimates of a local linear regression of the dependent variable on the running variable using a uniform kernel and a bandwidth of 20 points. The dots represent the mean of the outcome variable in 5-point bins. In the graph on the left, the dependent variable is the subindex of monitoring and target-setting practices. In the graph on the right, the outcome variable is the subindex of incentive practices.

by its standard deviation, and building the topic-level index as the average of those recentered and rescaled variables. Of the six categories of changes considered, I only find consistent evidence of impact for improvements in machinery and infrastructure. This can be seen in figure 1.6 and corroborated in table 1.8. The results for the other topics can be found in appendix A.7.2. Finding that the certification program led to an increase in the likelihood of improving machinery and infrastructure is consistent with the previous finding regarding managerial upgrading. As described in the literature review, the existing evidence suggests that the adoption of modern management practices boosts firm growth and that such effect occurs over long periods of time. If this process is indeed taking place, it would be reasonable to observe that, at least some the physical assets involved in the production process, are also been upgraded.

1.5.3 Effects on productivity and performance

Three years after treatment assignment, I find a positive, although statistically nonsignificant, effect of the certification program on productivity measured as log-sales per worker. The magnitude of this estimate is economically relevant. Columns 1 - 3 of table 1.9 report OLS and 2SLS estimates of the ITT and LATE using local linear regressions for bandwidths 10, 20, and

Table 1.6: Effect on monitoring & target-setting

	Usual inference				Randomized inference			
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.137* (0.082)	0.176*** (0.064)	0.211*** (0.059)	0.191*** (0.058)	0.170* (0.090)	0.153	0.152	0.156
rand. inf. p-val.						0.022	0.006	0.002
B. LATE estimates								
treated	0.164* (0.099)	0.214*** (0.080)	0.250*** (0.071)	0.226*** (0.070)	0.225* (0.124)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the the subindex of monitoring and target-setting practices. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

30. Columns 4 and 5 report similar results using full sample, and linear and cubic polynomials. These estimates are all positive, stable, but statistically nonsignificant at conventional levels even though their values are larger than their standard errors. Columns 6 - 8 in table 1.9 report randomized inference p-values for the ITT and indicate that the sharp null hypothesis of zero effect is rejected for the narrowest window around the cutoff, although this significance is lost for larger windows. Figure 1.7a shows the estimates of a local linear regression of log-sales per worker on the running variable. A small discontinuity can be seen at the cutoff, which is consistent with the OLS estimates of the ITT, and the randomized inference findings for observations close to the

Table 1.7: Effect on incentives

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.052 (0.094)	0.039 (0.074)	0.073 (0.070)	0.048 (0.066)	0.091 (0.112)	0.060	0.050	0.048
rand. inf. p-val.						0.446	0.386	0.339
B. LATE estimates								
treated	0.062 (0.111)	0.048 (0.090)	0.086 (0.082)	0.057 (0.078)	0.120 (0.147)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the subindex of incentives. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

threshold value of the running variable. Regarding sales and employment, I find no evidence of impact as seen in figures 1.7b and 1.7c. Tables in appendix A.4 report estimates and formal tests for these effects. These estimates are noisier than in the case of productivity.

The magnitude of the ITT estimates of the effect of the program on productivity oscillates around an increase of 40%.¹⁴ Figure 1.8 puts this value the context of other papers in the literature. The vertical axis of this graph shows the percentage change in productivity or profitability estimated by other studies on management-enhancing interventions. The horizontal axis indicates the


¹⁴ $100 * (\exp(0.337) - 1) = 40$

Table 1.8: Effect on likelihood of improving machinery and infrastructure

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.512* (0.261)	0.663*** (0.208)	0.616*** (0.191)	0.565*** (0.182)	0.524* (0.312)	0.373	0.490	0.528
rand. inf. p-val.						0.084	0.013	0.002
B. LATE estimates								
treated	0.612* (0.316)	0.810*** (0.261)	0.730*** (0.230)	0.670*** (0.219)	0.690 (0.424)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the index of improvement in machinery and infrastructure. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Figure 1.6: Effect on likelihood of improving machinery and infrastructure




This graph shows the estimates of a local linear regression of the Kling et al. (2007) index for improvements in machinery and infrastructure on the running variable using a uniform kernel and a bandwidth of 20 points. The dots represent the mean of the outcome variable in 5-point bins.

time elapsed since treatment assignment. The solid markers indicate estimates that are significant at conventional levels. The magnitude of the estimate in this paper is similar to those of previous studies.

The fact that the effect of the program on productivity is hard to detect despite the size and stability of the estimates is consistent with previous studies in the literature. Papers have shown that the effect of management-enhancing interventions on productivity and performance develops over several years and is difficult to detect shortly after treatment. For example, Giorcelli (2019) finds that the impact of the U.S. management support for Italian firms continued growing even 15 years after treatment assignment. Similarly, Bruhn et al. (2018) are able to detect the effect of a subsidy for management consulting on employment only six years after assignment.

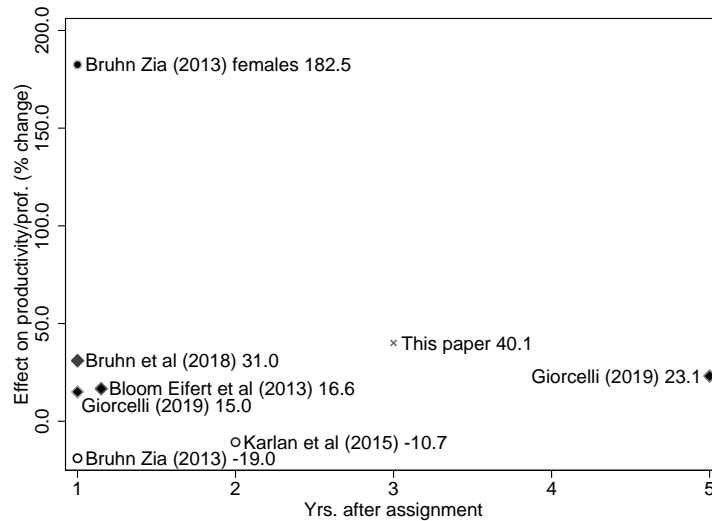
In previous sections, I have argued that the certification program helped improve management by reducing the cost of adopting modern management practices. However, the large estimates of the effect of the certification program on productivity raise a question about the mechanism connecting the intervention and the managerial improvement described above. More specifically, it could be possible that the certification program led to an increase in productivity through channels other than management quality and then, the increase in productivity induced the adoption of modern managerial practices. To address this concern, I reestimate the effect of the program

Figure 1.7: Effect on productivity and performance



Each of these graphs shows the estimates of a local linear regression of the dependent variable on the running variable using a uniform kernel and a bandwidth of 20 points. The dots represent the mean of the outcome variable in 5-point bins. In the graph on the left, the dependent variable is log-sales per worker. In the graph in the center, the outcome variable is log-employment (headcount), and in the graph on the right, the outcome variable is log-sales. Dependent variables were winsorized at the 10th and 90th percentiles to control for outliers.

Figure 1.8: Comparing effects on productivity



This figure presents results estimates from other studies about the effect of interventions to improve management on measures of productivity or profitability. The estimates have been transformed to represent percentage changes. Solid markers represent estimates significant at conventional levels, and hollow markers represent non-significant estimates. The estimates from Bloom et al. (2013), Bruhn et al. (2018), and Giorcelli (2019) refer to effects on TFP. Original results are coefficients from regression that take log-TFP as the dependent variable. These coefficients have been transformed by applying the exponential function, subtracting one, and multiplying by 100 to obtain percentage changes. The estimate from this paper comes from column (1) in table 1.9, and the same transformation has been applied to it. Bruhn and Zia (2013) report absolute changes in profits. These were transformed by reexpressing them as percentages of the value of this variable in the control group. Karlan et al. (2015) report absolute changes in revenue minus expenses. This value has been reexpressed as percentages of the value of this variable in the control group.

Table 1.9: Effect on log-sales per worker

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.337 (0.323)	0.413 (0.272)	0.310 (0.252)	0.303 (0.236)	0.504 (0.361)	0.444	0.238	0.289
rand. inf. p-val.						0.082	0.270	0.134
B. LATE estimates								
treated	0.402 (0.387)	0.505 (0.333)	0.367 (0.298)	0.359 (0.280)	0.664 (0.481)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-sales per worker, winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

on the management index introducing productivity measured as sales per worker as a control in equations (1.5) and (1.6). The estimates corresponding to this specification are presented appendix A.6. These estimates are similar in magnitude to those presented before in table 1.5. The significance of the estimates is also preserved, except in the case of the smallest bandwidth which is expected given that the number of estimated parameters has increased and the sample size has remained small. Appendix A.6 also reports similar estimates for the effect of the program on the management subindices and shows that they are also consistent with those presented above. These results suggest that the effect of the certification program on management practices is not operating through an increase in productivity.

1.6 Conclusion

In the last decade, several papers have shown that the use of certain modern management practices related to monitoring, target-setting, and provision of incentives to employees are important for firm growth. The lack of adoption of these practices has been regarded as one of the reasons behind the low firm productivity observed in developing countries. Until recently, the explanations for this lack of upgrading have mainly revolved around informational limitations: the fact that business-owners might not know about the existence of better practices or the profitability of adopting them. However, policy-feasible interventions addressing this constraint have shown mixed results regarding managerial upgrading and its effect on performance.

In this paper, I have tried to accomplish two goals. First, I have suggested that, in addition to informational constraints, adoption costs are also an important barrier to the adoption of better management practices. Second, I have proposed a policy-feasible way to reduce these costs: subsidizing the adoption of widely known process standards, such as ISO 9001. These standards require firms to standardize and document their internal processes. This requirement is different from adopting the modern management practices related with monitoring, target-setting, and incentives that have been associated with firm growth. However, I hypothesized that these standards might exhibit complementarity with at least some modern management practices. Under

this hypothesis, promoting the adoption of standards would also be useful to promote managerial upgrading beyond their scope. The existence of complementarity between different firm practices (not specifically process documentation and what I have called modern management practices) has been suggested previously (Ichniowski et al. (1997), Brynjolfsson and Milgrom (2013)). My findings suggest that this insight can be leveraged for policy design.

Using a newly created data set and a regression discontinuity design in a real policy setting, I have found evidence consistent with the complementarity hypothesis. The Peruvian certification program, which subsidized the adoption and certification of process standards, was successful at accomplishing its stated goal of promoting certification, and also promoted managerial improvement in other dimensions. The managerial improvement was concentrated in practices related with monitoring and target setting.

Three years after treatment assignment, I also found a positive, although statistically non-significant, increase in productivity measured as sales per worker. The magnitude of this change is similar to other findings in the literature. At the same time, I have found no evidence of changes in employment and sales. The lack of effect on these variables, as well as the lack of precision of the estimates for productivity, are consistent with previous studies, which have shown that the effects of interventions to promote better management materialize over several years and are hard to detect shortly after treatment (McKenzie and Woodruff (2014), Giorcelli (2019), Bruhn et al. (2018), Higuchi et al. (2019)).

Further research regarding the potential of process standards as a tool to improve management is warranted, given the advantages such standards offer with respect to other intervention techniques previously studied from an implementation point of view. Subsidies for adoption of standards are simple to manage because they can be treated like any other subsidy, a type of intervention governments have experience dealing with. Additionally, monitoring project completion is simple because the certification is easily verifiable. Moreover, the cost of this type of intervention is within the reach of a government or development agency. The program I studied granted, at most, USD 14,000. This value is similar to that of the Mexican intervention studied by Bruhn

et al. (2018) and significantly cheaper than those described by Bloom et al. (2013) and Giorcelli (2019). A natural next step in this research is to return to the field and collect a new round of data to determine whether the effect on productivity has consolidated, and whether improvements in performance have materialized.

Chapter 2: R&D Subsidies and Technology Adoption: Evidence from Peru¹

2.1 Introduction

R&D subsidies are one of the most popular policies to promote technological upgrading in developed and developing countries.² However, despite their popularity, existing rigorous studies on their effectiveness are scarce and focused on developed countries and high-tech firms. Critically, these studies have been focused in outcomes such as firm investment and growth, paying little attention to the impact of these intervention on the actual adoption of new technologies.³ Even though investment and growth are important in their own right, it is also necessary to know whether these effects (or lack thereof) are actually due to technological improvements.

In this paper, I assess the effect of a R&D subsidy program implemented in Peru. This program paid up to 50% of the cost of projects to develop new products and processes and was managed by *Innovate Perú*, a Peruvian government agency in charge of programs to promote productive development. I identify the impact of the intervention using a regression discontinuity strategy. To apply, firms had to submit a project which was evaluated and given a score by two reviewers. Only projects with scores above a certain threshold could get the subsidy. In a neighborhood of this cutoff value, the subsidy can be considered as randomly assigned.

¹I would like to thank professors Eric Verhoogen, Jonas Hjort, and Réka Juhász for their advice and support; professors Michael Best, Donald Davis, Christian Pop-Eleches, Amit Khandelwal, Suresh Naidu, David Weinstein, and Jack Willis for their valuable comments; Alejandro Afuso, Celia Cornejo, Rafael Castillo, Luis Mesías, Juan Manuel García, Santiago Cueto, and Eduardo Zegarra for their help during this project; and Andrea Quijano and Johnathon F. Cherry for their research assistance. I am also indebted to the *Grupo de Análisis para el Desarrollo* (GRADE) for providing a productive and stimulating work environment during my stay in Peru; and to *Innovate Perú* for providing information needed for this project. This research would not have been possible without the financial support provided by Columbia University's Department of Economics Development Colloquium and Program of Economic Research, the Center for Development Economics and Policy (CDEP), and the National Science Foundation (NSF, grant number 1829326). All views, opinions, and conclusions expressed are my own.

²Cirera and Maloney (2017)

³For studies in developed countries, see Bronzini and Iachini (2014), Bronzini and Piselli (2016), and Howell (2017). For studies in developing countries, see Castillo et al. (2019), the survey by Hall and Maffioli (2008), and some of the interventions reviewed in López-Acevedo and Tan (2011)

One of the main challenges when studying technology adoption is that this phenomenon is hard to observe. To address this issue, I ran a survey to collect information about technological upgrading among participant firms. More specifically, I asked the managers of these firms whether their companies had adopted the technology proposed in their projects.

In this paper, I find that the R&D subsidy program did not induce the adoption of new technologies by marginal firms. These firms carried out the projects described in their application, but failed to introduce the new products to the market or put the new processes in use. Qualitative evidence suggests that the firm managers found it too difficult and that the implementation challenges were initially underestimated. These results, together with those of the previous chapter, highlight the importance of considering firm capabilities and technological complexity when designing industrial policy. Beneficiaries of the R&D subsidy program received government support to implement complex technological projects, but were not able to adopt the proposed improvements. At the same time, beneficiaries of the certification program executed less complex projects and upgraded their technology.

This paper contributes to the literature on R&D subsidies by assessing, for the first time, whether they lead to the adoption of the products and process developed as part of the beneficiaries' projects. As mentioned before, existing papers on this field have not paid attention to the effect of these subsidies on the actual adoption of new technologies, despite that being their immediate goal. This study addresses this shortcoming by collecting data about technology adoption.

This study also contributes to the debate regarding the appropriateness of R&D subsidies in developing country contexts. Studies by Cirera and Maloney (2017)⁴ and Goñi and Maloney (2017) have pointed out that the lack of capabilities and complementary assets in developing country firms might prevent them from realizing the potential benefits of R&D subsidies. For example, firms might lack the capacity to identify profitable projects, or to introduce a new product to the market; or they may have a hard time finding the right type of human capital. A natural response to this critique, as suggested by Cirera and Maloney (2017)⁵ is to rebalance innovation policy priori-

⁴See chapter 3

⁵See chapter 4.

ties away from a focus centered around highly novel projects and towards the adoption of already tested technologies⁶, especially of established management practices as those described in Bloom and Van Reenen (2007) and Bloom and Van Reenen (2010). The results in this and the previous chapter are consistent with this observation.

2.2 Empirical strategy

2.2.1 The R&D subsidy program and the identification strategy

The operation of the R&D subsidy program is similar to the operation of the certification program. It is also managed by the same agency and subsidizes of up to 50% of the cost of projects to develop new products or processes, with a maximum of USD 90,000. Formal firms of any sector can apply. In this study, I use data from rounds 3 to 8 of the program (11/2011, 07/2012, 03/2013, 09/2013, 01/2014, 06/2014)⁷. They include 532 applications, of which 206 were funded.⁸ Table 2.1 presents descriptive statistics at the moment of applications. Applicants were, on average, almost 11 years old and were concentrated in the secondary and tertiary sectors. 78% of the applicants had annual sales of less than USD 2.1 million, which puts them in the category of micro and small enterprises. 21% and 34% of applications came from exporting and importing firms, respectively. 76% of applications were submitted by firms from the coastal regions, which are the most developed in the country.

To apply to the R&D program, firms submitted a document describing their project. Those projects went through a two-stage evaluation process. In the first stage, the project was evaluated by two reviewers, each of which assigned a score between 0 and 100 in nine criteria: innovative merit (15%), socio-economic potential (5%), positive externalities (15%), commercial viability (10%), technical quality (10%), environmental impact (5%), profitability (15%), relevance of co-financing (10%), and complementary factors (15%)⁹. The final score given by a reviewer is the

⁶Cirera and Maloney (2017) call this the "capacity escalator", the idea that different types of improvements are suitable for different stages of technological development.

⁷Information from rounds 1 and 2 was not available.

⁸Appendix B.1 provides details of the sample construction.

⁹Unlike the case of the certification program, these criteria were subdivided in different sub-criteria.

Table 2.1: Baseline descriptive statistics

	Total (1)	Nontreated (2)	Treated (3)
A. Sector			
primary	16.2%	17.8%	13.6%
secondary	36.1%	30.7%	44.7%
tertiary	47.7%	51.5%	41.7%
B. Size			
MSE	78.2%	81.5%	72.6%
medium	14.3%	12.5%	17.4%
large	7.5%	6.0%	10.0%
C. Trade			
exported	20.9%	18.1%	25.2%
imported	34.0%	28.5%	42.7%
D. Other characteristics			
coast	76.3%	73.3%	81.1%
age (yr)	10.9	10.0	12.2

Column (1) shows percentages corresponding to the total of received applications, column (2) refers only to nonsubsidized applications, and column (3) refers to subsidized applications. The size classification is based on sales. In U.S. dollars, micro and small firms (MSE) are those with sales between 0 and 2.1 million; medium firms are those with sales between 2.1 and 2.9 million. Large firms are those with sales of more than 2.9 million.

average of those scores. To pass the first evaluation stage, the project must pass the evaluation of at least one reviewer. For that to happen, the final score of that reviewer must be at least 70 and the score in the technical quality criterion must also be at least 70. Projects that pass the first evaluation stage move on to a second stage in which a committee reviews the project and either accepts or rejects it. I exploit the first stage of the evaluation process to carry out a regression discontinuity analysis. I use $s = \max(\min(s_1, s_{tec,1}), \min(s_2, s_{tec,2}))$ as the running variable. s_e is the final score of the evaluation of reviewer e ($e \in 1, 2$), and $s_{tec,e}$ is the score assigned on the technical quality criterion. Projects with $s \geq 70$ pass the first stage of the selection process and have a positive probability of getting the subsidy. Projects with $s \leq 70$ are rejected. Figure 2.1a shows the treatment probability as a function of the running variable. The treatment probability has a discontinuous increase of approximately 25% at the cutoff value. Figure 2.1b shows the density of the running variable.

As in the previous chapter, I estimate intent-to-treat effects (ITT) and local average treatment effects (LATE) using local linear regressions. The ITT estimate is given by the OLS estimate of β_1 in equations 2.1:

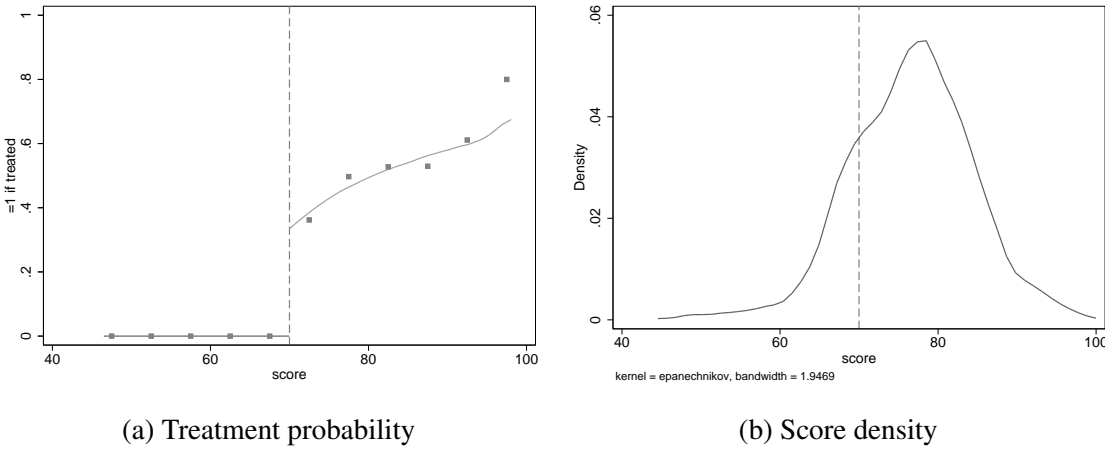
$$y = \beta_0 + \beta_1 \mathbf{1}(s \geq 70) + f(s) + u \quad (2.1)$$

where $\mathbf{1}(s \geq 70)$ is a dummy variable taking value one if the running variable s is to the right of the cutoff, $f(s)$ is polynomial of the running variable that allows for different slopes on different sides of the cutoff. The LATE estimate is given by the IV estimate of β_1 in equations 2.2

$$y = \beta_0 + \beta_1 \textit{treated} + f(s) + u \quad (2.2)$$

where *treated* is a dummy variable that takes value one for treated observation. I instrument this variable with $\mathbf{1}(s \geq 70)$. In both cases, I use observations restricted to different bandwidths around the cutoff and weighted according to a triangular kernel.

Figure 2.1: Treatment probability and score distribution



The graph on the left shows the estimates of a local linear regression of a dummy variable that takes value one if a subsidy for the project was granted and zero otherwise, on the running variable. The local linear regression uses triangular kernel and a bandwidth of 20 points. The graph on the right shows the density of the running variable.

2.2.2 Data sources

The data sources used in this study are similar to those from the previous chapter. The main source are the administrative records from *Innovate Perú*, the implementing agency. These data contain some baseline information, including the value of the running variable. I complement this data source with a survey I ran during the last quarter of 2018. In this survey, I asked the managers of applicant firms about the adoption of the technologies they had described in their projects. Gathering and using this information is the key contribution of this paper.

As in the case of the analysis in the previous chapter, I also use trade data provided by the Chamber of Peruvian Exporting Firms, and data manually collected from the tax authority's website.

2.2.3 Validity of the identification strategy

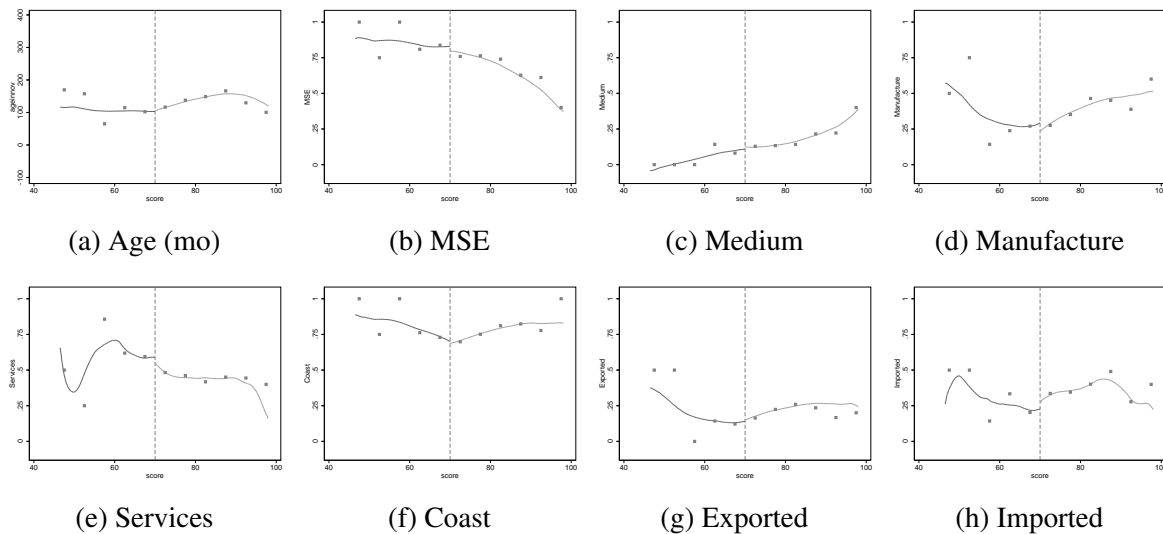
The identification assumption underlying the empirical analysis in this study is that the running variable has not been manipulated around the cutoff. Given the institutional context, manipulation of the running variable is unlikely due to two reasons. First, the identity of the reviewers is decided after the application has been submitted. This makes difficult the collusion between applicants and

reviewers. Second, the running variable depends on the scores of two reviewers. Again, this makes collusion even more difficult. Third, as in the case of the certification program studied in the first chapter, the implementing agency monitors the actual execution of the proposed project, reducing the profits of any potential collaboration between reviewers and applicant firms.

Even though the assumption of no manipulations cannot be directly tested, it is possible to check whether its observable implications hold in this specific context. More specifically, in absence of manipulation, the baseline covariates and the running variable density should not have any discontinuity at the cutoff.

Figure 2.2 shows local linear regressions of baseline covariates against the running variable and table 2.2 presents OLS estimates of changes at the cutoff using different bandwidths and polynomial degrees. None of these estimates is statistically significant, which is consistent with the absence of manipulation. Additionally, the McCrary (2008) test does not reject the null hypothesis of smoothness of the running variable density with a p-value of 0.907.

Figure 2.2: Continuity of baseline covariates



Each of these graphs shows the local linear regression of a baseline covariate on the running variable, using a triangular kernel and a bandwidth of 20 points. In graph (a), the outcome variable is the age of the firm in months. In graph (b), it is a dummy variable that indicates the firm is classified as micro and small firm. In graph (c), it is a dummy variable that indicates the firm is a medium firm. In graph (d), it is a dummy variable that indicates the firm is located in the coast region. In graph (e), it is a dummy variable that indicates the firm is classified in the manufacturing sector. In graph (f), it is a dummy variable that indicates the firm is classified in the service sector. In graph (g), it is a dummy variable that indicates that the firm had exported. In graph (h), it is a dummy variable that indicates that the firm had imported.

2.3 Effects of the R&D subsidy program

In this section I analyze the impact of the R&D program. More specifically, I show that the program did not accomplish its immediate goal of inducing the adoption of new products or processes by firm, at least at the cutoff. Moreover, consistent with this finding, I do not find evidence of impact on productivity.

2.3.1 Effects on technology adoption

The main finding of this study is that the R&D subsidy program did not induced the adoption of the proposed new technologies by firms. This is apparent in figure 2.3 that presents the local linear regression of a dummy variable taking value one is the technology proposed by the firm was adopted. As expected, the adoption probability is increasing in the running variable, which is a measure of project quality, but it does not increase discontinuously at the cutoff. Even though it is not possible to rule out that non-marginal firms are not being induced to adopt new technologies by the intervention, it is possible to say that there is no evidence of this happening among the marginal firms. The results in figure 2.2 are confirmed in table 2.3. Columns 1 - 3 in table 2.3 report OLS and 2SLS estimates of the ITT and LATE of interest using local linear regression for bandwidths 10, 20, and 30. Columns 4 and 5 report similar figures using a bandwidth of 30 and quadratic and cubic polynomials.


Why did the R&D program fail to promote technology adoption among marginal firms? With the available data, it is hard to make a definite statement. However, qualitative information collected in the survey and during interviews in the field suggest the that the difficulties were mainly related with the implementation of the improvements. In the survey, I asked treated firms reasons behind the lack of adoption. The most common reason was that, even though they had achieved what was initially described in the project, the managers realized that additional improvements were needed before actually taking the new products to the market and before starting to use the new processes. The second most common reason was that the cost of implementing the new tech-

Table 2.2: Continuity of baseline covariates

	(1) age	(2) mse	(3) medium	(4) manif.	(5) services	(6) coast	(7) exported	(8) imported
A. bw=10								
right	10.483 (21.087)	-0.087 (0.090)	0.051 (0.081)	-0.097 (0.106)	-0.037 (0.122)	0.012 (0.110)	-0.008 (0.084)	0.042 (0.112)
obs	339	339	339	339	339	339	339	339
B. bw=20								
right	2.383 (17.182)	-0.032 (0.072)	0.017 (0.061)	-0.058 (0.086)	-0.058 (0.098)	-0.012 (0.090)	0.007 (0.068)	0.077 (0.088)
obs	510	510	510	510	510	510	510	510
C. bw=30								
right	7.113 (16.276)	-0.012 (0.067)	0.000 (0.055)	-0.009 (0.079)	-0.091 (0.090)	-0.012 (0.083)	0.038 (0.064)	0.095 (0.080)
obs	532	532	532	532	532	532	532	532
D. bw=30, quadratic polynomial								
right	1.155 (21.310)	-0.079 (0.088)	0.049 (0.078)	-0.120 (0.108)	-0.023 (0.122)	-0.006 (0.110)	-0.028 (0.085)	0.053 (0.113)
obs	532	532	532	532	532	532	532	532
E. bw=30, cubic polynomial								
right	21.458 (27.032)	-0.141 (0.111)	0.078 (0.100)	-0.142 (0.135)	-0.012 (0.154)	0.033 (0.135)	-0.017 (0.102)	0.030 (0.142)
obs	532	532	532	532	532	532	532	532

***, **, * indicate significance at 1%, 5%, and 10%. The values reported in this table are the OLS estimates of β_1 in equation (2.1) taking the variable indicated at the top as the dependent variable. Values in parenthesis are standard errors clustered at the level of the firm. Panels A - C report estimates using a bandwidth of 10, 20, and 30 points around the cutoff, triangular kernel, and a linear polynomial of the running variable. Panels D and E report similar estimates using a bandwidth of 30 and quadratic polynomial and cubic polynomials of the running variable.

Figure 2.3: Effect on technology adoption



This graph shows a local linear regression of a dummy variable that takes value one if the new technology was adopted and zero otherwise, on the running variable using a triangular kernel and a bandwidth of 20 points. The dots represent the mean of the outcome variable in 5-point bins. In the case of new products, adoption means taking the product to market. In the case of new processes, adoption means using the new process in the regular operation of the firm.

nology was too high. Table 2.4 summarizes these responses.

The survey did not ask whether the need for additional improvements was foreseen before carrying out the project, but the information gathered through interviews suggest it was not. For example, consider the following excerpt from an interview with a manager of a treated firm that carried out a project to develop a new product.

Manager: "The product have not been sold. (...) The product ended up being more expensive than we expected, although the product itself is better than expected. We need to correct that. I just went China to see how this could be solved. (...) We did not foresee this problem. It was surprising."

The relevance of implementation costs was also frequently mentioned. For example, in the following interview excerpt, the manager of a software company indicates that introducing a new product to market is difficult due to marketing costs even though the new product is technically sound.

Interviewer: "Has this product been sold? [the one developed with support of the R&D program]."

Table 2.3: Effect on technology adoption

	bw=10 (1)	bw=20 (2)	bw=30 (3)	bw=30 (4)	bw=30 (5)
A. ITT estimates					
right	-0.057 (0.114)	0.078 (0.102)	0.091 (0.100)	-0.005 (0.124)	-0.123 (0.152)
B. LATE estimates					
treated	-0.218 (0.450)	0.233 (0.303)	0.257 (0.278)	-0.018 (0.450)	-0.528 (0.714)
C. First stage (dep. var.=treated)					
right	0.262*** (0.079)	0.333*** (0.065)	0.354*** (0.061)	0.277*** (0.078)	0.233** (0.094)
F-stat	10.917	26.661	34.051	12.470	6.132
obs	219	322	332	332	332

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a dummy variable that indicates if the proposed technology was adopted. In the case of new products, adoption means taking the product to market. In the case of new processes, adoption means using the new process in the regular operation of the firm. Panel A reports OLS estimates of β_1 in equation (2.1). Panel B reports the IV estimates of β_1 in equation (2.2). Panel C reports the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, triangular kernel, and linear polynomial. Columns 4 and 5 report similar estimates using a bandwidth of 30 and quadratic and cubic polynomials.

Manager: "Not yet. (...) The product, at this moment, is fully functional. (...) Last year we presented it to big companies, but they were not very enthusiast. They always ask 'who else use it?' (...) [To start selling it], we would need to have a successful first customer to generate a demonstration effect. It's a catch-22 situation. An adviser even told me that, I might need to give it for free at the beginning. Of course, that is quite expensive."

Table 2.4: Reasons for not adopting new technology


	% of non-adopters
Additional improvements needed	52.8
Implementation cost too high	35.8
Not a relevant improvement	15.1
Lack of time for implementation	5.7
Could not finish the project	22.6
Other	17.0

This table shows the reasons for not adopting the proposed technology. The percentage is calculated based on the total number of projects that that did not lead to technological adoption. Managers could indicate more than one reason, hence, the values do not need to add up to 100%.

2.3.2 Effects on productivity and performance

Consistent with the lack of technology adoption, I also find no significant effects of the R&D subsidy program on productivity measured as log-sales per worker. This result is apparent in figure 2.4a, which shows the results of a local linear regression of log-sales per worker on the running variable. Estimates of this effect are presented in table 2.5. Columns 1 - 3 show estimates using local linear regression and different bandwidths. Columns 4 and 5 report similar figures using a bandwidth of 30 and quadratic and cubic polynomials.

Figure 2.4: Effect on productivity and performance



Each of these graphs shows the estimates of a local linear regression of the dependent variable on the running variable using a triangular kernel and a bandwidth of 20 points. The dots represent the mean of the outcome variable in 5-point bins. In the graph on the left, the dependent variable is log-sales per worker. In the graph in the center, the outcome variable is log-employment (headcount), and in the graph on the right, the outcome variable is log-sales. Dependent variables were winsorized at the 10th and 90th percentiles to control for outliers.

Despite the lack of improvement in terms of productivity, I do find a positive and significant ef-

Table 2.5: Effect on log-sales per worker

	bw=10 (1)	bw=20 (2)	bw=30 (3)	bw=30 (4)	bw=30 (5)
A. ITT estimates					
right	0.002 (0.303)	0.083 (0.259)	0.141 (0.250)	-0.022 (0.318)	0.330 (0.404)
B. LATE estimates					
treated	0.007 (1.128)	0.245 (0.759)	0.392 (0.693)	-0.079 (1.121)	1.370 (1.762)
C. First stage (dep. var.=treated)					
right	0.269*** (0.080)	0.339*** (0.065)	0.359*** (0.061)	0.284*** (0.079)	0.241** (0.095)
F-stat	11.240	27.073	34.466	12.845	6.418
obs	219	322	332	332	332

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-sales per worker. Panel A reports OLS estimates of β_1 in equation (2.1). Panel B reports the IV estimates of β_1 in equation (2.2). Panel C reports the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, triangular kernel, and linear polynomial. Columns 4 and 5 report similar estimates using a bandwidth of 30 and quadratic and cubic polynomials.

fect on sales and employment, although these effects are impressively estimated. Table 2.6 reports and increase in employment. Columns 1 - 3 present OLS and 2SLS estimates of the ITT and LATE using local linear regression and different bandwidths. Columns 4 - 5 show similar estimates using quadratic and cubic polynomials. The magnitude of the ITT estimate is not quite stable across different bandwidths and oscillates around an increase of 81% in the number of employees¹⁰. At first, this seems as a strikingly large number. However, it is necessary to take into account the fact that the estimates are not very precise. For example, the 90% confidence interval for the ITT estimate in column 2 of table 2.6 is (0.13, 1.06), which represents and increase in the range of (13%, 188%)¹¹,

¹⁰ $80.9=(\exp(0.593)-1)*100$

¹¹The lower value comes from $0.13=0.593-1.645*0.281$ and $13.88=(\exp(.13)-1)*100$. The upper value comes from $1.06=0.593+1.645*0.281$ and $188.63=(\exp(1.06)-1)*100$.

which includes only positive values but is not very informative about their magnitude. Figure 2.4c provides the graphical representation of these results.

Table 2.7 reports similar estimates for sales. As in the case of employment, the magnitude of the estimates is large, but not very surprising once the range of the confidence intervals are considered. For example, the ITT estimate in column 2 implies an increase of 66%, but the 90% confidence interval of this estimate is (3%, 170%).

Even though the magnitude of the changes is not precisely estimated, they suggest an increase in sales and employment. Why did this happen despite the lack adoption of the new technologies? One possible explanation is that the subsidy eased some financial constraints that the firms were facing in their regular operations. Unlike the other R&D subsidy interventions, like the one studied by Howell (2017), *Innovate Perú* supervises that the money is spent in the activities indicated in the project. However, despite these efforts, it is possible that once the project's materials and personnel have been bought and hired, they are also partly use for other activities. For example, an engineer hired to work in the development of a new product could also devote part of her time to help with the regular operations of the firm. In fact, a situation like this was narrated in an interview by the manager of a company that received the subsidy. The following is an excerpt of that conversation:

Manager: "When you hire a highly trained worker [for the R&D project], they usually finish their duties very fast because they are very well trained and very motivated. And you end up with a worker that is having free time while others [not related with the R&D project] are still working on their assignments, and they comment 'oh, this person is working in only one thing, and now is having free time'. (...) [In those cases we] ended up asking the highly trained worker to help with other things."

Table 2.6: Effect on log-employment

	bw=10 (1)	bw=20 (2)	bw=30 (3)	bw=30 (4)	bw=30 (5)
A. ITT estimates					
right	0.882** (0.350)	0.593** (0.281)	0.508* (0.263)	0.948** (0.369)	1.243** (0.489)
B. LATE estimates					
treated	3.372** (1.643)	1.780** (0.904)	1.435* (0.773)	3.428** (1.624)	5.327* (2.889)
C. First stage (dep. var.=treated)					
right	0.262*** (0.079)	0.333*** (0.065)	0.354*** (0.061)	0.277*** (0.078)	0.233** (0.094)
F-stat	10.917	26.661	34.051	12.470	6.132
obs	219	322	332	332	332

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-employment (headcount). Panel A reports OLS estimates of β_1 in equation (2.1). Panel B reports the IV estimates of β_1 in equation (2.2). Panel C reports the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, triangular kernel, and linear polynomial. Columns 4 and 5 report similar estimates using a bandwidth of 30 and quadratic and cubic polynomials.

2.4 Conclusion

The popularity of R&D subsidies as a policy to promote firm-level technology adoption relies on the idea that financial constraints are an important barrier to upgrading. Indeed, some studies have found evidence consistent with this idea. However, those studies have been mainly focused on developed countries and high-tech firms. This rises the question of whether these results are also valid in a developing country context. In developing countries firms do face financial restrictions, but these are not necessarily the main obstacles they need to overcome. Human capital shortages, lack of expert advice, lack of firm capacity, among others, are obstacles that might be particularly acute among firms in low-income countries.

Table 2.7: Effect on log-sales

	bw=10 (1)	bw=20 (2)	bw=30 (3)	bw=30 (4)	bw=30 (5)
A. ITT estimates					
right	0.581* (0.340)	0.511* (0.293)	0.492* (0.280)	0.672* (0.365)	1.090** (0.468)
B. LATE estimates					
treated	2.219 (1.408)	1.534* (0.901)	1.388* (0.801)	2.432* (1.438)	4.674* (2.602)
C. First stage (dep. var.=treated)					
right	0.262*** (0.079)	0.333*** (0.065)	0.354*** (0.061)	0.277*** (0.078)	0.233** (0.094)
F-stat	10.917	26.661	34.051	12.470	6.132
obs	219	322	332	332	332

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-sales. Panel A reports OLS estimates of β_1 in equation (2.1). Panel B reports the IV estimates of β_1 in equation (2.2). Panel C reports the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, triangular kernel, and linear polynomial. Columns 4 and 5 report similar estimates using a bandwidth of 30 and quadratic and cubic polynomials.

In this paper, I have studied for the first time whether R&D subsidies induce the adoption of new technologies in firms in developing countries. To do this, I focused on the case of a R&D subsidy program in Peru. This intervention paid up to 50% of the cost of R&D projects and the assignment mechanism allows the use of a regression discontinuity to causally link the intervention with the findings.

The main result is that the program did not lead to the adoption of new technologies by the beneficiary firms and did not produce an increase in productivity. This suggests that, indeed, other restrictions beyond financial constraints are preventing technological upgrading in firms in developing countries. This result calls into questions the efficacy of this type of policy to promote innovation in developing countries.

Additionally, I also find some evidence of increases in employment and sales due to the subsidy. Even though these findings need to be interpreted with caution as the estimates are very imprecise, they suggest that firms are using part of the subsidy to pay for needs not related with innovation, for example, highly qualified employees or capital. In other words, firms might be using the resources of the program to solve financial constraints not related with technology upgrading.

Together with the findings of the previous chapter, the results suggest that policy-makers in developing countries should take into account the initial level of firms' capacity when designing policy to promote firm upgrading. More specifically, it is important to understand the obstacles to upgrading that the productive units face. If firms lack the capacity to successfully carry out high-novelty innovation projects, subsidizing these projects will not encourage upgrading. In these contexts, it might be more advisable to promote the adoption of less novel and less complex improvements, like better managerial practices, rather than more radical forms of innovation.

Chapter 3: Exports and Wage Premia: Evidence from Mexican Employer-Employee Data (with Judith A. Frías[†], David S. Kaplan[‡], and Eric Verhoogen[§])¹

3.1 Introduction

In textbook models of perfectly competitive labor markets, the wage a worker earns does not depend on the particular firm in which she is employed; she would receive the same wage in any firm that would employ her in equilibrium. A long-established and recently resurgent literature has challenged this view. Following the seminal work of Abowd et al. (1999) (hereafter AKM), a number of recent studies have fit simple models with individual and firm fixed effects and have found that the firm effects account for a substantial fraction of overall wage variation (Card et al. (2013), Card et al. (2015), Abowd et al. (2018), Alvarez et al. (2018), Barth et al. (2018), Sorkin (2018), Song et al. (2019)). Firms appear to pay different wage premia, defined as wages above what individuals would earn elsewhere on the labor market, and understanding how those premia evolve appears to be important for understanding the broader evolution of wage structures.

Despite the recent advances, there has been relatively little empirical work on *why* firms pay

¹We would like to thank, without implicating, Rita Almeida, Paula Bustos, David Card, Bo Honoré, Leonardo Iacovone, Wojciech Kopczuk, Thomas Lemieux, Arik Levinson, Steven Lich-Tyler, Leigh Linden, Osmel Manzano, Marc Melitz, Marcelo Moreira, Marc-Andreas Muendler, Daniel Parent, Cristian Pop-Eleches, Jesse Rothstein, Matthias Schuendeln, Matthew Slaughter, John Strauss, John Sutton, Catherine Thomas, Miguel Urquiola, Till von Wachter, and many seminar participants for helpful comments. We would also like to thank Gerardo Durand, José Antonio Mejía, Jorge Alberto Reyes, Gerardo Leyva, Abigail Durán, Adriana Ramírez, and Gabriel Romero of INEGI for help with the first versions of the paper, and Natalia Volkow and other staff of the INEGI Microdata Laboratory for assistance after a long hiatus due in part to a change in data-access regime. All analysis of confidential data was carried out on-site and in accordance with Mexican confidentiality laws. Guru Sethupathy, Hitoshi Shigeoka, Kensuke Teshima and Matt Wai-Poi provided excellent research assistance. Verhoogen is indebted to the National Science Foundation (SES-0721068) for funding. The views expressed are the opinions of the authors and do not reflect the official positions of the Instituto Mexicano del Seguro Social, the Secretaría de Hacienda y Crédito Público, or the Inter-American Development Bank. [†] Consejo de la Judicatura Federal. [‡] Inter-American Development Bank. [§]Columbia University.

wage premia. What leads some firms to pay more than others to similar workers? Much of the recent literature has aimed to characterize the extent of rent-sharing with different groups of workers or the contribution of firms to overall wage inequality, but not to identify the causal determinants of wage premia.² The recent review by Card et al. (2018) concludes with a call for research to fill this gap.³

In this paper, we draw on a combination of employer-employee and longitudinal plant data from Mexico to investigate one possible determinant of wage premia: firms' engagement in export markets. It is well documented that exporting firms tend to pay higher wages on average, within narrow industries.⁴ There is growing evidence that this relationship is causal — that exogenous increases in exporting lead firms to pay higher wages on average Verhoogen (2008), Álvarez and López (2009), Kandilov (2009), Bustos (2011a), Brambilla et al. (2012). But it remains unclear to what extent these plant-level results reflect changes in wage premia as opposed to changes in workforce composition. Several studies (discussed briefly below) have examined the relationship between exporting and wages in employer-employee data, but to our knowledge none has employed quasi-experimental variation to estimate the effect of exporting on wage premia as we have defined them.

This paper proceeds in two steps. First, following Card et al. (2013) (hereafter CHK), we fit simple AKM-type models in different periods and decompose plant-level wages into a plant component, which we interpret as a plant-specific average wage premium, and a person component reflecting individual characteristics, including individual effects capturing time-invariant ability.⁵

²See for example Card et al. (2015) on differences in rent-sharing for different groups, and Card et al. (2013), Barth et al. (2016), citeBarthDavisFreeman2018, and Song et al. (2019) on characterizing the role of firms in overall wage inequality. Lazear and Shaw (2018) provide a useful overview.

³“The field continues to rely almost exclusively on observational studies predicated on plausible, but ultimately debatable, identifying assumptions. More research is needed applying (quasi-)experimental research designs that convincingly tease out the mechanisms through which firm shocks are transmitted to workers.” Card et al. (2018)

⁴The first papers to document this pattern are Bernard and Jensen (1995), Bernard and Jensen (1999). Brambilla et al. (2017) review the literature and present evidence from 61 developing countries of a robust positive relationship between exporting and firm-level wages.

⁵In Mexican manufacturing, the vast majority of firms are single-establishment. For instance, the 2009 Economic Census in Mexico reports that more than 75% of manufacturing firms with 51 or more employees and more than 90% of manufacturing firms with 11 or more employees are single-establishment. The data we use are at the plant level and unfortunately we do not know which firm a given plant is part of. Hereafter we will treat plants as decision-makers.

Because of data constraints, discussed below, we focus on the periods 1992-1994, 1996-1998, and 2000-2002 in our baseline specification. We conduct diagnostic exercises similar to CHK and find that the data appear to be well described by a model with additive individual and plant fixed effects within several-year periods.

Second, following Verhoogen (2008), we use the late-1994 peso devaluation as a source of exogenous variation in the incentive to export, to identify the effect of exporting on the estimated wage premia. Verhoogen (2008) motivated this approach with a model in which firms in a developing country are heterogeneous in entrepreneurial ability (or “capability”), sell higher-quality varieties when exporting in order to appeal to richer consumers in the export market, and pay higher wages when producing higher-quality goods. A devaluation leads more-capable firms to increase export share and hence to raise average product quality, average wages, and (possibly) wage premia more than less-capable firms in the same industry.⁶ In this paper, we rely on the same reduced-form prediction for the relationship between exporting and wage outcomes without taking a stand on the underlying theoretical mechanism. (We discuss several possible mechanisms below.) As in the earlier paper, we use plant size as the primary proxy for the underlying plant heterogeneity and interpret the interaction of initial size and an indicator for the devaluation period as a proxy for the differential inducement to export generated by the peso devaluation. We estimate an instrumental-variables (IV) regression of within-plant changes in wage premia on changes in export share instrumented by this interaction. We also present the corresponding reduced-form regressions of the changes in wage premia on the interaction term and other covariates.

Our main finding is that exporting has a robust, positive, statistically significant effect on wage premia. The magnitude is economically significant: a one percentage point increase in a plant’s export share (on an initial mean of 4% overall, or 14% conditional on exporting) is associated with a 3% increase in wage premia at the plant. Our estimates suggest that essentially all of the effect of exporting on average wages at the plant level is explained by the increase in wage premia, rather than by changes in workforce composition, at least for the medium-term time frame we are able to

⁶Section III.B in Verhoogen (2008) suggested that the plant-level wage effects could reflect changes either in wage premia or in workforce composition.

study. These results are robust to using total factor productivity (TFP) and predicted export share as alternative proxies for plant capability and to including an earlier pre-crisis period (1988-1990) in the analysis. We conclude that exporting is an important determinant of wage premia among Mexican manufacturing firms.

Two broader implications of these findings seem particularly salient. First, for the recently resurgent literature on the role of firms in wage-setting, they underline the causal role of product markets in shaping the evolution of wage distributions through their effect on firms' wage policies. Product-market shocks are potentially related to, but are conceptually distinct from, shocks to productivity (i.e. technical efficiency), which have been the focus of much of the rent-sharing literature.⁷ Second, the paper contributes to a growing body of evidence about how firms' responses to trade shocks contribute to wage inequality. Previous work has argued that trade liberalization can increase wage dispersion across firms within industries in developing countries, as larger, more-productive firms, which already tend to be higher-wage, take greater advantage of export opportunities, which leads them to raise wages further (Verhoogen (2008), Helpman et al. (2017)). This paper finds that the increase in across-plant dispersion in our setting is largely driven by plants' wage policies, rather than sorting on worker ability. The limited impact on workforce composition we find suggests that the export shock had limited effects on the general-equilibrium return to skill, at least in the short to medium run. The trade-induced increase in wage dispersion may be confined to manufacturing (and other tradable) sectors.

As mentioned above, a number of theoretical mechanisms are consistent with a causal relationship between exports and wage premia. Differential quality upgrading is one possibility, if producing higher-quality products for richer consumers in the export market requires a particularly motivated workforce, which in turn requires paying high efficiency wages Verhoogen (2008). A second possibility is fair-wage concerns: if the exogenous increase in exporting is associated with an increase in profitability (as in many Melitz (2003)-type models), firms may share profits in order to induce workers to reciprocate with effort Akerlof and Yellen (1990), Egger and Kre-

⁷See e.g. Van Reenen (1996), Kline et al. (2019), and the review by Card et al. (2018). These papers are discussed further below.

ickemeier (2009), Egger and Kreickemeier (2012), Amiti and Davis (2012). A third possibility is that the labor market is characterized by search frictions and firm-worker bargaining; several papers have posited mechanisms through which the bargained wage may increase with exports in the context of heterogeneous-firm models.⁸ A fourth possibility is that managers share rents with workers in order to ensure themselves a quiet life, rather than maximizing profits Bertrand and Mullainathan (2003).⁹ These mechanisms (and others) are often classified under the general heading of “rent-sharing.” In this paper, we do not take a stand on precisely which mechanism is correct. Nevertheless, we believe that our results carry an important implication for trade theory, by providing evidence against the traditional assumption of perfectly competitive, frictionless labor markets, in which any relationship between exporting and wages at the firm level is explained only by the sorting of workers by skill.

In addition to the research cited above, this paper is related to several streams of literature. Early papers investigating the effects of exogenous shocks on rent-sharing at the firm level include Abowd and Lemieux (1993) and Van Reenen (1996). In firm-level data from collective bargaining contracts in Canada, Abowd and Lemieux (1993) employed industry-level import and export prices as instruments for quasi-rents per worker (a measure of profitability when labor is valued at its alternative wage), and found that greater quasi-rents led to higher wages. In panel data on British firms, Van Reenen (1996) related major innovations, which are arguably exogenous, to wage changes at the firm level, and found that quasi-rents due to innovation were passed through to higher wages. In both cases, the researchers lacked data on individual workers and hence were not able to definitively answer questions about whether the wage changes reflect changes in wage premia or changes in skill composition.¹⁰

⁸In Helpman et al. (2010), firms employ a fixed-cost screening technology; as scale increases with increased exports, firms screen more intensively, which increases the revenue per worker to be bargained over, which in turn increases wages. In Coşar et al. (2016), firms face labor adjustment costs and the benefit of filling a position is greater for expanding firms than shrinking firms, leading the bargained wage to be higher in such firms; to the extent that increased exporting leads firms to expand, wages will also increase. See also Davidson et al. (2008), Felbermayr et al. (2011) Fajgelbaum (2016) and Bellon (2017).

⁹This list of possible mechanisms is not exhaustive. Another is simply that workers are unionized and are able to bargain for a share of firm profits.

¹⁰The current paper is also related to the earlier literature on *industry* wage differentials (Krueger and Summers (1988), Katz and Summers (1989)) and in particular to Dickens and Katz (1987), which examined correlations between

A recent paper by Kline et al. (2019) also examines innovation and rent-sharing, with employer-employee data and rich patent information from the U.S. Under the identifying assumption that the U.S. Patent Office's initial decision on a patent application is as good as random, conditional on observable characteristics of the application and firm, the authors find that each patent-induced additional dollar of operating surplus yields a 29-cent increase in a firm's wage bill. Using a matching estimator in Finnish data, Aghion et al. find that patents are associated with higher wages for co-workers of inventors. They also make the important point that the higher wages for co-workers may reflect compensation for efforts to operationalize and commercialize an invention, rather than a pure shock to profitability or productivity.¹¹ The current paper is complementary to these papers and focused on a very different type of shock — a shock to product-market conditions rather than to the technical capabilities of a firm.

As noted above, this paper is also related to a number of papers that have related exporting to wages in employer-employee data, including Schank et al. (2007), Munch and Skaksen (2008), Davidson et al. (2014), Baumgarten (2013), Klein et al. (2013), Irarrazabal et al. (2013), Hummels et al. (2014), Krishna et al. (2014), Araújo and Paz (2014), Macis and Schivardi (2016), Helpman et al. (2017), Barth et al. (2018), and Garin and Silvério (2018).¹² To our knowledge, this paper is the only one to relate quasi-experimental variation in exports to differences in wage premia defined as wages above what a given individual would earn elsewhere. Schank et al. (2007), Munch and Skaksen (2008), Hummels et al. (2014), Araújo and Paz (2014), and Krishna et al. (2014) estimate models with job-spell effects, which absorb wage differences across firms and hence do not address the question of why a given individual earns a higher wage at one firm than another.¹³ Klein et al.

various industry characteristics and industry wage differentials.

¹¹In this sense, it is not clear that a patent can be considered a pure shock to firm productivity or profitability; commercializing an invention may require changes in how a firm allocates workers' time or organizes production.

¹²An earlier version of the current paper (Frías et al. (2009)) estimated a model that allowed more flexibly for changes over time in the return to individual ability, using an approach pioneered by Holtz-Eakin et al. (1988) and found results broadly similar to those reported here. In this version, delayed in part by a change in data-access regime, we focus on simple AKM-type specifications estimated in separate periods, because the method is arguably more transparent and is more comparable to the literature as it has developed more recently. Many of the papers cited above refer to the earlier version of this paper. In a related short paper, Frías et al. (2012) show (without controlling for individual effects) that exports increase within-plant wage inequality.

¹³The paper by Hummels et al. (2014), perhaps the highest-profile in this set, focuses on the wage effects of offshoring in Denmark but also examines the wage effects of exporting. As an instrument for exports, the authors

(2013), Macis and Schivardi (2016), and Barth et al. (2018) relate firm (or plant) effect estimates to exports (and other firm-level variables), but do not employ a design-based strategy to isolate the effect of exogenous variation in exports.¹⁴ A recent working paper by Garin and Silvério (2018) constructs firm-specific export shocks due to differential changes in the import demands of Portuguese firms' export destinations and relates the shocks to the wage changes of workers initially employed in the corresponding firm, but does not consider the change in AKM-type firm effects as an outcome.¹⁵

The next section describes our econometric strategy in more detail. Section 3.3 describes the data and briefly provides background on the peso crisis. Section 3.4 presents results from the first step of our econometric procedure, estimating wage premia at the plant level. Section 3.5 presents results from the second step, estimating the effect of exports on the estimated wage premia. Section 3.6 concludes.

3.2 Econometric Strategy

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

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

construct a firm-specific weighted average of imports of particular goods by a firm's trading partners, using the firm's initial export shares as weights. As the authors acknowledge, this strategy is subject to the concern that demand shocks are correlated in Denmark and the trading partners. In addition, the paper estimates the effect of exporting only within job spells, as noted above.

¹⁴Macis and Schivardi (2016) attempt to implement a design similar to Verhoogen (2008) and this paper using the 1992 devaluation of the Italian lira, but they find no evidence of a differential effect of the devaluation by firm size. In the end, they regress estimates of wage premia on current export share, without instrumenting, letting the coefficient on export share differ pre- and post-devaluation.

¹⁵Garin and Silvério (2018)'s strategy for constructing export shocks is similar to that used by Bastos et al. (2018), also in Portuguese data, but instead of constructing firm-specific averages of real-exchange-rate movements, the authors construct firm-specific averages of import-demand changes due to the 2008-09 recession, in the spirit of Hummels et al. (2014). As noted above (footnote 13), correlation of demand shocks between the home country and trading partners (in this case, correlation of the demand responses to the recession) would violate the required exclusion restriction.

by:

$$w_{it} = \alpha_i + \psi_{J(i,t)} + X'_{it}\beta + \varepsilon_{it} \quad (3.1)$$

where α_i is a time-invariant individual effect; $J(i, t)$ indicates the plant j in which person i is employed in year t ; $\psi_{J(i,t)}$ is the corresponding plant effect; X_{it} is a vector of time-varying observables; and ε_{it} is a mean-zero error term. This specification follows AKM and CHK.¹⁶

We interpret the individual effect α_i as portable ability, invariant during a given several-year period, compensated equally in all firms. The plant effect $\psi_{J(i,t)}$ reflects a premium (or discount) paid by plant j to all employees. In X_{it} , we include tenure and its square, year effects, and a polynomial in age. Including a linear term for age is not possible because it would be collinear with the person and year effects. Following CHK and Card et al. (2013), we drop the linear term and include the square and cube of recentered age (age - 40). We define the “person” component to be the individual effect plus the contribution of the other observables: $s_{it} \equiv \alpha_i + X'_{it}\beta$. We will use the term “average person component” to refer to the plant-level mean of the (individual-level) person component. (We will use the terms plant effect and plant component interchangeably.)

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

$$E(\varepsilon_{it} | X_{i1} \dots X_{iT}, \psi_1 \dots \psi_J, \alpha_i) = 0 \quad (3.2)$$

In the employer-employee literature, (3.2) is referred to as a *conditional random mobility* assumption, since it requires that, conditional on observables, an individual’s current-period idiosyncratic shock is uncorrelated with which plant she is employed in.¹⁷ While (3.2) rules out sorting across plants on the basis of contemporaneous shocks to individual ability, it allows for many salient forms of sorting. For instance, as CHK point out, it allows for sorting on time-invariant ability

¹⁶Card et al. (2018) provide a theoretical justification for this specification, in a model in which workers have idiosyncratic attachments to particular plants and either different worker skill groups are substitutes in the production function or the plants face the same supply elasticity from these different groups.

¹⁷In the terminology of the panel-data literature, (3.2) requires the covariates to be strictly exogenous, i.e. uncorrelated with all past and future error terms. Intuitively, after a within transformation, the transformed covariates and errors contain within-individual averages across years, which must be uncorrelated.

(e.g. higher turnover among lower-ability workers) as well as for the possibility that workers are more likely on average to move from low- to high-wage plants than vice-versa (since we condition on $\psi_1 \dots \psi_J$). In Section 3.4 below we present evidence, following CHK, that the model (3.1)-(3.2) appears to summarize well the wage and mobility patterns we observe.

We face an important choice about how to define the several-year periods in which to estimate this model. As discussed in Section 3.3 below, the employer-employee data are available for 1985-2005, but the plant panel with export information is available on a consistent basis only for 1993-2003. The peso devaluation occurred in December 1994, and the crisis played out over the next several months. We need the first period to be clearly pre-devaluation, but we also need information on export status, the earliest of which is from 1993. Our preferred solution is to use 1992-1994 as the pre-crisis period. We do not use the crisis year 1995. To maintain equally spaced periods, we use 1996-1998 as the first post-crisis period and 2000-2002 as the second post-crisis period. We will refer to 1992-1994, 1996-1998, and 2000-2002 as periods 1, 2, and 3 respectively.¹⁸

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

Under assumption (3.2), for plant $j = j'$ in year $t = t'$ the expected plant-level wage can be expressed as the sum of two components:

$$\begin{aligned} E(w_{it}|j = j', t = t') &= \psi_{j'} + E(\alpha_i + X'_{it}\beta|j = j', t = t') \\ &= \psi_{j'} + E(s_{it}|j = j', t = t') \end{aligned} \tag{3.3}$$

where the conditioning on $X_{i1} \dots X_{iT}, \psi_1 \dots \psi_J, \alpha_i$ is omitted but should be understood. The first

¹⁸In a robustness check using the employer-employee data alone, we will also present estimates for 1988-1990, which we refer to as period 0.

term is the plant component, the wage premium paid by the plant. The second term is the expected person component, which captures average workforce skill. To generate sample analogues of these components, we fit model (3.1) and recover the estimated plant effects and a plant-year-level average of the estimated person components, which we refer to as the average person component. Computationally, we estimate (3.1) using the pre-conditioned conjugate gradient (PCG) algorithm in Matlab, following CHK. We estimate the model separately by (3-year) period.

It is worth emphasizing that although the individual effects, α_i , are assumed to be constant within 3-year periods, they are allowed to vary across periods. Hence if the peso devaluation generated an increase in the general-equilibrium return to skill in the economy,

Once we have recovered these estimates, the next step is to determine the effect of exporting on the plant and average person components. As mentioned above, our approach is motivated by the Melitz (2003)-type theoretical framework of Verhoogen (2008). In this framework, firms differ in an underlying “capability” parameter, λ , there is a fixed cost of exporting, and firm capability and worker effort are complements in determining product quality. In equilibrium, more-capable firms are larger, pay higher wages, and produce higher-quality goods in cross-section. In response to an exogenous inducement to export such as a devaluation, more-capable firms increase exports, increase average product quality, and raise wages relative to less-capable firms in the same industry. If producing high-quality products requires especially motivated workers, which in turn requires paying high efficiency wages, then the increase in exports will be accompanied by an increase in wage premia. As noted in the introduction, subsequent papers have developed heterogeneous-firm models that emphasize different mechanisms but that generate similar predictions for the relationship between firm size, exporting and wages in response to trade shocks, and a subset of these predict similar effects on wage premia.

Motivated by this framework, we adopt an econometric specification of the following form:

$$\Delta y_{jp} = \theta \Delta e_{jp} + \gamma \widehat{\lambda}_{jp-1} + \xi_{kp} + \psi_{rp} + u_{jp} \quad (3.4)$$

where j indexes establishments, p indexes periods; Δy_{jp} is the change in an outcome variable (e.g. wages, the plant component or average person component from our AKM-type model, or the capital-labor ratio) between period $p - 1$ and period p ; Δe_{jp} is the change in export share between period $p - 1$ and p ; $\widehat{\lambda}_{jp-1}$ is a proxy for the capability parameter, λ , discussed below; and ξ_{kp} , ψ_{rp} and u_{jp} are an industry-period effect, a region-period effect and a mean-zero disturbance, respectively. For variables that vary by year, we average over years within a period. Note that (3.4) is effectively in first-differences, with time-invariant firm characteristics differenced out.

There are a number of reasons why the change in export share, Δe_{jp} , may be correlated with the error term, u_{jp} , and OLS estimation of (3.4) may plausibly lead to either positive or negative bias in the estimate of the coefficient of interest, θ . On one hand, positive productivity shocks at the plant level may lead both to greater exports and to higher wages, generating a positive bias in the OLS coefficient. On the other hand, a positive labor supply shock to a plant would be expected to lead to lower wages and greater exports, generating a negative bias in OLS. Other biases are also possible.¹⁹ To address the endogeneity of export changes, we instrument Δe_{jp} with the interaction between an indicator for the period immediately following the peso crisis ($p = 2$) and the value of the capability proxy in the previous period: $\widehat{\lambda}_{jp-1} * T_2$. This term can be interpreted as a proxy for the differential inducement to export created by the peso devaluation, which we take to be exogenous. Note that any stable relationship between the capability proxy and the change in wage outcomes will be captured by the uninteracted $\widehat{\lambda}_{jp-1}$ term. The IV estimate of θ will only reflect the effect of the change in exporting induced by the differential impact of the devaluation on more- vs. less-capable firms between period 1 and period 2.

A potential concern with our IV approach is that exporting and wages can be thought of as simultaneous outcomes of the same firm/plant optimization problem. In this view it is not clear it makes sense to estimate the effect of one choice on the other. This difficulty can be avoided by

¹⁹For instance, a positive demand shock specific to the domestic market would likely lead to greater investment, a higher capital-labor ratio, higher wages, and a lower export share (since domestic sales appear in the denominator of export share), generating a negative bias in OLS.

focusing on the reduced form corresponding to the IV model:

$$\Delta y_{jp} = \varphi(\widehat{\lambda}_{jp-1} * T_2) + \widetilde{\gamma} \widehat{\lambda}_{jp-1} + \widetilde{\xi}_{kp} + \widetilde{\psi}_{rp} + \widetilde{u}_{jp} \quad (3.5)$$

where $\widetilde{\xi}_{kp}$, $\widetilde{\psi}_{rp}$, and \widetilde{u}_{jp} are again an industry-period effect, a region-period effect, and a mean-zero error, respectively. Here the $\widehat{\lambda}_{jp-1} * T_2$ term, the proxy for the differential inducement to export, is clearly predetermined with respect to a plant's period-2 decisions, and the coefficient on that term can be interpreted as capturing the effect of the export inducement on the outcome variable.²⁰

A key step in the implementation is to choose the proxy for the capability parameter, $\widehat{\lambda}$. Following Verhoogen (2008), we use plant size as our primary proxy. Intuitively, since capability leads plants to become large, we can infer plants' underlying capability from their size. Plant size has the advantage that it is directly observed, is relatively well measured, and does not require the strong assumptions imposed by standard residual-based methods of estimating total factor productivity (TFP). Two natural measures of plant size are domestic sales and employment. In Melitz (2003)-type theoretical frameworks (such as the one used by Verhoogen (2008)), domestic sales have the attractive theoretical property that they bear a smooth, continuously differentiable relationship to the latent capability term; by contrast, employment (like total sales) jumps discontinuously at the extensive margin for exports. On the other hand, domestic sales have the empirical disadvantage that measurement error can generate spurious correlation with the export share, since domestic sales appear in the denominator of the latter. Measurement error in domestic sales in period $p - 1$ will generate a spurious negative relationship with the level of the export share in $p - 1$ and a spurious positive relationship with the change in export share from $p - 1$ to p .²¹ Because our primary

²⁰Note that the reduced form is equivalent to the approach of Verhoogen (2008). That paper estimated OLS models of the form:

$$\Delta y_{jp} = \mu + \widehat{\lambda}_{jp-1} \pi_p + \widetilde{\xi}_k + \widetilde{\psi}_r + v_j \quad (3.6)$$

separately by period, with $p = 1993, 1997, 2001$, and then compared π_{1997} (for the peso crisis period 1993-1997) to π_{2001} . Except for the fact that here we look at three-year periods 1992-1994, 1996-1998, and 2000-2002 in place of individual years 1993, 1997, 2001, the difference in coefficients $\pi_{2001} - \pi_{1997}$ in that paper is equivalent to the OLS estimate of φ in our reduced-form model, (3.5).

²¹A similar issue arises when log employment is used as the proxy and the capital-labor ratio or average hourly wage (from the EIA data), which are calculated by dividing by employment, are dependent variables.

specification is an IV specification with export share on the left-hand side of the first stage, we prefer (log) employment as our primary proxy for the capability parameter, $\hat{\lambda}$. But we also present results using (log) domestic sales in the reduced form (equation (3.5)) where export share is not on the left-hand side and spurious correlation due to measurement error is not a concern.²²

To check robustness, we also present results using two additional proxies: a measure of total factor productivity using the Levinsohn and Petrin (2003) method, and an index of export propensity, calculated as the predicted values from a tobit regression of export share on log sales, log employment and log capital-labor ratio. Since both rely in part on variation in sales, they are subject to the spurious correlation issue discussed above; for this reason, we only estimate the reduced-form model when we use them. Qualitatively, the results do not depend on the choice of proxy.

We conduct two further robustness checks. First, using only the IMSS data we include an earlier period, 1988-1990 (which we refer to as period 0), in addition to the three periods mentioned above, 1992-1994 (period 1), 1996-1998 (period 2), and 2000-2002 (period 3). This allows us to compare changes between two pre-shock periods (periods 0 and 1) to changes over periods that span the shock (periods 1 and 2). Because we do not observe export share when using only the IMSS data, we again use the reduced-form model for this check. Second, we estimate the effect of exporting on the wages of stayers, workers who are continuously employed in a given firm over two periods. These results do not use the AKM methodology and in particular do not require the conditional random mobility assumption, (3.2). They are analogous to estimating a model with job-spell fixed effects, as in several papers cited above. We will see that both of these robustness checks yield results qualitatively similar to our baseline estimates.

In both the IV and reduced-form approaches, our interpretation relies on the assumption that the devaluation affected wage outcomes differentially within industries only through its impact on the incentives of plants to export. Verhoogen (2008) considered a number of reasons why this as-

²²Verhoogen (2008) used log domestic sales as the preferred proxy. But that paper used a reduced-form framework (see footnote 20) where estimates did not rely on a first stage with export share as the dependent variable. (When export share was on the left-hand side, that paper instrumented log domestic sales with its lag.)

sumption might be violated, if for instance the devaluation affected larger and smaller firms within industries differently for reasons unrelated to exporting. Readers are referred to that paper for more extensive discussion, but two points are worth re-emphasizing. First, the assembly-for-export (*maquiladora*) sector in Mexico provides a sort of placebo test. The *maquiladora* sector exported essentially all of its output both before and after the devaluation, and hence we would not expect the devaluation to have generated a differential within-industry shock to exporting. At the same time, if the macro shock had a differential within-industry impact through a channel other than exporting, one would expect it to show up in the *maquiladora* sector as well. Consistent with our interpretation, there was no differential change in wages between larger and smaller *maquiladora* plants during the peso-crisis period. Second, it does not appear that differential access to credit markets can explain the empirical patterns. While there is evidence that exporting plants faced a lower cost of capital than non-exporters, likely due to greater access to foreign capital, they also had a greater share of dollar-denominated loans before the crisis and hence their balance sheets were more adversely affected by the devaluation; these effects appear to offset. There was no differential within-industry change in the cost of capital that would explain the differential wage changes.²³

3.3 Data and Background


The employer-employee data we use are from the administrative records of the *Instituto Mexicano del Seguro Social (IMSS)*, the Mexican social security agency.²⁴ In principle, all private Mexican employers are required to report wages for their employees to IMSS and to pay social-security

²³A further possible concern is that plants (within industries) differ in the extent to which they import inputs and that they were differentially affected by the devaluation for this reason. But the import share of inputs is positively correlated with plant size and export status, so this effect, if present, would suggest that larger plants were more negatively impacted by the devaluation relative to smaller plants, making it more difficult for us to find the positive (relative) wage effects that we find; in this sense, our estimates are conservative. Another possible concern is that we do not allow the return to ability to change within period and the peso devaluation might have affected the general-equilibrium return to skill. A first response to this concern is that we estimate the model (3.1) separately by periods before and after the devaluation, allowing the individual effect to differ across periods. A one-time effect of the devaluation on the return to ability will be captured by the period-specific individual effects. A second response is that a previous version of this paper (Frías et al. (2009) allowed for year-by-year changes in the return to ability (see footnote 12) and found broadly similar results.

²⁴Previous papers using these data include Castellanos et al. (2004) and Kaplan et al. (2004), Kaplan et al. (2010).

taxes on the basis of their reports. In practice, only about half of private-sector remunerated employees are reported to IMSS (and hence considered formal), and half are considered informal; see Appendix Table C.2.1.²⁵ While the size of the informal sector seems large by developed-country standards, it is not out of line for countries at Mexico's income level.²⁶ At the level of individuals, the IMSS data contain information on age, sex, daily wage (including benefits), and state and year of the individual's first registration with IMSS. Unfortunately, they do not contain information on individuals' education levels. At the establishment level, the data contain only industry and location. We have access to the IMSS data from 1985 to 2005.

Figure 3.1. Wage percentiles, top- and bottom-codes, IMSS data, 1985-2005



Notes: Wage percentiles calculated from raw IMSS data, after very basic cleaning (steps 1-3 of cleaning procedure described in Appendix C.1). There are three minimum wages in Mexico, corresponding to different geographic regions (zones A, B, C). Displayed is the minimum wage for Zone C, the lowest of the three. The top-code was 10 times the minimum wage in Mexico City (Zone A) from 1985-1993, 18 times in 1994, and 25 times from 1995-2005. Prior to 1991, establishments were allowed to report wages below the corresponding minimum wage to IMSS. Beginning in 1991, this practice was disallowed. Average 2002 exchange rate: 9.60 pesos/US\$1.

²⁵Most public-sector workers and employees of the state-run petroleum company are covered by separate systems.

²⁶See e.g. Schneider and Enste (2000). The employment figures for manufacturing in the IMSS data differ by less than 10% from independently reported figures in the 1993 Industrial Census, suggesting that underreporting of employment is not especially severe in the social security data Kaplan et al. (2004), Kaplan et al. (2010).

An important practical issue is that the bottom- and top-codes in the IMSS data have changed over time. Prior to 1991, IMSS allowed establishments to report wages below the corresponding regional minimum wage;²⁷ beginning in 1991, this practice was disallowed (even if actual wages presumably continued to be below the legal minimum in some cases). The changes in the top-code are illustrated in Figure 3.1, which also displays several wage quantiles from the IMSS data.²⁸ (Note also that average real wages dropped significantly following the peso devaluation in late 1994; we will see below that they simply dropped less in plants that increased exports during the crisis period.) To reduce biases due to changes in top- and bottom-codes, we “winsorize” the wage data, by replacing wages below the 10th percentile by the wage at the 10th percentile, and wages above the 90th percentile by the wage at the 90th percentile. This process also reduces the influence of outliers due to misreporting and other forms of measurement error.²⁹ Another potentially important measurement issue is that employers’ payroll tax burdens depend on the wages they report to IMSS and hence they have an incentive to under-report wages. Kumler et al. (2015) document such under-reporting. But they also show that the under-reporting appears to be minimal for the set of larger manufacturing plants that can be linked to the EIA plant panel; see Figure 10 of that paper and the corresponding discussion in the main text. Moreover, as long as any under-reporting is constant over time within plant, it will be differenced out in our estimation procedure.

Our cleaning procedure for the IMSS data is described in detail in Appendix C.1. Briefly, we focus on workers classified by IMSS as permanent, who earn a positive wage, in an establishment for which we know municipality and industry, and who are age 14-64. If wages for more than one job are reported for a particular worker, we select only the highest-wage job. Although in principle

²⁷There are three minimum-wage regions in Mexico, with the minimum wage in Mexico City and other urban areas generally 10-20% higher than in poorer rural areas.

²⁸Prior to 1993, the top-code was 10 times the minimum wage in Mexico City; in 1994, it was 18 times; and since 1995 it has been 25 times the minimum wage in Mexico City.

²⁹Our procedure for dealing with the top- and bottom-coding differs from that of CHK, who impute wages above the top-codes in their data using a series of Tobit models. Given the large changes in the top-codes over time, we worry that such an imputation would run the risk of introducing significant errors in our setting. Given also that our aim is not to characterize firms’ contributions to overall wage inequality, the advantages of such imputation seem small relative to the risk of errors they might introduce.

we observe a wage for every individual for every day of every year, we select the wage on a single day of each year, Sept. 30. After cleaning, we observe between 5.26 and 10.15 million individuals per year over the period. Summary statistics for the cleaned IMSS data are reported in Appendix Table C.2.2.

The plant-level data are from the *Encuesta Industrial Anual (EIA)* [Annual Industrial Survey], conducted by the *Instituto Nacional de Estadísticas y Geografía (INEGI)*, the Mexican statistical agency.³⁰ The variables are standard for plant surveys: employment, total wage bill, total hours worked, investment, capital stock, domestic and export sales, among others. The sampling design is less standard. INEGI has periodically drawn a deterministic sample of plants from a subset of manufacturing industries and followed those plants over time, with no refreshing of the sample. Plants with more than 100 workers are included with certainty in the initial sample. INEGI has created separate panels for 1984-1994, 1993-2003, 2003-2009, and 2009-present.³¹ Although it is possible to link some plants across panels, the number that can be linked both across panels and to the IMSS data is too small for our purposes. We therefore focus on the 1993-2003 panel. During this period, the EIA did not include information on *maquiladora* plants, assembly-for-export plants located mainly along the U.S. border, which were covered by a different dataset. In our cleaning procedure, we do imputations for missing information following the procedure described in Appendix II of Verhoogen (2008). We then require plants to have complete information at the plant level on employment (hours and number employed), hourly wage, total sales, export share, and capital-labor ratio. After cleaning (also following Appendix II of Verhoogen (2008)), there are 3,529 plants in a balanced panel with complete EIA information in every year over the 1993-2003 period. We refer to this balanced panel as the “EIA panel.” Summary statistics by export status for the EIA panel for 1993 are in Appendix Table C.2.3. The differences between exporters and non-exporters are similar to those documented for the U.S. by Bernard and Jensen (1999), and sub-

³⁰The name of INEGI’s main plant panel has changed over time; for recent years, it is referred to as the *Encuesta Anual de la Industria Manufacturera (EAIM)* [Annual Survey of Manufacturing Industry].

³¹INEGI has also conducted a monthly version of the survey, formerly called the *Encuesta Industrial Mensual (EIM)* [Monthly Industrial Survey] and now called the *Encuesta Mensual de la Industria Manufacturera (EMIM)* [Monthly Survey of Manufacturing Industry], with slightly different periodization. Wage bill and hours for two occupational categories, white-collar (*empleados*) and blue-collar (*obreros*), are available from this companion survey.

sequently for many other countries: exporters are larger, more capital-intensive, and higher-wage than non-exporters, and they make up a minority of plants in each industry.

The EIA data have been linked to the IMSS employer-employee data using establishment name, location (municipality and state) and street address. Although in principle all plants appearing in the EIA should also appear in the IMSS data, it has only been possible to link approximately 2,800 of the 3,529 plants from the EIA panel to the IMSS data in each year.

As noted above, we estimate wage premia using the largest connected set in each period. In our setting, the largest connected set does not include as large a share of establishments as has typically been the case in developed countries. Table 3.1 reports summary statistics from the IMSS data for the largest connected sets in each period — period 1 (1992-1994), period 2 (1996-1998), and period 3 (2000-2002), as well as period 0 (1988-1990) which will be used in a robustness check below. Comparing to Table C.2.2, we see that the largest connected sets typically include roughly half the number of establishments in the full IMSS dataset in each year. This is in part because of the high rate of informality in Mexico and in part because we focus on three-year periods, in contrast to AKM (and Abowd et al. (2002)), who focus on a single 12-year period, and CHK, who focus on 7-year periods; both factors tend to reduce the number of observable connections and hence the size of the largest connected sets. In addition, the establishment size distribution in Mexico is skewed to the left relative to, for instance, the U.S.;³² this in itself tends to reduce the share of establishments in the largest connected set. But it does not appear that the limited coverage of the largest connected sets is a severe problem for our purposes, for two reasons. First, although they cover a minority of establishments, the largest connected sets nonetheless cover a large majority of workers — approximately 90%. Second, and perhaps more importantly, the largest connected sets include almost all of the larger plants in manufacturing that appear in the EIA panel, for which we observe exporting behavior. Appendix Table C.2.4 reports the number of EIA panel plants that can be linked to the IMSS data, by connected set status, for each of our three main periods. Of the EIA plants that can be linked to the IMSS data, 98% or more are in the largest connected set in each

³²See e.g. Figure XIV of Hsieh and Klenow (2014), which plots the establishment size distribution for Mexico, the U.S., and India.

year.

Once the EIA panel has been linked to the IMSS data, we impose the requirement that plants be in the largest connected set in periods 1-3. The resulting panel contains 2,625 plants. We refer to this balanced panel as the “EIA-IMSS panel.” Table 3.2 reports summary statistics for the EIA-IMSS panel for 1993. Comparing to Appendix Table C.2.3, this panel is on average quite similar to the EIA panel. As mentioned above, once the plant and person components have been calculated in the IMSS data, we collapse the EIA-IMSS panel to the plant-period level, averaging across years within plant-period.³³

It is important to recognize that our estimates of the effect of exporting on wage premia are valid conditional on firms being in the EIA-IMSS panel, in which large firms are over-represented. Moreover, the wage premia estimates themselves are identified by workers who switch between formal plants (i.e. plants that are registered and report regularly to the social security agency); these are likely to be workers with more stable attachment to the formal sector. At the same time, larger manufacturing plants and workers with strong attachment to the formal sector are those for whom exporting is a realistic possibility. In this sense, the EIA-IMSS panel is a relevant sample for investigating the effect of exporting on wages. We believe that the results are likely to generalize to larger manufacturing plants in other countries, especially those approximately at Mexico’s level of industrial development.

As discussed in Verhoogen (2008), the devaluation of the peso in December 1994 represented an enormous shock to the Mexican economy. The Mexican peso lost approximately 50% of its nominal value in a few days. Figure 3.2 plots the real exchange rate over the 1989-2004 period. GDP fell by 6.7% from 1994 to 1995. Exports rose sharply, with approximately 85% destined for the U.S. market. Using the EIA panel of 3,529 plants, Figure 3.3 illustrates the shift toward the export market: the export share for the panel as a whole jumped sharply, and the number of plants with positive exports rose from approximately 30% to 45% of the sample.³⁴ Nominal wages,

³³For the plant and person components, we average over 1992-1994, 1996-1998, and 2000-2002; for the EIA variables, we average over 1993-1994, 1996-1998, and 2000-2002, since the EIA data we use only begin in 1993.

³⁴The peso crisis was a much larger shock than the North American Free Trade Agreement (NAFTA), which took effect in January 1994. Mexico’s main trade liberalization came with its entrance into the General Agreement on

Table 3.1. Summary statistics, IMSS individual-level data, largest connected sets

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

Notes: Sample is from IMSS employer-employee records after cleaning steps 1-7 in Appendix C.1. (Sample includes individuals who earn a positive wage in an establishment for which municipality and industry are observed, are age 14-64, and are employed in an establishment in the largest connected graph of establishments in their respective period (among 1988-1990, 1992-1994, 1996-1998 and 2000-2002). Winsorization is at 10th and 90th percentiles. Wages are reported both in “raw” (i.e. pre-winsorized) and winsorized form. See Section 3.3 and Appendix C.1 for further details. Average 2002 exchange rate: 9.60 pesos/US\$1.

perhaps surprisingly, remained nearly constant on average through the crisis; the sudden increase in the price of imports, and the inflation that it generated, meant that real wages fell dramatically. As noted above, the wage gains among exporting plants relative to other plants in the same industry should be understood in this context of large real wage declines overall; real wages in exporting plants fell by less than in non-exporting plants.

Table 3.2. Summary statistics, EIA-IMSS plant panel, 1993


	non-exporters (1)	exporters (2)	all plants (3)
Total revenue	153.10 (9.55)	417.40 (58.12)	232.02 (18.74)
Employment	188.32 (5.55)	333.43 (13.40)	231.66 (5.73)
K/L	140.14 (5.35)	177.64 (9.19)	151.34 (4.66)
Export share of sales		0.14 (0.01)	0.04 (0.00)
Avg. hourly wage (EIA)	42.75 (0.55)	58.14 (1.15)	47.35 (0.53)
N	1841	784	2625

Notes: Table reports statistics using 1993 data for the EIA-IMSS (balanced) panel of plants. Standard errors of means in parentheses. Exporter defined as export sales > 0. Export share is fraction of total sales derived from exports. Sales are measured in millions of 2002 Mexican pesos, capital-labor ratio in thousands of 2002 pesos, and average daily wage in 2002 pesos. Average 2002 exchange rate: 9.60 pesos/US\$1. For further details, refer to Section 3.3 and Appendix C.2.

3.4 Estimating Wage Premia

As described in Section 3.2 above, the first step of our approach is to estimate the AKM-type model (3.1) in the IMSS individual-level data, separately for periods 1992-1994, 1996-1998, and Tariffs and Trade in the mid-1980s, and by 1994 the vast majority of Mexican imports were covered by tariffs of 20% or less. Average U.S. tariffs on goods from Mexico were on the order of 3-5%. In the majority of cases, NAFTA phased out existing tariffs slowly over time. Relative to the exchange-rate devaluation, the year-by-year tariff changes were quite small.


Figure 3.2. Real exchange rate



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

2000-2002. Before presenting the estimates, we provide evidence for the validity of the conditional random mobility assumption, (3.2). Following CHK, we show simple event-study plots of average wage changes between firms in different quartiles of the distribution of wages paid to co-workers. For these plots, to ensure that we can observe two years before and two years after a job transition (again following CHK), we focus on four-year periods, 1992-1995 and 2000-2003. (The requirement that we observe two years before and two years after means that we must focus on transitions between the middle two years of each period, i.e. 1993-1994 or 2001-2002.) Figures 3.4 and 3.5 show the mean real wages of movers for these periods; to reduce visual clutter, we focus on workers leaving firms in quartiles 1 and 4. As in CHK, two features stand out. First, the wage trends before and after job switches are parallel across the different types of transitions; there are no Ashenfelter (1978)-type dips or rises before transitions. If the individual time-varying productivity shocks, ε_{it} , were determining job transitions, we would expect positive wage changes prior to movement to a higher-quartile plant, and negative changes prior to movement to a lower-quartile plant. Second,

Figure 3.3. Shift to export market, EIA panel



Notes: Data from 1993-2003 EIA (balanced) panel of 3,529 plants. Export percentage of sales calculated as (total exports for all plants)/(total sales for all plants). Plants with exports greater than zero classified as exporters. See data appendix for details.


the wage change for workers moving from one quartile to another has approximately the same magnitude and opposite sign of the change for workers moving in the opposite direction.³⁵ If the ε_{it} contained an employer-employee match-specific effect that also affected workers' mobility, we would expect the gain for workers moving to a higher quartile to be larger than the loss for those making the opposite move.³⁶ Both features suggest that the data are well described by a model with a fixed individual component, a fixed plant component, time-varying observables, and a random shock uncorrelated with mobility.

We now turn to estimation of the AKM-type model, (3.1). We estimate plant and person components for all individuals in the largest connected sets described by Table 3.1, separately by period, for periods 1-3. Table 3.3 reports key statistics. There are several points to notice. First,

³⁵Appendix Figures C.2.1 and C.2.2 make the symmetry clear visually. Appendix Figures C.2.3 and C.2.4 are similar to figures 3.4 and 3.5, but use nominal wages instead of real wages. The patterns are similar.

³⁶As Card et al. (2018) point out, models with a worker-firm match component tend to predict that workers who move will see wage increases, irrespective of whether they move to higher or lower quartiles.

Figure 3.4. Movers' mean real wages, IMSS data, 1992-1995




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

despite being parsimonious, the model has a high in-sample predictive power, with adjusted R^2 above 90%, as in other studies in the literature. Second, we see already that the standard deviation of the individuals effects remains relatively constant but the standard deviation of the plant component increases over time, foreshadowing our results below that the plant-level response to the export shock is explained primarily by changes in wage premia. Third, there is little evidence of negative correlation between estimated person and plant effects.

One might be concerned that the use of three-year periods and the high rate of informality in our context might limit the number of observed plant-to-plant switches and exacerbate what Abowd et al. (2004) call "limited mobility bias." Small numbers of switchers can generate a negative bias in the correlation between individual and plant effects (Andrews et al. (2008)) which can be substantial (Maré and Hyslop (2006), Andrews et al. (2008)). Intuitively, if plant effects are calculated first, and individual effects are calculated by subtracting the plant effects and contribu-

Figure 3.5. Movers' mean real wages, IMSS data, 2000-2003



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

tions of observables from log wages (as is common), then over-estimation of plant effects will tend to lead to under-estimation of individual effects (and vice-versa).³⁷ The smaller is the number of switchers, the less precisely the plant effects are estimated and the more severe is the bias. Andrews et al. (2008) suggest a correction, but (as CHK and Card et al. (2018) note) the correction requires strong assumptions on the covariance structure of the time-varying error terms. Following CHK, we instead rely on the assumption that the limited mobility bias is constant across periods and will be differenced out when we look at within-firm changes in the plant and average person components.

³⁷In our case, the Matlab PCG algorithm solves simultaneously for the plant effects and the individual effects for switchers, and the individual effects for non-switchers are then calculated using the estimated plant effects.

Table 3.3. Estimation results for AKM-type model, per period

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

Notes: The table shows statistics from the estimation of model (3.1) for periods 1992-1994, 1996-1998 and 2000-2002 separately. Sample is described in notes to Table 3.1. Individual effects are estimates of α_i , plant effects are estimates of $\psi_{J(i,t)}$, and Xb are estimates of $X'_{it}\beta$ in (3.1). Covariates included in X_{it} are age squared, age cubed (both recentered at 40), tenure, tenure squared, and year effects.

3.5 Estimating the Effect of Exporting on Wage Premia

This section examines the relationship between exporting and the estimated plant and average person components. Before turning to regression tables, we provide a visual illustration of the key reduced-form relationships. Using the EIA-IMSS panel, Figures 3.6a-c present cross-sectional non-parametric regressions for outcome variables from the EIA — export share, log capital-labor ratio, and log plant-average hourly wage — against log plant size for period 1 (1992-1994). Because of the potential spurious correlation due to measurement error discussed in Section 3.2 above, we use log employment (hours) as the capability proxy when export share is the y-axis variable (Figure 3.6a) and log domestic sales as the proxy when log capital-labor ratio and log hourly wage

are the y-axis variables (Figures 3.6b-c). All variables have been deviated from industry means. All three graphs display a clear positive relationship: within industries, larger plants are more likely to export, are more capital intensive, and pay higher wages. The graphs in the second row, Figures 3.6d-f, plot non-parametric regressions of *changes* in the same variables between two consecutive periods against log employment in the first of those periods (i.e. the x-axis variable is log plant size in period 1 (1992-1994) for the period 2 vs. period 1 curve, and log plant size in period 2 (1996-1998) for the period 3 vs. period 2 curve). Given that all variables have been deviated from industry means, the key features are the relative slopes of the curves. We see that larger plants saw a greater increase in export share, capital intensity, and average hourly wages than smaller plants from period 1 (1992-1994) to period 2 (1996-1998), and these differential changes were greater than the corresponding changes from period 2 (1996-1998) to period 3 (2000-2002). The graphs show that the peso crisis affected plants differently within industries, with more positive (or less negative) effects on larger plants. These figures are similar to Figures A4-A5 of Verhoogen (2008), which used a broader EIA panel without linking to the IMSS data.

Turning to outcomes from the IMSS employer-employee data, Figures 3.7a-c present non-parametric regressions similar to those in Figures 3.6a-f, with average log daily wage from IMSS, the plant component, and the average person component as y-axis variables and log employment (hours) from the EIA as the x-axis variable.³⁸ In the first row (Figures 3.7a-c), we see that all three variables bear a positive relationship to log domestic sales. Recall that in expectation the plant component and average person component sum to the plant-average log wage (refer to equation (3.3)); the equality very nearly holds true for the sample analogues. The y-axes have the same scale in all graphs and the slope for the plant component is approximately twice the slope for the average person component. Intuitively, this suggests that approximately two thirds of the within-plant *cross-sectional* correlation between plant-average wages and plant size can be attributed to wage premia. Figures 3.7d-f plot regressions of changes against log hours, similarly to Figures 3.6d-f. The similarity of Figure 3.7d, which uses average log daily wage from the IMSS data,

³⁸The spurious correlation issue is less of a concern here, because the IMSS-derived wage variables are not calculated by dividing by the EIA employment variable.

to Figure 3.6f, which uses log average hourly wage from the EIA data, suggests that there was not much differential change in hours worked in response to the devaluation. Comparing Figures 3.7e and 3.7f to Figure 3.7d illustrates the main point of the paper: the difference in slopes for plant-average wages of the period 1-2 and 2-3 curves in Figure 3.7d is largely explained by the difference in slopes for wage premia (Figure 3.7e), not the difference in slopes for the average person component (Figure 3.7f). This suggests that the differential effect of the devaluation on wages can be largely attributed to the effect of the devaluation on wage premia, rather than on workforce composition.

We now turn to the estimation of our parametric IV model, equation (3.4). Table 3.4 presents our baseline results, with the first stage in Panel A Column 1, the reduced-form results in Panel A Columns 2-6, and the IV results in Panel B. As discussed in Section 3.2 above, we use log employment (hours) as the capability proxy in our baseline IV estimates.³⁹ In the first stage (Panel A, Column 1), the coefficient on the excluded instrument, initial log employment \times devaluation dummy (written as $\hat{\lambda}_{p-1} * T_2$ in equation (3.5) above), is significant at the 1% level. The coefficient indicates that comparing two plants that differ in initial employment by a factor of 2.7 (100 log points), the larger plant experienced approximately a 1 percentage point greater increase in export share due to the devaluation, above and beyond the trend of .4 percentage point greater increase between periods. This is a modest but highly statistically significant difference. Consistent with Figures 3.6-3.7, the reduced form results show that the relationships between the changes in the capital-labor ratio, log average hourly wage, average log daily wage, and the plant component are steeper for periods 1-2 than for periods 2-3, and are statistically highly significant. There is no significant difference in these slopes for the person component. The IV estimates in Panel B indicate a positive effect of exporting on the capital-labor ratio, log average hourly wage, average log daily wage, and the plant component, but not on the person component. The coefficient in Column 4, for the change in average log daily wage, indicates that a .01 (1 percentage point) increase in export share (on a mean value of 4 percentage points; see Table 3.2) generates a 3.172

³⁹The spurious-correlation concern discussed in Section 3.2 above still applies for the capital-labor ratio and log average hourly wage from the EIA. We present similar results using alternative proxies below.

log point (3.22%) increase in the average log daily wage. This response is large but not implausible.

Comparing the estimates in Columns 4 and 5, we see that the effect on the plant component is not statistically different from that of the effect on plant-level average wages. The point estimate for the plant component (Column 5) is larger than the coefficient for the plant-level average wage (Column 4) in this specification, but this difference will not persist across specifications below. The robust fact is that the plant component effect is approximately as large as the effect on plant-level average wages. Since the plant and person components sum to the plant-level average in expectation, and very nearly in our sample, the coefficients in Columns 5 and 6 very nearly sum to the Column 4 coefficient; the point estimate for the person component (Column 6) is negative in this specification, and will be robustly close to zero across specifications. Our interpretation of the IV results is that essentially all of the within-plant effect of exporting on plant-level average wages driven by the peso devaluation can be explained by the effect on wage premia, at least over the time frame we are able to study.

Weakness of our instrument is a potential concern. The robust first-stage F-statistic, reported in Column 1 of Panel A, is 7.320, which is below the Staiger and Stock (1997) rule-of-thumb level of 10. It is important to acknowledge that weak instruments can magnify the bias resulting from any violations of the exclusion restriction (in our case that the devaluation affected plants differentially within industry only through its differential effect on export shares) (Bound et al. (1995)). But two arguments suggest that the weak-instrument concern is not prohibitively severe in this context. First, in the presence of weak instruments, assuming the exclusion restriction is valid, we would expect the IV estimates to be biased toward the OLS estimates. (See e.g. section 4.6.4 of Angrist and Pischke (2008).) Appendix Table C.2.5 reports the OLS estimates, and reveals that the OLS estimates are smaller than the IV estimates, suggesting that weakness of the instrument would bias the IV estimates down in our case, and make it more difficult to reject the null of no effect of exporting on our outcomes.⁴⁰ Second, in this just-identified context with a single endogenous

⁴⁰Panel A of Appendix Table C.2.5 reports OLS estimates omitting the initial level of the capability proxy (log employment). The coefficients for the changes in capital-labor ratio, plant-average wages, and the plant component are positive but not significant. Once the initial level of the capability proxy is included, in Panel B, the estimates are even smaller. Comparing the Panel B results to the IV results in Table 3.4, we see that the OLS coefficients are

regressor and a single instrument, the Anderson et al. (1949) Wald test, which is robust to weak instruments, is equivalent to a test that the coefficient on the instrument is zero when we include it in place of the endogenous covariate in the second stage. We effectively do that in the reduced-form specifications of Columns 2-6, and t-tests comfortably reject the hypothesis that the coefficients on the instrument are zero in Columns 2-5.

To examine robustness of these results, we present estimates using three alternative proxies: log domestic sales, log TFP, and predicted export share. To estimate log TFP, we use a version of the Levinsohn and Petrin (2003) methodology, with value-added as the outcome, employment (white collar and blue collar separately) and log capital as covariates, and log materials and log electricity as proxies, separately by 2-digit sector. We predict export share using a tobit regression of export share on log employment (hours), log total sales, and log capital-labor ratio. Coefficient estimates from the TFP and predicted export share estimation are reported in Appendix Table C.2.6. In Melitz (2003)-type theoretical frameworks, all three proxies are expected to be positively correlated to underlying plant capability. Unsurprisingly, the proxies are highly (but not perfectly) correlated with one another (Appendix Table C.2.7). Table 3.5 reports cross-sectional regressions of the outcome variables from Figures 3.6-3.7 and Table 3.4 against these three proxies, as well as log employment (hours) for comparison. The results are almost entirely consistent with the positive relationships we observed in Figures 3.6a-c and 3.7a-c; the one exception is the relationship between export share and domestic sales, which is subject to the spurious negative bias (in levels) due to measurement error discussed in Section 3.2 above.

Table 3.6 presents first-stage, reduced-form, and IV results using the three alternative proxies, similar to Columns 1, 4, 5, and 6 of Table 3.4. The results are very consistent with the baseline results. The first-stage coefficient remains highly significant and the IV estimates are statistically indistinguishable from those of Table 3.4. The reduced-form estimates naturally are of different magnitudes, because of the different scaling of the proxies, but the qualitative patterns are similar.

significantly smaller than the IV estimates. This is consistent with the hypothesis that plants are subject to idiosyncratic factor-supply shocks that move factor prices and generate a negative correlation between exporting and factor prices (e.g. a positive labor-supply shock reduces wages and increases exports), generating a negative bias in the OLS coefficients.

These results reinforce the finding that essentially all of the differential increase in wages due to the peso devaluation can be attributed to changes in the plant component (i.e. wage premia), rather than sorting on skill. It is reassuring that the key results do not depend on the choice of capability proxy.

The share of the plant-average wage effect that can be explained by the effect on wage premia is large, indeed larger than we originally expected, and merits further discussion. Before the devaluation, we observe a positive, significant relationship between plant size (a strong predictor of export status) and the average person component in cross-section, which presumably reflects a long-run equilibrium. The contrast with the zero differential response of the average person component in response to the devaluation suggests that there may be a difference between the short-term and long-term responses to an exogenous increase in exporting. As discussed in the introduction, in models with search frictions in the labor market, an expansion of exporting may raise the value to firms of avoiding vacancies, and may lead firms to raise wage premia in the short term to retain existing workers and/or to attract new ones. These wage premia are then expected to fall over time, as firms replace existing workers with higher-skill workers from outside the firm or simply as the value of avoiding vacancies falls (Coşar et al. (2016), Fajgelbaum (2016), Bellon (2017)). There is suggestive evidence of this dynamic in Table 3.4: the coefficients on (uninteracted) initial log employment for the average person component in Column 6, Panels A and B (which are identified by differential trends between smaller and larger plants in the post-devaluation periods (between period 2 (1996-1998) and period 3 (2000-2002))) are positive and significant, while the corresponding coefficients for the plant component are close to zero, suggesting that the person component catches up in the post-devaluation period. However, this finding is not robust to using log domestic sales or predicted export share as alternative capability proxies in Table 3.6. It seems difficult to draw strong conclusions about the overshooting phenomenon based on this evidence.

Another possible hypothesis about the dynamic response to the devaluation is simply that the response unfolded slowly over time and was not complete by the end of our period 2 (1998). Figure 3.3 illustrated that entry into exporting continued to increase through 1997; it is possible that wage

effects were more protracted. In this case, the differential change in wages in wage premia between period 2 (1996-1998) and period 3 (2000-2002) would capture some of the delayed effect of the devaluation. We note that in this case our estimates comparing the period 1-2 change and the period 2-3 change would understate the true effect of the shock.

To shed further light on the dynamics, we compare the period 1-2 change to a change over similarly spaced periods before the devaluation. Due to the data constraints discussed in Section 3.3, we cannot use the EIA plant-level data for this exercise, and hence we cannot estimate the IV model with export share as an endogenous covariate. However, using only the IMSS employer-employee data, we are able to estimate the corresponding reduced form over a longer period. We define period 0 to be 1988-1990 and estimate the AKM-type model (3.1) for the largest connected set for this period. We then select establishments in the EIA-IMSS panel which which AKM estimates are also available in period 0. There are 2,316 plants in this balanced panel, observed over four periods. We use log employment (headcount) as the capability proxy, since we do not observe hours in the IMSS data. Panel A of Table 3.7 reports cross-sectional correlations similar to Table 3.5 for this modified sample and capability proxy; the results are similar to those in Table 3.5. To show the time pattern across periods, Panel B of Table 3.7 reports regressions period-by-period, of changes in the dependent variables on the level of the capability proxy separately by period. As discussed in footnote 20 above, the difference in coefficients between Panels B.2 and B.3 is precisely the estimate we would obtain of the coefficient φ in the reduced-form model, (3.5).⁴¹ The advantage of viewing the results period-by-period is that we can also compare the coefficient in Panel B.2 (for the change from period 1 (1992-1994) to period 2 (1996-1998)) to the coefficient in Panel B.1 (for the change from period 0 (1988-1990) to period 1 (1992-1994)). Panel C reports the differences in coefficients and their standard errors. Overall, the table is consistent with our previous findings. The coefficients for (changes in) both plant-average log wages and the plant component are significantly larger for the periods spanning the peso crisis period (periods 1-2, π_2) than for the two earlier periods (periods 0-1, coefficient π_1) or the two later periods (periods

⁴¹The coefficients in Panels B.1-B.3 correspond to estimates of π_p in equation (3.6) in footnote 20.

2-3, π_3). It is important to note that the coefficients for (changes in) plant-average log wages, the plant component, and the average person component are significantly positive (marginally so for the average person component) even in the earlier periods (periods 0-1, π_1). This suggests that there are differential background trends between larger and smaller plants, independent of the devaluation. The fact that the coefficient for the pre-devaluation periods (π_1) is larger than the coefficient for the post-devaluation periods (π_3) is consistent with the hypothesis that there was overshooting in wages and wage premia and that wages and the plant components declined over periods 2-3 in response, as discussed above. But the key points for our interpretation is that the differential response of both wages and wage premia was clearly larger in the periods spanning the pesos crisis (periods 1-2), at 99% confidence, and that effectively all of the effect in plant-level wages can be explained by the effect for wage premia.

As a final robustness check, we consider the effect of the export shock on the wages for stayers, employees continuously employed in a given plant for consecutive periods. Table 3.8 reports first-stage, reduced-form, and IV results similar to Columns 1 and 4 of Table 3.4, but where the change in the average log daily wage of stayers is the outcome of interest. As mentioned in Section 3.2 above, this is analogous to estimating a model with job-spell fixed effects, as in several existing papers.⁴² The estimates are identified by variation in wages within firm-worker matches in response to changes in export share at the plant level. Despite the very different approach, the reduced-form and IV coefficients in Column 2 are very similar to the baseline estimates for changes in wage premia in Table 3.4, Column 5. The coefficient from this table falls within the range of the point estimates for the four capability proxies from Table 3.4, Column 5 and Table 3.6 Column 3. It is consistent with our expectations, and reassuring, that the estimates for stayers are similar to the estimates based on switchers between plants.

⁴²This specification is not exactly equivalent to an individual-year-level specification with job-spell effects, because we have collapsed the data at the period level, and define as a stayer any worker who appears in a given plant in at least one year of two consecutive periods.

3.6 Conclusion

This paper has investigated the role of exporting in shaping Mexican plants' wage policies – in particular the payment of wage premia, defined as wages above what individuals would earn elsewhere on the labor market. Following CHK, we have estimated simple AKM-type models separately by three-year periods, to decompose plant-average wages into a plant component (i.e. wage premium) and an average person component. Following Verhoogen (2008), we have argued that the peso devaluation of late 1994 generated a greater effective export inducement to larger than to smaller plants within each industry. Using this differential shock as a source of exogenous variation in exports, we find that exporting led plants to increase wage premia.⁴³ The increase in the plant component can explain essentially all of the differential increase in plant-average wages arising from the export shock over the period we study. The findings highlight the causal importance of trade shocks — and product-market shocks more generally — in determining wage premia at the firm level.

Three questions seem particularly worthy of further research. First, how persistent are the effects of exporting on wage premia? We have found that the effects persist over 6-8 years, but it remains an open question whether firms will gradually reduce premia over a longer period, as several recent models with search frictions suggest. Answering this question would require a longer data series following the trade shock.

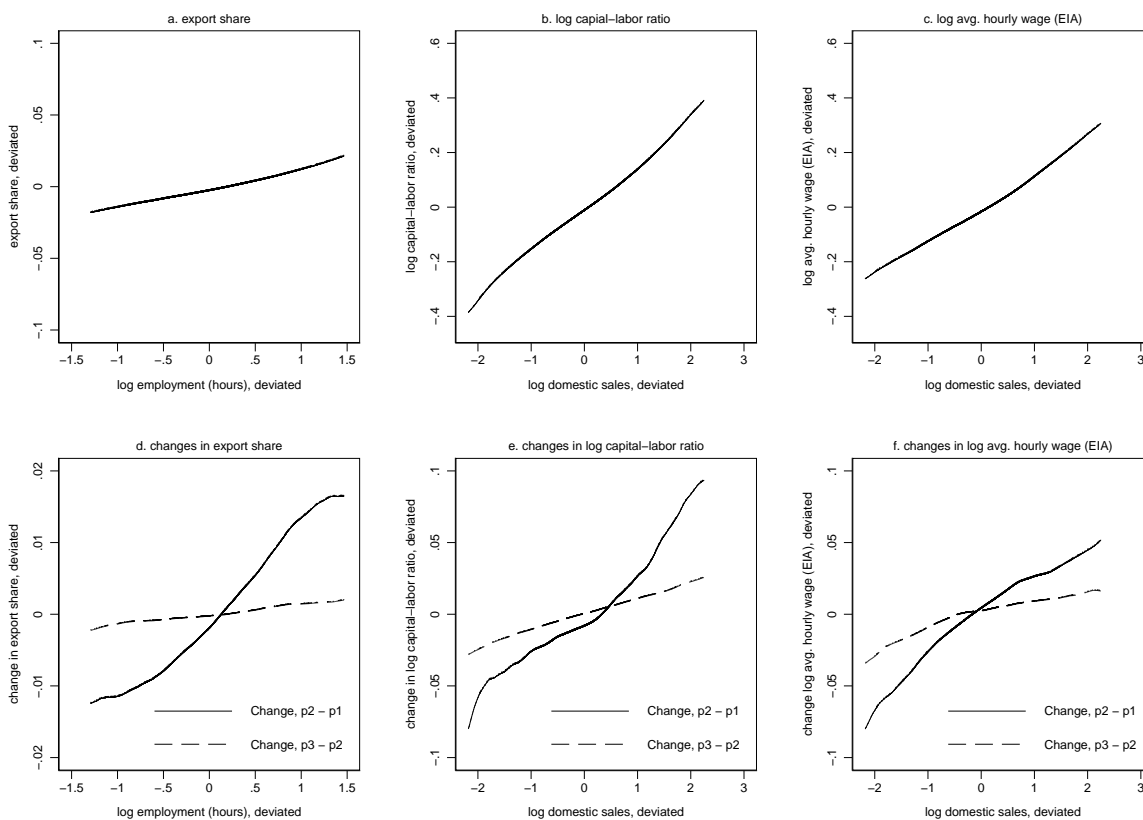
Second, to what extent did the differential increase in wage premia in response to the peso devaluation contribute to aggregate wage inequality in Mexico? Our paper suggests a mixed response to this question. On one hand, the export shock led initially larger plants within each industry, which already paid higher wages and wage premia, to increase wages and wage premia further. In this sense, the effect of exports on wage premia tended to increase inequality. On the other hand, the fact that we find limited effects on skill composition suggests that the effect of the export shock on the general-equilibrium return to skill is likely to be limited. The within-industry disequalizing

⁴³For readers who prefer the reduced-form specification over the IV, for the reasons discussed in Section 3.2 above, the results can be interpreted as showing that the differential inducement to export created by the devaluation increased wage premia.

effect of the trade shock may be confined to manufacturing (and possibly other tradable) sectors. A fuller accounting of the role of firms' wage policies in overall inequality in Mexico is a topic for future work.

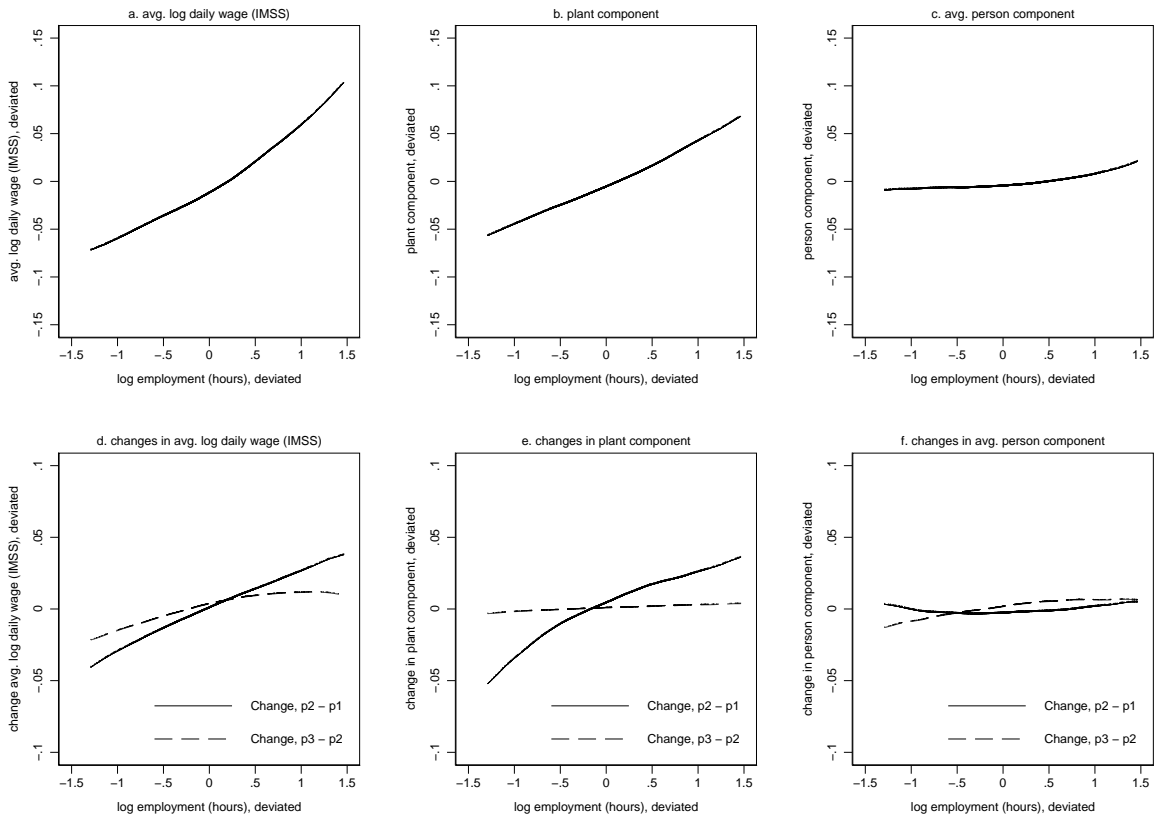
Third, to what extent do the income levels (or other characteristics) of destination markets matter for how export demand shocks get transmitted into wage premia? There is increasing evidence that destination income, as opposed to the volume of exports per se, matters for plant-average wages (Brambilla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

Figure 3.6. Non-parametric regressions, variables from EIA plant data



Notes: Non-parametric regressions are estimated using Stata `lpoly` command, Epanechnikov kernel, and the default rule-of-thumb bandwidth. Figures 3.6a-c plot non-parametric regressions in levels for period 1 (1992-1994). Figures 3.6d-f plot non-parametric regressions of changes between two consecutive periods against log employment in the earlier period (i.e. the x-axis variable is log plant size in period 1 (1992-1994) for the “Change, p2 - p1” curve, and log plant size in 1996-1998 for the “Change, p3 - p2” curve. Sample is EIA-IMSS panel (with variables averaged by plant-period). Variables are from EIA, and are deviated from their industry-period means. Log avg. hourly wage calculated as (total wage bill/total hours worked). Further data details are in Section 3.3 and Appendix ??.

Figure 3.7. Non-parametric regressions, wage components from IMSS employer-employee data



Notes: Non-parametric regressions are estimated using Stata `lpoly` command, Epanechnikov kernel, and the default rule-of-thumb bandwidth. Figures 3.7a-c plot non-parametric regressions in levels for period 1 (1992-1994). Figures 3.7d-f plot non-parametric regressions of changes between two consecutive periods against log employment in the earlier period (i.e. the x-axis variable is log plant size in period 1 (1992-1994) for the “Change, p2 - p1” curve, and log plant size in 1996-1998 for the “Change, p3 - p2” curve). Sample is EIA-IMSS panel (with variables averaged by plant-period). Plant component is $\hat{\psi}_{J(i,t)}$ and average person component is the plant-level average of $\hat{s}_{it} = \hat{\alpha}_i + X'_{it}\hat{\beta}$ from estimation of equation (3.1) in IMSS data for corresponding period. Further data details are in Section 3.3 and Appendix ??.

Table 3.4. Baseline estimates, exports and plant-level outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
	Δ export share	Δ log K/L	Δ log avg. hourly wage (EIA)	Δ avg. log daily wage (IMSS)	Δ plant component	Δ avg. person component
A. First stage and reduced form						
init. log employ. \times devaluation	0.009*** (0.003)	0.040** (0.018)	0.045*** (0.009)	0.029*** (0.006)	0.040*** (0.008)	-0.011 (0.007)
init. log employ.	0.004* (0.004)	0.053*** (0.053)	0.024*** (0.024)	0.012*** (0.012)	0.002 (0.002)	0.010** (0.010)
F-stat	7.320					
B. IV						
Δ export share		4.449* (2.495)	5.041** (2.008)	3.172** (1.280)	4.425** (1.755)	-1.237 (0.852)
init. log employ.		0.037* (0.021)	0.006 (0.016)	0.001 (0.010)	-0.014 (0.014)	0.015** (0.007)
6-digit industry \times period effects	Y	Y	Y	Y	Y	Y
region (state) \times period effects	Y	Y	Y	Y	Y	Y
N (plants)	2625	2625	2625	2625	2625	2625
N (obs)	5250	5250	5250	5250	5250	5250

Notes: The dependent variables (at top) are changes between periods (after averaging within periods). The excluded instrument is the interaction between initial log employment (i.e. in the earlier period) and an indicator for the devaluation that equals 1 for change from period 1 to period 2, and 0 otherwise. Panel A reports the first stage in column (1) and reduced-form results in columns (2)-(6). Panel B reports the corresponding IV regressions. Sample is EIA-IMSS panel. Export share is fraction of total sales derived from exports. Robust standard errors in parentheses. * 10% level, **5% level, ***1% level.

Table 3.5. Cross-sectional correlations, alternative proxies, period 1 (1992-1994)

	(1)	(2)	(3)	(4)	(5)	(6)
Capability proxy	export share	log K/L	log avg. hourly wage (EIA)	avg. log daily wage (IMSS)	plant component	person component
log employment (hours)	0.029*** (0.003)	0.182*** (0.026)	0.126*** (0.011)	0.095*** (0.007)	0.074*** (0.006)	0.020*** (0.004)
log domestic sales	0.004 (0.003)	0.364*** (0.020)	0.190*** (0.007)	0.119*** (0.005)	0.082*** (0.004)	0.037*** (0.003)
log TFP (Levinsohn-Petrin)	0.020*** (0.003)	0.097*** (0.029)	0.184*** (0.011)	0.111*** (0.007)	0.076*** (0.005)	0.035*** (0.004)
predicted export share	0.274*** (0.029)	6.150*** (0.211)	2.005*** (0.084)	1.257*** (0.056)	0.893*** (0.044)	0.360*** (0.037)
6-digit industry effects	Y	Y	Y	Y	Y	Y
region (state) effects	Y	Y	Y	Y	Y	Y

Notes: Each panel reports six regressions, all corresponding to period 1 (1992-1994). Log TFP estimated using Levinsohn and Petrin (2003) method, with value-added as outcome, employment (white collar and blue collar separately) and log capital as covariates, and log materials and log electricity as proxies, separately by 2-digit sector. Predicted export share estimated from tobit regression of export share on log sales, log employment and log capital-labor ratio. Sample is EIA-IMSS panel. All regressions have N=2,625 except those with log TFP proxy, which have N=2,588 (because observations with non-positive value-added have been dropped and only plants with TFP estimates in all periods included). See Appendix Table C.2.6 for coefficient estimates from estimation of log TFP and predicted export share. Robust standard errors in parentheses. *10% level, **5% level, ***1% level.

Table 3.6. First stage, reduced-form, and IV estimates, alternative proxies

	(1)	(2)	(3)	(4)
	Δ export share	Δ avg. log daily wage	Δ plant comp.	Δ avg. person comp.
A. Log domestic sales as proxy				
1. First stage and reduced form				
init. log domestic sales \times devaluation	0.008*** (0.003)	0.032*** (0.004)	0.030*** (0.006)	0.002 (0.005)
init. log domestic sales	0.006*** (0.006)	0.006*** (0.006)	0.004 (0.004)	0.002 (0.002)
2. IV				
Δ export share		3.773*** (1.266)	3.554*** (1.269)	0.247 (0.585)
init. log domestic sales		-0.016 (0.012)	-0.017 (0.012)	0.001 (0.006)
B. Log TFP (Levinsohn-Petrin) as proxy				
1. First stage and reduced form				
init. log TFP (Levinsohn-Petrin) \times devaluation	0.009*** (0.003)	0.028*** (0.005)	0.035*** (0.007)	-0.007 (0.006)
init. log TFP (Levinsohn-Petrin)	0.001 (0.001)	0.009*** (0.009)	-0.001 (-0.001)	0.010** (0.010)
2. IV				
Δ export share		2.989*** (1.067)	3.740*** (1.377)	-0.731 (0.678)
init. log TFP (Levinsohn-Petrin)		0.007 (0.007)	-0.002 (0.009)	0.010** (0.005)
C. Predicted export share as proxy				
1. First stage and reduced form				
init. Predicted export share \times devaluation	0.111*** (0.030)	0.333*** (0.046)	0.346*** (0.064)	-0.012 (0.060)
init. Predicted export share	0.027 (0.027)	0.067** (0.067)	0.020 (0.020)	0.050 (0.050)
2. IV				
Δ export share		3.013*** (0.874)	3.132*** (0.967)	-0.113 (0.515)
init. Predicted export share		-0.015 (0.072)	-0.066 (0.085)	0.053 (0.050)
6-digit industry \times period effects	Y	Y	Y	Y
region (state) \times period effects	Y	Y	Y	Y

Notes: Each panel is similar to Columns 1, 4, 5 and 6 of Table 3.4 for indicated alternative capability proxy. See notes to Table 3.5 for construction of log TFP and predicted export share. Sample is EIA-IMSS panel. Regressions in Panels A and C have 2,625 plants and 5,250 observations; Regressions in Panel B have 2,588 plants and 5,176 observations (because observations with negative value-added have been dropped and only plants with TFP estimates in all periods included). Robust standard errors in parentheses. *10% level, **5% level, ***1% level.

Table 3.7. Differential effect of devaluation, IMSS 1988-2002 panel

	(1) avg. log daily wage (IMSS)	(2) plant component	(3) person component
A. Cross-sectional correlations, period 1			
log employment, period 1	0.091*** (0.007)	0.077*** (0.005)	0.013*** (0.004)
	Δ avg. log daily wage (IMSS)	Δ plant component	Δ person component
B.1 Outcomes: change from period 0 to 1			
log employment, period 0 (π_1)	0.026*** (0.004)	0.017*** (0.005)	0.009* (0.005)
B.2 Outcomes: change from period 1 to 2			
log employment, period 1 (π_2)	0.042*** (0.004)	0.037*** (0.005)	0.005 (0.005)
B.3. Outcomes: change from period 2 to 3			
log employment, period 2 (π_3)	0.009*** (0.003)	0.003 (0.005)	0.006 (0.005)
C. Differences in coefficients			
$\pi_2 - \pi_1$	0.017*** (0.006)	0.021*** (0.007)	-0.004 (0.007)
$\pi_3 - \pi_2$	-0.033*** (0.005)	-0.034*** (0.007)	0.002 (0.007)
6-digit industry effects	Y	Y	Y
region (state) effects	Y	Y	Y
N	2316	2316	2316

Notes: Sample is plants in EIA-IMSS panel that are also in largest connected set in period 0 (1988-1990). Variables are from IMSS data. Period 1 is 1992-1994, period 2 is 1996-1998, period 3 is 2000-2002. Panel A reports three separate cross-sectional regressions of the dependent variable at top on average log employment in period 1, analogous to Table 3.5. Panels B1-B3 report three separate regressions each, of the dependent variable at the top (i.e. change for indicated period) on average log employment in the previous period. See discussion in Section 5 for details. All regressions include industry and state effects and have sample size indicated at bottom. Panel C reports differences in coefficients from Panel B and the standard errors on the differences. Robust standard errors in parentheses. *10% level, **5% level, ***1% level.

Table 3.8. Differential effect of devaluation, IV estimates, stayers only

	(1) △ export share	(2) △ avg. log daily wage, stayers
A. First stage and reduced form		
initial log employ. × devaluation	0.009*** (0.003)	0.024*** (0.005)
initial log employ.	0.004* (0.004)	0.019*** (0.019)
B. IV		
change in export share		2.703** (1.061)
initial log employ.		0.010 (0.009)
6-digit industry × period effects	Y	Y
region (state) × Period effects	Y	Y
N	2623	2623

Notes: Table is similar to Columns 1 and 4 of Table 3.4, but here (in Column 2) the plant-level average wage has been computed using only workers who did not change jobs between periods (stayers). Sample is plants in EIA-IMSS panel for which a change in average log wage for stayers could be constructed both between periods 1 and 2 and between periods 2 and 3. (Of the 2,625 EIA-IMSS panel plants, two did not have stayers between one pair of these periods.) Export share is fraction of total sales derived from exports. Robust standard errors in parentheses. *10% level, **5% level, ***1% level.

Bibliography

- John A Abowd and Thomas Lemieux. The Effects of Product Market Competition on Collective Bargaining Agreements: The Case of Foreign Competition in Canada. The Quarterly Journal of Economics, 108(4):983–1014, 1993.
- John M Abowd, Francis Kramarz, and David N Margolis. High wage workers and high wage firms. Econometrica, 67(2):251–333, 1999.
- John M. Abowd, Robert H. Creecy, and Francis Kramarz. Computing Person and Firm Effects Using Linked Longitudinal Employer-Employee Data. LEHD Technical Paper 2002-06, 2002.
- John M Abowd, Francis Kramarz, Paul Lengermann, and Sébastien Pérez-Duarte. Are Good Workers Employed by Good Firms? A Test of a Simple Assortative Matching Model for France and the United States. Unpub. paper, Cornell., 2004.
- John M Abowd, Kevin L McKinney, and Nellie L Zhao. Earnings Inequality and Mobility Trends in the United States: Nationally Representative Estimates from Longitudinally Linked Employer-Employee Data. Journal of Labor Economics, 36(S1):S183–S300, 2018.
- Philippe Aghion, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen. On the returns to invention within firms: Evidence from finland.
- George A Akerlof and Janet L Yellen. The Fair Wage-Effort Hypothesis and Unemployment. The Quarterly Journal of Economics, 105(2):255–283, 1990.
- Jorge Alvarez, Felipe Benguria, Niklas Engbom, and Christian Moser. Firms and the Decline in Earnings Inequality in Brazil. American Economic Journal: Macroeconomics, 10(1):149–89, 2018.
- Roberto Álvarez and Ricardo A López. Skill upgrading and the real exchange rate. World Economy, 32(8):1165–1179, 2009.
- Mary Amiti and Donald R Davis. Trade, Firms, and Wages: Theory and Evidence. The Review of economic studies, 79(1):1–36, 2012.
- Theodore W Anderson, Herman Rubin, et al. Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations. The Annals of Mathematical Statistics, 20(1):46–63, 1949.
- Martyn J Andrews, Len Gill, Thorsten Schank, and Richard Upward. High wage workers and low wage firms: negative assortative matching or limited mobility bias? Journal of the Royal Statistical Society: Series A (Statistics in Society), 171(3):673–697, 2008.
- Joshua D Angrist and Alan B Krueger. Empirical Strategies in Labor Economics. In Handbook of labor economics, volume 3, pages 1277–1366. Elsevier, 1999.

- Joshua D Angrist and Jörn-Steffen Pischke. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton university press, 2008.
- Bruno César Araújo and Lourenço S Paz. The Effects of Exporting on Wages: An Evaluation Using the 1999 Brazilian Exchange Rate Devaluation. Journal of Development Economics, 111:1–16, 2014.
- Orley Ashenfelter. Estimating the effect of training programs on earnings. The Review of Economics and Statistics, pages 47–57, 1978.
- Erling Barth, Alex Bryson, James C Davis, and Richard Freeman. It Where You Work: Increases in the Dispersion of Earnings Across Establishments and Individuals in the United States. Journal of Labor Economics, 34(S2):S67–S97, 2016.
- Erling Barth, James Davis, and Richard B Freeman. Augmenting the Human Capital Earnings Equation with Measures of Where People Work. Journal of Labor Economics, 36(S1):S71–S97, 2018.
- Paulo Bastos, Joana Silva, and Eric Verhoogen. Export Destinations and Input Prices. American Economic Review, 108(2):353–92, February 2018. doi: 10.1257/aer.20140647. URL <http://www.aeaweb.org/articles?id=10.1257/aer.20140647>.
- Daniel Baumgarten. Exporters and the rise in wage inequality: Evidence from german linked employer–employee data. Journal of International Economics, 90(1):201–217, 2013.
- Matthieu Bellon. Trade Liberalization and Inequality: A Dynamic Model with Firm and Worker Heterogeneity. Unpub. paper, Columbia University, 2017.
- Andrew B Bernard and J Bradford Jensen. Exporters, Jobs, and Wages in US Manufacturing: 1976-1987. Brookings papers on economic activity. Microeconomics, 1995:67–119, 1995.
- Andrew B Bernard and J Bradford Jensen. Exceptional exporter performance: cause, effect, or both? Journal of international economics, 47(1):1–25, 1999.
- Federico Bernini, Lucas Figal Garone, and Alessandro Maffioli. International Quality Certification: Signaling to Whom? Impact on Firm Performance in Latin America and the Caribbean. Mimeo, 2017.
- Marianne Bertrand and Sendhil Mullainathan. Enjoying the Quiet Life? Corporate Governance and Managerial Preferences. Journal of political Economy, 111(5):1043–1075, 2003.
- Nicholas Bloom and John Van Reenen. Measuring and Explaining Management Practices across Firms and Countries. The Quarterly Journal of Economics, 122(4):1351–1408, 2007.
- Nicholas Bloom and John Van Reenen. Why Do Management Practices Differ across Firms and Countries? Journal of Economic Perspectives, 24(1):203–24, 2010.
- Nicholas Bloom, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts. Does Management Matter? Evidence from India. The Quarterly Journal of Economics, 128(1):1–51, 2013.

- Nicholas Bloom, Raffaella Sadun, and John Van Reenen. Management as a Technology? NBER Working Paper 22327, 2016.
- Nicholas Bloom, Aprajit Mahajan, David McKenzie, and John Roberts. Do Management Interventions Last? Evidence from India. NBER Working paper 24249, 2018.
- Nicholas Bloom, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen. What Drives Differences in Management Practices? American Economic Review, 109(5):1648–83, 2019.
- Olivier Boiral, Laurence Guillaumie, Iñaki Heras-Saizarbitoria, and Christian Valery Tayo Tene. Adoption and Outcomes of ISO 14001: A Systematic Review. International Journal of Management Reviews, 20(2):411–432, 2018.
- John Bound, David A Jaeger, and Regina M Baker. Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. Journal of the American statistical association, 90(430):443–450, 1995.
- Irene Brambilla, Daniel Lederman, and Guido Porto. Exports, export destinations, and skills. American Economic Review, 102(7):3406–38, 2012.
- Irene Brambilla, Nicolas Depetris Chauvin, and Guido Porto. Examining the export wage premium in developing countries. Review of International Economics, 25(3):447–475, 2017.
- Boris Braun. Building Global Institutions: The Diffusion of Management Standards in the World Economy – An Institutional Perspective. In Linking Industries Across the World. Ashgate, 2005.
- Raffaello Bronzini and Eleonora Iachini. Are incentives for r&d effective? evidence from a regression discontinuity approach. American Economic Journal: Economic Policy, 6(4):100–134, November 2014. doi: 10.1257/pol.6.4.100. URL <http://www.aeaweb.org/articles?id=10.1257/pol.6.4.100>.
- Raffaello Bronzini and Paolo Piselli. The impact of r&d subsidies on firm innovation. Research Policy, 45(2):442–457, 2016.
- Miriam Bruhn and Bilal Zia. Stimulating Managerial Capital in Emerging Markets: The Impact of Business Training for Young Entrepreneurs. Journal of Development Effectiveness, 5(2): 232–266, 2013.
- Miriam Bruhn, Dean Karlan, and Antoinette Schoar. The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico. Journal of Political Economy, 126(2):635–687, 2018.
- Erik Brynjolfsson and Paul Milgrom. Complementarity in Organizations. In The Handbook of Organizational Economics, pages 11–55. Princeton University Press Princeton, 2013.
- Paula Bustos. The Impact of Trade Liberalization on Skill Upgrading: Evidence from Argentina. Evidence from Argentina, 2011a.

- Paula Bustos. Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms. American Economic Review, 101(1):304–40, 2011b.
- Jing Cai and Adam Szeidl. Interfirm Relationships and Business Performance. The Quarterly Journal of Economics, 133(3):1229–1282, 2017.
- Elisa Calza, Micheline Goedhuys, and Neda Trifković. Drivers of Productivity in Vietnamese SMEs: The Role of Management Standards and Innovation. Economics of Innovation and New Technology, 28(1):23–44, 2019.
- David Card, Jörg Heining, and Patrick Kline. Workplace heterogeneity and the rise of west german wage inequality. The Quarterly journal of economics, 128(3):967–1015, 2013.
- David Card, Ana Rute Cardoso, and Patrick Kline. Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. The Quarterly Journal of Economics, 131(2):633–686, 2015.
- David Card, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. Firms and labor market inequality: Evidence and some theory. Journal of Labor Economics, 36(S1):S13–S70, 2018.
- Sara G Castellanos, Rodrigo García-Verdú, and David S Kaplan. Nominal Wage Rigidities in Mexico: Evidence from Social Security Records. Journal of Development Economics, 75(2): 507–533, 2004.
- Victoria Castillo, Lucas Figal Garone, Alessandro Maffioli, Sofia Rojo, and Rodolfo Stucchi. Knowledge spillovers through labour mobility: An employer–employee analysis. The Journal of Development Studies, pages 1–20, 2019.
- Matias D Cattaneo, Brigham R Frandsen, and Rocio Titiunik. Randomization Inference in the Regression Discontinuity Design: An Application to Party Advantages in the US Senate. Journal of Causal Inference, 3(1):1–24, 2015.
- Matias D Cattaneo, Nicolás Idrobo, and Rocio Titiunik. A Practical Introduction to Regression Discontinuity Designs: Volume II. Mimeo, 2018.
- Xavier Cirera and William Maloney. The Innovation Paradox. The World Bank, 2017.
- A Kerem Coşar, Nezih Guner, and James Tybout. Firm Dynamics, Job Turnover, and Wage Distributions in an Open Economy. American Economic Review, 106(3):625–63, 2016.
- Carl Davidson, Steven J Matusz, and Andrei Shevchenko. Globalization and Firm Level Adjustment with Imperfect Labor Markets. Journal of international Economics, 75(2):295–309, 2008.
- Carl Davidson, Fredrik Heyman, Steven Matusz, Fredrik Sjöholm, and Susan Chun Zhu. Globalization and Imperfect Labor Market Sorting. Journal of International Economics, 94(2):177–194, 2014.
- William T Dickens and Lawrence F Katz. Interindustry wage differences and industry characteristics. In Kevin Lang and Jonathan Leonard, editors, Unemployment and the Structure of Labor Markets. Blackwell, 1987.

- Alejandro Drexler, Greg Fischer, and Antoinette Schoar. Keeping It Simple: Financial Literacy and Rules of Thumb. American Economic Journal: Applied Economics, 6(2):1–31, 2014.
- Esther Duflo, Pascaline Dupas, and Michael Kremer. The Impact of Free Secondary Education: Experimental Evidence from Ghana. Mimeo. Stanford University, 2019.
- Hartmut Egger and Udo Kreickemeier. Firm Heterogeneity and the Labor Market Effects of Trade Liberalization. International Economic Review, 50(1):187–216, 2009.
- Hartmut Egger and Udo Kreickemeier. Fairness, Trade, and Inequality. Journal of International Economics, 86(2):184–196, 2012.
- Pablo D Fajgelbaum. Labor market frictions, firm growth, and international trade. Unpub. paper, UCLA, 2016.
- Gabriel Felbermayr, Julien Prat, and Hans-Jörg Schmerer. Globalization and labor market outcomes: Wage bargaining, search frictions, and firm heterogeneity. Journal of Economic theory, 146(1):39–73, 2011.
- Andrew D Foster and Mark R Rosenzweig. Microeconomics of Technology Adoption. The Annual Review of Economics, 2(1):395–424, 2010.
- Brigham R Frandsen. Party Bias in Union Representation Elections: Testing for Manipulation in the Regression Discontinuity Design when the Running Variable is Discrete. Mimeo, 2016.
- Judith A. Frías, David S. Kaplan, and Eric A. Verhoogen. Exports and wage premia: Evidence from Mexican employer-employee data. Unpub. paper, Columbia University, 2009.
- Judith A Frías, David S Kaplan, and Eric Verhoogen. Exports and Within-Plant Wage Distributions: Evidence from Mexico. American Economic Review Papers and Proceedings, 102(3): 435–40, 2012.
- Andrew Garin and Filipe Silvério. How Responsive are Wages to Demand within the Firm? Evidence from Idiosyncratic Export Demand Shocks. Unpub. paper, 2018.
- Xavier Giné and Ghazala Mansuri. Money or Ideas? A Field Experiment on Constraints to Entrepreneurship in Rural Pakistan. Mimeo. World Bank., 2014.
- Michela Giorcelli. The Long-Term Effects of Management and Technology Transfers. American Economic Review, 109(1):121–52, 2019.
- Edwin Goñi and William F Maloney. Why Don't Poor Countries do R&D? Varying Rates of Factor Returns Across the Development Process. European Economic Review, 94:126–147, 2017.
- Bronwyn H Hall and Alessandro Maffioli. Evaluating the impact of technology development funds in emerging economies: evidence from latin america. The European Journal of Development Research, 20(2):172–198, 2008.
- Elhanan Helpman, Oleg Itskhoki, and Stephen Redding. Inequality and unemployment in a global economy. Econometrica, 78(4):1239–1283, 2010.

- Elhanan Helpman, Oleg Itskhoki, Marc-Andreas Muendler, and Stephen J Redding. Trade and inequality: From theory to estimation. The Review of Economic Studies, 84(1):357–405, 2017.
- Iñaki Heras-Saizarbitoria and Olivier Boiral. ISO 9001 and ISO 14001: Towards a Research Agenda on Management System Standards. International Journal of Management Reviews, 15(1):47–65, 2013.
- Iñaki Heras-Saizarbitoria, Olivier Boiral, German Arana, and Erlantz Allur. OHSAS 18001 Certification and Work Accidents: Shedding Light on the Connection. Journal of Safety Research, 68:33–40, 2019.
- Yuki Higuchi, Edwin P Mhede, and Tetsushi Sonobe. Short-and Medium-run Impacts of Management Training: An Experiment in Tanzania. World Development, 114:220–236, 2019.
- Douglas Holtz-Eakin, Whitney Newey, and Harvey S Rosen. Estimating Vector Autoregressions with Panel Data. Econometrica, pages 1371–1395, 1988.
- Sabrina T. Howell. Financing innovation: Evidence from r&d grants. American Economic Review, 107(4):1136–64, April 2017. doi: 10.1257/aer.20150808. URL <http://www.aeaweb.org/articles?id=10.1257/aer.20150808>.
- Chang-Tai Hsieh and Peter J Klenow. Misallocation and manufacturing tfp in china and india. The Quarterly journal of economics, 124(4):1403–1448, 2009.
- Chang-Tai Hsieh and Peter J Klenow. The Life Cycle of Plants in India and Mexico. The Quarterly Journal of Economics, 129(3):1035–1084, 2014.
- David Hummels, Rasmus Jørgensen, Jakob Munch, and Chong Xiang. The wage effects of offshoring: Evidence from danish matched worker-firm data. American Economic Review, 104(6):1597–1629, 2014.
- Casey Ichniowski, Kathryn Shaw, and Giovanna Prennushi. The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines. The American Economic Review, pages 291–313, 1997.
- Alfonso Irarrazabal, Andreas Moxnes, and Karen Helene Ulltveit-Moe. Heterogeneous Firms or Heterogeneous Workers? Implications for Exporter Premiums and the Gains from Trade. Review of Economics and Statistics, 95(3):839–849, 2013.
- Ivan T Kandilov. Do Exporters Pay Higher Wages? Plant-Level Evidence from an Export Refund Policy in Chile. The World Bank Economic Review, 23(2):269–294, 2009.
- David Kaplan, Gabriel Martinez, and Raymond Robertson. Worker and job flows in mexico. Unpub. paper, ITAM, 2004.
- David Kaplan, Gabriel Martinez, and Raymond Robertson. What Happens to Wages after Displacement? Economia, pages 197 – 242, 2010.

- David S Kaplan, Gabriel Martínez González, and Raymond Robertson. Mexican employment dynamics: Evidence from matched firm-worker data. World Bank Policy Research working paper No. 4433, 2007.
- Dean Karlan and Martin Valdivia. Teaching Entrepreneurship: Impact of Business Training on Microfinance Clients and Institutions. Review of Economics and statistics, 93(2):510–527, 2011.
- Dean Karlan, Ryan Knight, and Christopher Udry. Consulting and Capital Experiments with Microenterprise Tailors in Ghana. Journal of Economic Behavior & Organization, 118:281–302, 2015.
- Lawrence F Katz and Lawrence H Summers. Industry Rents: Evidence and Implications. Brookings Papers on Economic Activity. Microeconomics, 1989:209–290, 1989.
- Michael W Klein, Christoph Moser, and Dieter M Urban. Exporting, Skills and Wage Inequality. Labour Economics, 25:76–85, 2013.
- Patrick Kline, Neviana Petkova, Heidi Williams, and Owen Zidar. Who profits from patents? rent-sharing at innovative firms. The Quarterly Journal of Economics, 134(3):1343–1404, 2019.
- Jeffrey R Kling, Jeffrey B Liebman, and Lawrence F Katz. Experimental Analysis of Neighborhood Effects. Econometrica, 75(1):83–119, 2007.
- Michal Kolesár and Christoph Rothe. Inference in Regression Discontinuity Designs with a Discrete Running Variable. American Economic Review, 108(8):2277–2304, 2018.
- Pravin Krishna, Jennifer P Poole, and Mine Zeynep Senses. Wage Effects of Trade Reform with Endogenous Worker Mobility. Journal of International Economics, 93(2):239–252, 2014.
- Alan B Krueger and Lawrence H Summers. Efficiency Wages and the Inter-Industry Wage Structure. Econometrica: Journal of the Econometric Society, pages 259–293, 1988.
- Todd Johnson Kumler, Eric A Verhoogen, and Judith Frias. Enlisting workers in monitoring firms: Payroll tax compliance in Mexico. NBER working paper no. 19385, updated April 2015., 2015.
- Edward P Lazear and Kathryn L Shaw. Introduction: Firms and the Distribution of Income: The Roles of Productivity and Luck. Journal of Labor Economics, 36(S1):S1–S12, 2018.
- David S Lee and David Card. Regression Discontinuity Inference with Specification Error. Journal of Econometrics, 142(2):655–674, 2008.
- James Levinsohn and Amil Petrin. Estimating Production Functions Using Inputs to Control for Unobservables. The review of economic studies, 70(2):317–341, 2003.
- Gladys López-Acevedo and Hong W Tan. Impact evaluation of small and medium enterprise programs in Latin America and the Caribbean. The World Bank, 2011.
- Mario Macis and Fabiano Schivardi. Exports and Wages: Rent Sharing, Workforce Composition, or Returns to Skills? Journal of Labor Economics, 34(4):945–978, 2016.

- Kalina Manova. Credit Constraints, Heterogeneous Firms, and International Trade. Review of Economic Studies, 80(2):711–744, 2012.
- David C Maré and Dean Hyslop. Worker-Firm Heterogeneity and Matching: An Analysis Using Worker and Firm Fixed Effects Estimated from LEED. Unpub. paper., 2006.
- Oliver Masakure, Spencer Henson, and John Cranfield. Standards and Export Performance in Developing Countries: Evidence from Pakistan. The Journal of International Trade & Economic Development, 18(3):395–419, 2009.
- Justin McCrary. Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test. Journal of econometrics, 142(2):698–714, 2008.
- David McKenzie and Christopher Woodruff. What are We Learning from Business Training and Entrepreneurship Evaluations around the Developing World? The World Bank Research Observer, 29(1):48–82, 2014.
- David McKenzie and Christopher Woodruff. Business Practices in Small Firms in Developing Countries. Management Science, 63(9):2773–3145, 2017.
- Marc J Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. econometrica, 71(6):1695–1725, 2003.
- Jakob Roland Munch and Jan Rose Skaksen. Human Capital and Wages in Exporting Firms. Journal of International Economics, 75(2):363–372, 2008.
- Steven Olley and Ariel Pakes. The Dynamics of Productivity in the Telecommunications Equipment Industry. Econometrica, 64(6):1263–1297, 1996. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2171831>.
- Hammad Riaz, Abubakr Saeed, Muhammad Saad Baloch, Zeeshan Ahmad Khan, et al. Valuation of Environmental Management Standard ISO 14001: Evidence from an Emerging Market. Journal of Risk and Financial Management, 12(1):21, 2019.
- Paulo Sampaio, Pedro Saraiva, and António Guimarães Rodrigues. ISO 9001 Certification Research: Questions, Answers and Approaches. International Journal of Quality & Reliability Management, 26(1):38–58, 2009.
- Thorsten Schank, Claus Schnabel, and Joachim Wagner. Do exporters really pay higher wages? first evidence from german linked employer–employee data. Journal of international Economics, 72(1):52–74, 2007.
- Friedrich Schneider and Dominik H Enste. Shadow Economies: Size, Causes, and Consequences. Journal of Economic Literature, 38(1):77–114, 2000.
- Jae Song, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. Firming Up Inequality. The Quarterly Journal of Economics, 134(1):1–50, 2019.
- Isaac Sorkin. Ranking firms using revealed preference. The quarterly journal of economics, 133(3):1331–1393, 2018.

- Douglas Staiger and James H. Stock. Instrumental Variables Regression with Weak Instruments. Econometrica, 65(3):557–586, 1997. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/2171753>.
- Ying Sun and Wenjing Ouyang. International Standards for Exporting Firms: Evidence from China. Journal of Applied Business Research, 30(6):1753, 2014.
- Chad Syverson. Product Substitutability and Productivity Dispersion. Review of Economics and Statistics, 86(2):534–550, 2004.
- Chad Syverson. What Determines Productivity? Journal of Economic literature, 49(2):326–65, 2011.
- Juan José Tarí, José Francisco Molina-Azorín, and Iñaki Heras. Benefits of the ISO 9001 and ISO 14001 Standards: A Literature Review. Journal of Industrial Engineering and Management, 5(2):297–322, 2012.
- Mustafa V Uzumeri. ISO 9000 and Other Metastandards: Principles for Management Practice? Academy of Management Perspectives, 11(1):21–36, 1997.
- Martin Valdivia. Business Training Plus for Female Entrepreneurship? Short and Medium-Term Experimental Evidence from Peru. Journal of Development Economics, 113:33–51, 2015.
- John Van Reenen. The Creation and Capture of Rents: Wages and Innovation in a Panel of UK Companies. The Quarterly Journal of Economics, 111(1):195–226, 1996.
- Eric A Verhoogen. Trade, quality upgrading, and wage inequality in the mexican manufacturing sector. The Quarterly Journal of Economics, 123(2):489–530, 2008.
- Christian Volpe Martincus, Sebastián Castresana, and Tomás Castagnino. ISO standards: A Certificate to Expand Exports? Firm-level Evidence from Argentina. Review of International Economics, 18(5):896–912, 2010.

Appendix A: Additional information for chapter 1

A.1 Certification Program sample

The administrative records of *Innovate Perú* indicate that 250 applications were received and sent to reviewers to be evaluated in the first three rounds of the Certification Program. Firms could not request reconsideration of their applications, but could submit different projects. In the main text of the article, I consider all the applications received independently and use cluster standard errors at the level of the firm to avoid overestimating the precision of the estimates. This is a conservative approach as it makes more difficult to find any effect and avoids arbitrarily selecting observations. In the appendix A.5.1, I show that the results are the same if they are estimated using a restricted sample that includes only the first round in which the never-treated firms applied, and the first round in which the ever-treated won.

A possible concern could be that reviewers manipulated the score to benefit or obstruct firms applying more than once. This possibility can be tested by checking that the number of previous applications does not change discontinuously at the cutoff. As shown in table A.1, the number of previous applications is not different between firms just to the left and right of the cutoff, suggesting that there is no manipulation. This table reports OLS estimates of the difference in this variable on different sides of the the cutoff using local linear regressions with different bandwidths (columns 1 - 3), and the full sample using a linear and a cubic polynomial (columns 4 and 5, respectively). The table also reports randomized inference p-values for different windows around the cutoff (columns 6 - 8).

Table A.1: Balance of number of previous applications

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
right	-0.033 (0.198)	0.004 (0.138)	0.061 (0.143)	0.024 (0.112)	-0.158 (0.220)	0.012	0.030	0.041
rand. inf. p-val.						1.000	0.870	0.774
obs	172	219	240	250	250	75	129	153

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in this table correspond to estimates of the difference in a variable indicating the number of previous applications by the same firm. Columns 1 - 5 report OLS estimates of β_1 in equation (1.5). The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 report three values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the difference is zero. The third value is the number of observations. Estimates in column 6 use observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

A.2 Management index and complementary data sources

The management index was build using information collected in the survey using a questionnaire suited for firms in the sample. The starting point to develop the questionnaire was the one used by the U.S. Census Bureau in the Management and Organizational Practices Survey (MOPS). The main changes with respect to that instrument were: In the MOPS, there were two different questions regarding the use of bonuses for managers, and two for non-managers. One asked whether they were used and the percentage of managers/non-managers that received it. The other asked about the criteria used to determine their value. I reduced these questions to one per type of worker asking whether bonuses were used and the criteria they were based on. Additionally, in the MOPS, there were questions regarding the use of production display boards, difficulty achieving targets, and review of performance indicators by non-managers. These questions were omitted to accommodate the smaller size of the firms in the sample and the fact that the sample includes non-manufacturing firms.

With the collected management information, I built nine variables with support $[0, 1]$ measuring the adoption of different practices. These variables were grouped in categories according to whether they relate with monitoring, target-setting, or incentives. The category-level average of those variables is the score for that category, and the average of those categories is the value of the management index.

The variables created and the way in which the answers were scored is as follows (variables 1 and 2 were used to measure monitoring, variable 3 was used to measure target-setting, and variables 4 to 9 were used to measure incentives).

1. Use and number of performance indicators. Does not use=0, use between 1-2 indicators=1/4, use between 3-5=1/2, use between 6-9=3/4, use 10 or more=1.
2. Frequency with which performance indicators are reviewed. Does not use=0, annually=1/6, once per semester=2/6, quarterly=3/6, monthly=4/6, weekly=5/6, daily=1.
3. Use of explicit targets and their time horizon. Does not set explicit goals=0, only short-run (less than a year)=1/3, only long run (more than a year)=2/3, both short- and long-run=1.
4. Use of bonuses for managers and criteria they are based on. Does not use=0, use and they based on firm performance=1/2, use and they are based on individual performance=1.
5. Use of bonuses for non-managers and criteria they are based on. Does not use=0, use and they based on firm performance=1/2, use and they are based on individual performance=1.
6. Promotion methods for managers. Did not promote=0, promoted based on factors other than performance (e.g. tenure, owners trust, etc.)=1/3, promoted based on performance and other factors (e.g. tenure, owners trust, etc.)=2/3, promoted based on ability and individual performance=1.
7. Promotion methods for non-managers. Did not promote=0, promoted based on factors other than performance (e.g. tenure, owners trust, etc.)=1/3, promoted based on performance and

other factors (e.g. tenure, owners trust, etc.)=2/3, promoted based on ability and individual performance=1.

8. Time to reassign or to fire an under-performing manager. Did not reassign=0, more than six months after identifying the problem=1/2, six months or less after identifying the problem=1.
9. Time to reassign or to fire an under-performing non-manager. Did not reassign=0, more than six months after identifying the problem=1/2, six months or less after identifying the problem=1.

In the questions regarding promotion of personnel and firing of under-performing employees, the majority of firms indicated that such situation had not occurred in 2018. To construct the management index, I treated this answers as zeros. This is the same decision taken by the Census Bureau in 2015 with respect to skipped questions. The results are the same using an alternative specification in which these questions are simply omitted.

Table A.2 shows that the management index is positively correlated with performance measures. This exercise is useful to validate the management index. More specifically, the table reports the estimates of γ_1 in the regression $y = \gamma_0 + \gamma_1 mgt + v$, where y is a measure of performance in 2018 (log-employment, log-sales, or a dummy variable indicating that the firm exported) and mgt is the management index.

Table A.2: Correlation between performance and mgt. index

	log-employ.	log-sales	exported
mgt. index	2.071*** (0.307)	2.012*** (0.371)	0.248* (0.134)
Obs.	183	183	183

***, **, * indicate significance at 1%, 5%, and 10%. This table reports the OLS estimate in of γ_1 in the regression $y = \gamma_0 + \gamma_1 mgt + v$ where y is the variable indicated at the top of the corresponding column and mgt is the management index. The values in parentheses are standard errors clustered at the level of the firm.

As mentioned in the main body of the paper, in addition to the main sources of data, I also used trade information from the commercial service ADEX Data Trade and from the service *Consulta RUC* provided by the Tax Authority. From ADEX Data Trade, I manually downloaded information about the trade behavior of the firms in the sample because the service does not allow the use of the full customs data at once. Instead, firm-level information has to be looked for using the name or tax id of the company and manually downloaded in Excel spreadsheets. The trade data used in this paper comes from this source. From *Consulta RUC*, I manually downloaded information about the date in which the firms started operating and firm sector (which is reported using the ISIC rev. 3 classification). With the first piece of information I calculated the age of the firm at the moment of application. With the second piece of information I created the sector variables. The tax authority allows reporting more than one sector. This leads to some level of missclassification as some firms report activities in sectors that do not correspond to their main activity. For example, firms that manufacture goods and sell them might report wholesale or retail activities in addition to manufacturing. To overcome this problem, I considered that firms belonged to the service sector only if that was the only sector they reported, otherwise I gave priority to the non-service sector.

A.3 Balance tests using honest confidence intervals

Table A.3: Continuity of baseline covariates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	age (mo)	mse	manuf.	services	coast	export	import
A. Using M=0 (full sample)							
right	14.137	0.034	0.082	-0.017	0.020	0.041	0.042
	[-3.650	[-0.058	[-0.029	[-0.131	[-0.051	[-0.044	[-0.068
	31.925]	0.125]	0.194]	0.098]	0.091]	0.127]	0.152]
obs.	250	250	250	250	250	250	250
B. Using data-driven selection for M							
right	3.502	0.034	0.082	-0.017	0.066	0.041	0.042
	[-37.452	[-0.058	[-0.029	[-0.131	[-0.112	[-0.044	[-0.068
	44.456]	0.125]	0.194]	0.098]	0.244]	0.127]	0.152]
half. bw.	8.9	45.0	45.0	45.0	9.3	45.0	45.0
obs.	164	250	250	250	164	250	250

This table reports the estimates of β_1 in equation (1.5) taking as dependent variable the one indicated at the top of the column, and the corresponding 90% honest confidence intervals computed with the method of Kolesár and Rothe (2018). The method requires the researcher to provide the value of a parameter M that measures the non-linearity of the CEF. In turn, this parameter determines the bandwidth accepted by the method. Kolesár and Rothe (2018) proposed a data-driven method to estimate a lower bound for M and panel B reports results using that value. In many cases the selected value was zero which implies using the full sample. Panel A reports results using the $M = 0$ for all dependent variables.

A.4 Effect on employment and sales

Table A.4: Effect on log-employment

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	-0.290 (0.347)	0.063 (0.280)	0.201 (0.257)	0.202 (0.245)	-0.322 (0.391)	-0.127	0.132	0.175
rand. inf. p-val.						0.638	0.537	0.380
B. LATE estimates								
treated	-0.347 (0.420)	0.077 (0.342)	0.238 (0.304)	0.239 (0.290)	-0.425 (0.523)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-employment (headcount), winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.5: Effect on log-sales

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	-0.020 (0.369)	0.314 (0.300)	0.368 (0.277)	0.350 (0.264)	0.060 (0.411)	0.125	0.223	0.328
rand. inf. p-val.						0.690	0.371	0.162
B. LATE estimates								
treated	-0.024 (0.441)	0.383 (0.365)	0.436 (0.327)	0.415 (0.313)	0.079 (0.540)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-sales, winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

A.5 Robustness of results

A.5.1 Results using the restricted sample

Table A.6: Effect on certification probability

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.649*** (0.157)	0.740*** (0.120)	0.714*** (0.105)	0.691*** (0.095)	0.494** (0.195)	0.694	0.733	0.728
rand. inf. p-val.						0.000	0.000	0.000
B. LATE estimates								
treated	0.759*** (0.175)	0.855*** (0.135)	0.802*** (0.115)	0.776*** (0.103)	0.641*** (0.241)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a dummy variable that indicates if the proposed certification was obtained. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.7: Effect on management index

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.175* (0.091)	0.172** (0.075)	0.200*** (0.069)	0.177** (0.069)	0.248** (0.097)	0.176	0.141	0.140
rand. inf. p-val.						0.016	0.020	0.008
B. LATE estimates								
treated	0.205* (0.108)	0.199** (0.088)	0.225*** (0.078)	0.199** (0.078)	0.321** (0.134)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the management index. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.8: Effect on monitoring & target-setting

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.171* (0.090)	0.215*** (0.073)	0.255*** (0.066)	0.220*** (0.066)	0.195* (0.100)	0.188	0.178	0.182
rand. inf. p-val.						0.009	0.004	0.001
B. LATE estimates								
treated	0.200* (0.108)	0.249*** (0.085)	0.287*** (0.075)	0.248*** (0.076)	0.253* (0.135)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the subindex of monitoring and target-setting practices. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.9: Effect on incentives

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.205 (0.128)	0.148 (0.110)	0.107 (0.099)	0.134 (0.102)	0.381*** (0.146)	0.192	0.092	0.126
rand. inf. p-val.						0.126	0.431	0.205
B. LATE estimates								
treated	0.240 (0.149)	0.171 (0.126)	0.120 (0.111)	0.150 (0.114)	0.493** (0.194)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the subindex of incentive practices. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.10: Effect on log-sales per worker

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.398 (0.401)	0.508 (0.333)	0.423 (0.304)	0.344 (0.289)	0.443 (0.443)	0.508	0.343	0.380
rand. inf. p-val.						0.093	0.204	0.123
B. LATE estimates								
treated	0.466 (0.471)	0.587 (0.387)	0.475 (0.343)	0.387 (0.325)	0.574 (0.580)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-sales per worker, winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.11: Effect on log-employment

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	-0.056 (0.366)	0.269 (0.299)	0.354 (0.279)	0.366 (0.265)	0.095 (0.434)	0.094	0.270	0.312
rand. inf. p-val.						0.699	0.291	0.187
B. LATE estimates								
treated	-0.066 (0.429)	0.310 (0.343)	0.398 (0.313)	0.412 (0.296)	0.123 (0.561)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-employment (headcount), winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.12: Effect on log-sales

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.189 (0.439)	0.557 (0.370)	0.602* (0.347)	0.517 (0.331)	0.325 (0.504)	0.345	0.405	0.507
rand. inf. p-val.						0.332	0.177	0.066
B. LATE estimates								
treated	0.222 (0.513)	0.644 (0.427)	0.677* (0.389)	0.581 (0.371)	0.421 (0.652)			
C. First stage								
right	0.854*** (0.078)	0.866*** (0.051)	0.890*** (0.044)	0.890*** (0.044)	0.772*** (0.098)			
F-stat	120.165	284.976	418.150	418.150	62.182			
obs	102	137	154	157	157	46	80	92

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on log-sales, winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

A.5.2 Results using honest confidence intervals

Table A.13: Results using honest confidence intervals

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	certif.	mgt. index	monit. & targ.	incentives	log-sales per wkr.	log-employ.	log-sales
A. Using M=0 (full sample)							
right	0.636	0.153	0.192	0.064	0.324	0.127	0.305
	[0.519	[0.086	[0.123	[-0.018	[0.027	[-0.175	[-0.026
	0.752]	0.219]	0.261]	0.146]	0.621]	0.429]	0.636]
obs.	183	183	183	183	183	183	183
B. Using data-driven selection for M							
right	0.636	0.286	0.254	0.064	0.324	-0.057	0.328
	[0.519	[0.097	[0.021	[-0.018	[0.027	[-0.588	[-0.577
	0.752]	0.476]	0.486]	0.146]	0.621]	0.473]	1.232]
half. bw.	38.8	5.7	5.1	38.8	38.8	17.5	6.9
obs.	183	96	96	183	183	152	99

This table reports the estimates of β_1 in equation (1.5) taking as dependent variable the one indicated at the top of the column, and the corresponding 90% honest confidence intervals computed with the method of Kolesár and Rothe (2018). The method requires the researcher to provide the value of a parameter M that measures the non-linearity of the CEF. In turn, this parameter determines the bandwidth accepted by the method. Appendix S.2 in Kolesár and Rothe (2018) proposed a data-driven method to estimate a lower bound for M and panel B reports results using that value. In many cases the selected value was zero which implies using the full sample. Panel A reports results using the $M = 0$ for all dependent variables.

A.6 Effect on management index and subindices controlling for productivity

Table A.14: Effect on management index controlling for productivity

	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)
A. ITT estimate					
right	0.120 (0.080)	0.118* (0.063)	0.160*** (0.058)	0.145** (0.057)	0.177** (0.089)
B. LATE estimates					
treated	0.145 (0.096)	0.145* (0.077)	0.189*** (0.068)	0.171** (0.067)	0.232* (0.120)
C. First stage					
right	0.831*** (0.080)	0.818*** (0.054)	0.846*** (0.047)	0.846*** (0.047)	0.762*** (0.099)
F-stat	107.770	231.539	329.624	329.652	59.214
obs	122	160	179	183	183

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the management index controlling for log-sales per worker winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5) including log-sales per worker as a control. In panel B, these columns report the IV estimates of β_1 in equation (1.6) including log-sales per worker as a control. In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial.

Table A.15: Effect on monitoring & target-setting controlling for productivity

	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)
A. ITT estimate					
right	0.124 (0.080)	0.164*** (0.063)	0.204*** (0.058)	0.186*** (0.057)	0.163* (0.089)
B. LATE estimates					
treated	0.149 (0.098)	0.200** (0.079)	0.242*** (0.070)	0.220*** (0.069)	0.213* (0.121)
C. First stage					
right	0.831*** (0.080)	0.818*** (0.054)	0.846*** (0.047)	0.846*** (0.047)	0.762*** (0.099)
F-stat	107.770	231.539	329.624	329.652	59.214
obs	122	160	179	183	183

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the subindex of monitoring and target-setting practices controlling for log-sales per worker winsorized at the 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5) including log-sales per worker as a control. In panel B, these columns report the IV estimates of β_1 in equation (1.6) including log-sales per worker as a control. In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial.

Table A.16: Effect on incentives controlling for productivity

	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)
A. ITT estimate					
right	0.045 (0.093)	0.037 (0.074)	0.078 (0.070)	0.054 (0.066)	0.101 (0.112)
B. LATE estimates					
treated	0.054 (0.112)	0.046 (0.090)	0.092 (0.082)	0.063 (0.078)	0.132 (0.146)
C. First stage					
right	0.831*** (0.080)	0.818*** (0.054)	0.846*** (0.047)	0.846*** (0.047)	0.762*** (0.099)
F-stat	107.770	231.539	329.624	329.652	59.214
obs	122	160	179	183	183

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on the subindex of incentive practices controlling for log-sales per worker winsorized at 10th and 90th percentiles. In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5) including log-sales per worker as a control. In panel B, these columns report the IV estimates of β_1 in equation (1.6) including log-sales per worker as a control. In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parentheses are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial.

A.7 Changes in recent years

A.7.1 Questions per category of change

In the survey, I asked whether the following changes had occurred in the firm. The possible answers were yes or no.

Changes in organizational structure:

- Some positions were eliminated.
- Some positions were created.
- Responsibilities of some positions were changed.

Changes in the workforce:

- Some position were eliminated.
- Some positions were created.
- Responsibilities of some positions were changed.

Improvement in machinery and infrastructure:

- New machines were bought.
- Business infrastructure was renewed.

Reductions in cost:

- Labor cost was reduced.
- Materials cost was reduced.
- Energy cost was reduced.

Change in customers:

- New local customer.
- New foreign customer.
- Inserted into the mining, forestry, or fishing industry.
- Started selling to the state.

Change in suppliers:

- Stopped dealing with a supplier.
- Started dealing with a new supplier.

A.7.2 Results for other changes

Table A.17: Effect on changes in customers

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.224 (0.177)	0.288** (0.138)	0.245** (0.124)	0.188 (0.117)	0.202 (0.208)	0.212	0.190	0.225
rand. inf. p-val.						0.117	0.083	0.026
B. LATE estimates								
treated	0.268 (0.211)	0.352** (0.169)	0.290** (0.148)	0.223 (0.138)	0.266 (0.275)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a Kling et al. (2007)-type index summarizing the questions related with changes in customers (see appendix A.7.1 for details about these questions). In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.18: Effect on changes in suppliers

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.306 (0.272)	0.400* (0.205)	0.321 (0.199)	0.315* (0.180)	0.485 (0.320)	0.339	0.234	0.299
rand. inf. p-val.						0.078	0.160	0.051
B. LATE estimates								
treated	0.365 (0.325)	0.488* (0.251)	0.380 (0.235)	0.374* (0.212)	0.639 (0.428)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a Kling et al. (2007)-type index summarizing the questions related with changes in suppliers (see appendix A.7.1 for details about these questions). In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.19: Effect on changes in the workforce

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.118 (0.227)	0.164 (0.178)	0.183 (0.157)	0.168 (0.146)	0.111 (0.265)	0.101	0.161	0.202
rand. inf. p-val.						0.590	0.284	0.145
B. LATE estimates								
treated	0.141 (0.269)	0.201 (0.216)	0.217 (0.185)	0.199 (0.172)	0.147 (0.347)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a Kling et al. (2007)-type index summarizing the questions related with changes in the workforce (see appendix A.7.1 for details about these questions). In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.20: Effect on cost reduction

	Usual inference					Randomized inference		
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	-0.053 (0.267)	-0.190 (0.193)	-0.049 (0.186)	-0.147 (0.172)	-0.218 (0.284)	-0.178	-0.073	-0.149
rand. inf. p-val.						0.321	0.605	0.294
B. LATE estimates								
treated	-0.064 (0.320)	-0.232 (0.237)	-0.058 (0.221)	-0.174 (0.205)	-0.287 (0.379)			
C. First stage								
right	0.837*** (0.080)	0.819*** (0.054)	0.844*** (0.047)	0.844*** (0.047)	0.759*** (0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a Kling et al. (2007)-type index summarizing the questions related with reductions in cost (see appendix A.7.1 for details about these questions). In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

Table A.21: Effect on changes in organizational structure

	Usual inference				Randomized inference			
	bw=10 (1)	bw=20 (2)	bw=30 (3)	full (4)	full (5)	w=2mp (6)	w=5 (7)	w=7.5 (8)
A. ITT estimate								
right	0.370*	0.368**	0.401***	0.315**	0.304	0.232	0.313	0.292
	(0.212)	(0.162)	(0.154)	(0.141)	(0.241)			
rand. inf. p-val.						0.134	0.023	0.022
B. LATE estimates								
treated	0.442*	0.449**	0.475***	0.373**	0.401			
	(0.252)	(0.196)	(0.181)	(0.165)	(0.317)			
C. First stage								
right	0.837***	0.819***	0.844***	0.844***	0.759***			
	(0.080)	(0.054)	(0.047)	(0.047)	(0.099)			
F-stat	108.961	232.313	329.115	329.115	58.810			
obs	122	160	179	183	183	58	96	111

***, **, * indicate significance at 1%, 5%, and 10%. The estimates reported in panels A and B of this table correspond to the effect of the treatment on a Kling et al. (2007)-type index summarizing the questions related with changes in the organizational structure (see appendix A.7.1 for details about these questions). In panel A, columns 1 - 5 report OLS estimates of β_1 in equation (1.5). In panel B, these columns report the IV estimates of β_1 in equation (1.6). In panel C, they report the estimates of the first stage corresponding to the IV estimates in panel B. The F-stat of that first stage is reported below panel C. The values in parenthesis are standard errors clustered at the level of the firm. Estimates in columns 1, 2 and 3 use bandwidths of 10, 20, and 30 points respectively, uniform kernel, and linear polynomial. Column 4 reports similar estimates using the full sample. Column 5 reports similar estimates using the full sample and a cubic polynomial. Columns 6 - 8 in panel A report two values. The first value is the mean difference of the dependent variable between observations located to the left and right of the cutoff. The second value is the randomized inference p-value of a test of the sharp null hypothesis that the effect of the dependent variable is zero. Estimates in column 6 uses observations within two mass points of the cutoff. Estimates in columns 7 and 8 use observations within 5 and 7.5 points of the cutoff.

A.8 Exploiting complementarity under financial constraints

In the main body of the paper, I illustrated a case in which a subsidy was used to modify the incentives of a firm. In this subsection, I show how a subsidy can be used to unleash a technological upgrading process desired by the firm, but obstructed by the presence of financial constraints. To be more specific, I study a case in which the unrestricted optimal action is to adopt t_1 and t_2 , but financial constraints impede these actions. I introduce financial frictions in the model by imposing the restriction that the firm cannot spend more than R in technology adoption in any given period. A lower value of R represents a tighter financial restriction¹. If R is small enough, the firm would not be able to afford the adoption cost of technology t_1 , $(\phi_{t_1} - 1)^2$, in any period and would keep t_0 until production time.

First, let's describe the set of possible value of possible values of (ϕ_{t_1}, ϕ_{t_2}) that would make adopting t_1 and t_2 the optimal action path in the absence of financial constraints. This baseline situation would occur if conditions (A.1) and (A.2) hold.

$$\phi_{t_2}b - (\phi_{t_2} - \phi_{t_1})^2 - (\phi_{t_1} - 1)^2 > \phi_{t_1}b - (\phi_{t_1} - 1)^2 \quad (\text{A.1})$$

$$\phi_{t_2}b - (\phi_{t_2} - \phi_{t_1})^2 - (\phi_{t_1} - 1)^2 > b \quad (\text{A.2})$$


Condition (A.1) ensures that, if t_1 is reached, the firm will want to continue upgrading to t_2 . Condition (A.2) guarantees that the option of upgrading to t_1 and then to t_2 is better than simply keeping t_0 ². Figure A.1 represents graphically these conditions in a (ϕ_{t_1}, ϕ_{t_2}) -plane. They are satisfied in the area ABCGD. Condition (A.1) is satisfied between the two parallel lines L1 and

¹Note that the financial restriction applies to expenses in technology adoption, but not to the purchase of inputs. This is reasonable as inputs can be financed with credit lines for working capital, which are easier to access than credit for technology adoption. A similar assumption is made in the baseline model presented in Manova (2012), in which the credit constraint affects the payment of the export fixed costs, but not the purchase of inputs.

²When writing condition (A.2), I have assumed that the optimal action in lower node of the second period in figure 1.1 is keeping t_0 . Another possibility would be that upgrading to t_2 were optimal in that node. In that case condition (A.2) would need to be substituted with $\phi_{t_1}b - (\phi_{t_1} - 1)^2 > b$. Under condition (A.1), this implies condition (A.2). The analysis here focuses on the more general case.

L2, and condition (A.2) holds to the left of the curve L3. (ϕ_{t_1}, ϕ_{t_2}) cannot be below L2 because technological downgrading is never optimal. If (ϕ_{t_1}, ϕ_{t_2}) were above L1 and to the left of L3, the firm would upgrade to t_1 in either the first or second period, but would not continue to t_2 . If (ϕ_{t_1}, ϕ_{t_2}) were to the right of L3, reaching t_1 from the initial t_0 would be too costly, and the firm would just keep t_0 until production time.

Figure A.1: Possible parameter configurations



Second, let's introduce financial constraints. This means putting a restriction to the sequence of technologies that are possible for the firm. In particular, if the firm cannot spend more than R in any given period; then it would not be able to adopt any technology with productivity higher than $1 + \sqrt{R}$ in the first period, restriction represented with the line L4. Similarly, given a first period technology with productivity ϕ_{t_1} ; the firm would not be able to reach any technology beyond $\phi_{t_2} = \phi_{t_1} + \sqrt{R}$, restriction represented by L5.

Consider a situation in which the parameters ϕ_{t_1} and ϕ_{t_2} are as in point M. In this case, the

unrestricted optimal action path of the firm would be to adopt t_1 and then t_2 ; but due to the financial constraint, that would not be possible as the firm would not be able to upgrade to t_1 in the first period. A subsidy for projects to upgrade to t_1 in the first period would move the restriction L4 to the right up to the point in which upgrading to t_1 becomes feasible. With this support the firm would adopt t_1 and then it would upgrade to t_2 without additional subsidy. This second upgrade is possible without additional support because technology t_2 is close to t_1 (M is below L5 in figure A.1), hence the financial constraint would not be able to prevent the adoption of t_2 once t_1 has been reached. This policy action would be welfare improving because, with the subsidy, the firm would get $\phi_{t_2}b - (\phi_{t_2} - \phi_{t_1})^2 - (\phi_{t_1} - 1)^2 + S$ instead of b . Because $\phi_{t_2}b - (\phi_{t_2} - \phi_{t_1})^2 - (\phi_{t_1} - 1)^2 > b$ by assumption, the firm would be able to pay back the subsidy if required, and still be better off.

A.9 Proof of supermodularity

Let's call T_1 and T'_1 to two possible values for the technology adopted in the first period such that $(\phi_{T'_1} > \phi_{T_1})$. Similarly T_2 and T'_2 are two possible values of the technology in the second period $(\phi_{T'_2} > \phi_{T_2})$. In this case, supermodularity of the profit function means that:

$$\begin{aligned} \pi(\phi_{T'_1}, \phi_{T'_2}) - \pi(\phi_{T'_1}, \phi_{T_2}) &> \pi(\phi_{T_1}, \phi_{T'_2}) - \pi(\phi_{T_1}, \phi_{T_2}) \iff \\ \phi_{T'_2}b - (\phi_{T'_2} - \phi_{T'_1})^2 - \phi_{T_2}b + (\phi_{T_2} - \phi_{T'_1})^2 &> \phi_{T'_2}b - (\phi_{T'_2} - \phi_{T_1})^2 - \phi_{T_2}b + (\phi_{T_2} - \phi_{T_1})^2 \iff \\ -(\phi_{T'_2}^2 - 2\phi_{T'_2}\phi_{T'_1} + \phi_{T'_1}^2) + (\phi_{T_2}^2 - 2\phi_{T_2}\phi_{T'_1} + \phi_{T'_1}^2) &> -(\phi_{T'_2}^2 - 2\phi_{T_1}\phi_{T'_2} + \phi_{T_1}^2) + (\phi_{T_2}^2 - 2\phi_{T_1}\phi_{T_2} + \phi_{T_1}^2) \iff \\ 2\phi_{T'_2}\phi_{T'_1} - 2\phi_{T'_1}\phi_{T_2} &> 2\phi_{T_1}\phi_{T'_2} - 2\phi_{T_1}\phi_{T_2} \iff \\ 2\phi_{T'_2}(\phi_{T'_1} - \phi_{T_1}) - 2\phi_{T_2}(\phi_{T'_1} - \phi_{T_1}) &> 0 \end{aligned}$$

and the last conditions is true because $\phi_{T'_1} > \phi_{T_1}$ and $\phi_{T'_2} > \phi_{T_2}$.

Appendix B: Additional information for chapter 2

B.1 R&D program sample

The administrative records indicate that 532 applications were received and sent to reviewers to be evaluated. As in the case of the certification program, firms could not request reconsideration, but could submit a new project. In the main text of the article, I consider all the applications received and use cluster standard errors at the level of the firm to avoid overestimating the precision of the estimates. This is a conservative approach as it makes more difficult to find any effect and avoids arbitrarily selecting observations.

Table B.1 presents estimates of the difference in the number of previous applications between the observations on both sides of the cutoff. As shown in the table, the number of previous applications is not different between firms just to the left and right of the cutoff. This can be interpreted as evidence that reviewers did not manipulate the score to benefit or obstruct firms applying more than once.

Table B.1: Balance of number of previous applications

	bw=10	bw=20	bw=30	bw=30
	(1)	(2)	(3)	(4)
right	-0.142 (0.273)	-0.158 (0.211)	-0.135 (0.203)	-0.121 (0.336)
obs	355	512	532	532

***, **, * indicate significance at 1%, 5%, and 10%. The values reported in this table are the OLS estimates of β_1 in equation (2.1). The dependent variable is the number of previous application by the same firm. Values in parenthesis are standard errors clustered at the level of the firm. Estimates use a bandwidth of 10, 20, and 30 points around the cutoff, triangular kernel and a linear polynomial of the running variable.

Appendix C: Additional information for chapter 3

C.1 IMSS individual-level data

All private Mexican employers are legally required to report wages for their employees to the Mexican social security agency, *Instituto Mexicano del Seguro Social (IMSS)*. Not all employers comply; those that do not are commonly considered to be in the informal sector. The raw IMSS data can thus be considered a census of private, formal-sector establishments and their workforces for 1985-2005. (Most public-sector workers and employees of the state-run oil company are covered by other insurance programs.)

The IMSS data contain information on the daily wage of individuals. The wages are a measure of total compensation, called the *salario base de cotización*, which includes both earnings and benefits, including payments made in cash, bonuses, commissions, room and board, overtime payments, and in-kind benefits. The data are reported as a sequence of spells for each worker, with beginning and end dates. In principle it is possible to recover a wage for every individual for every day of every year. We extracted data for September 30 for each year. At the level of individuals, the data also contain information on age, sex, and state and year of the individual's first registration with IMSS. At the establishment level, the data contain information only on location and industry (using the IMSS's own 4-digit industrial categories, of which there are 276.)

We impose the following criteria in cleaning the data. (1) In its internal records, IMSS classifies wage records by types referred to as *modalidades*. We use only *modalidades* corresponding to permanent workers and for which consistent, reliable wage figures are available.¹ (2) We require that an individual have a positive wage. (3) We require that municipality and industry are reported

¹In the IMSS internal classification system, we use *modalidades* 10, 13 and 17. This excludes rural casual laborers, self-employed individuals who are insured through IMSS, employees of rural agricultural cooperatives and credit unions, freelance workers, taxi drivers, domestic workers, miscellaneous public-sector workers insured through IMSS, and a number of smaller categories.

for establishment. (4) We winsorize wages within year, assigning wages above the 90th percentile to the 90th percentile and wages below the 10th percentile to the 10th percentile, for the reasons discussed in Section 3.3. (5) If wages for more than one establishment are observed simultaneously for a given individual, we keep only the highest-wage observation. (6) We require that individuals be 14 years or older and 64 years or younger. (7) We require that workers be employed in an establishment in the largest connected set of establishments, as described in Section 3.2 above.

The total number of workers with wage data in the “raw” IMSS files (i.e. the sample size after step 3 of the cleaning procedure described in the previous paragraph) ranges from approximately 4 million in 1985 to approximately 10 million in 2005. The numbers of individuals in the cleaned data, after step (6) above but before limiting to the largest connected sets, are in Appendix Table C.2.2. The numbers after limiting to the largest connected sets are in Table 3.1. Additional details on the IMSS data are available in Castellanos et al. (2004) and Kaplan et al. (2010), Kaplan et al. (2007).

C.2 EIA plant-level data

The cleaning procedure for the plant-level data from the *Encuesta Industrial Anual (EIA)* [Annual Industrial Survey] is the same as described in Appendix II of Verhoogen (2008), and rather than repeat the entire description we focus here on key points.

For the reasons discussed in Section 3.3, we focus in this paper on the EIA data from 1993-2003. The sample was drawn in 1993, to include the largest plants in 205 of the 309 6-digit industries (*clases*) in the Mexican industrial classification system, covering 85% of the value of production in each industry. These plants were followed over time, with minimal refreshing of the sample.

Capital stock was constructed using the perpetual-inventory method. Capital was classified into three types: machinery and equipment, land and buildings, and transportation equipment and other fixed assets. Following Olley and Pakes (1996), each type of capital was assumed to evolve according to $K_{jt} = (1 - \delta_j)K_{jt-1} + i_{jt-1}$, where j indexes the three types of capital. Following


Levinsohn and Petrin (2003), the depreciation rates, δ_j for machinery and equipment, land and buildings, and transportation equipment were assumed to be 10%, 5% and 20% respectively. Total capital stock is the sum of the three types of capital. The book value of capital stock in 1993 was taken as the initial value.

The following cleaning procedures were implemented. (1) Plants in multi-plant firms for which complete information was not reported separately by plant were dropped. (2) Plants owned in whole or in part by government entities were dropped. (3) Establishments that appeared to be *maquiladoras*, because they derived more than 95% of their income from exports or subcontracting, were dropped. (4) Variables that changed within a plant by more than a factor of 10 from one year to the next were set to missing. (5) Missing values of variables were imputed following the procedure described in Appendix II of Verhoogen (2008). (6) After imputation, plants with incomplete information on any key variable (employment, hours, wage bill, total costs, domestic sales, total sales, capital stock) were dropped. (7) The key variables listed in the previous point were “winsorized” at the 1st and 99th percentiles, following a suggestion Angrist and Krueger (1999).

We then selected a balanced panel of plants with complete data in all years 1993-2003, which we refer to as the EIA panel. 3,529 plants are included in this balanced panel. We then linked the EIA panel to the IMSS data and collapsed to the period level (period 1 is 1992-1994, period 2 is 1996-1998, period 3 is 2000-2002; see Section 3.3 for justification), averaging variables within period.² We then selected plants with estimated plant and average person components for all three periods. 2,625 plants satisfied this requirement. We refer to this balanced plant-period-level panel as the EIA-IMSS panel.


²For period 1, we averaged EIA variables for 1993-1994, since the EIA variables are not available in 1992.

Figure C.2.1. Comparing upward and downward moves, IMSS data, 1992-1995




Notes: Sample is all workers observed in 1992-1995 in the IMSS database (after cleaning steps 1-6 described in Appendix C.1) who changed job between 1993 and 1994 and held both the preceding and new job for at least two years. The dashed line is at -45 degrees. Each dot plots upward and downward transitions between two types of firms, classified according to quartiles of average coworkers' wage. Wage changes are changes in log real wage, averaged over workers making same transition, between 1993 and 1994.

Figure C.2.2. Comparing upward and downward moves, IMSS data, 2000-2003




Notes: Sample is all workers observed in 2000-2003 in the IMSS database (after cleaning steps 1-6 described in Appendix C.1) who changed job between 2000 and 2001 and held both the preceding and new job for at least two years. The dashed line is at -45 degrees. Each dot plots upward and downward transitions between two types of firms, classified according to quartiles of average coworkers' wage. Wage changes are changes in log real wage, averaged over workers making same transition, between 2000 and 2001.

Figure C.2.3. Movers' mean nominal wages, IMSS data, 1992-1995



Notes: Figure is similar to Figure 3.4 but shows *nominal* wage changes. Sample is all workers observed in 2000-2003 in the IMSS database (after cleaning steps 1-6 described in Appendix C.1) who changed job between 2001 and 2002 and held both the preceding and new job for at least two years. Each line corresponds to a transition between types of firms classified by quartiles of the average coworkers' wage.

Figure C.2.4. Movers' mean nominal wages, IMSS data, 2000-2003



Notes: Figure is similar to Figure 3.5 but shows *nominal* wage changes. Sample is all workers observed in 2000-2003 in the IMSS database (after cleaning steps 1-6 described in Appendix C.1) who changed job between 2001 and 2002 and held both the preceding and new job for at least two years. Each line corresponds to a transition between types of firms classified by quartiles of the average coworkers' wage.

Table C.2.1. Aggregate labor force statistics

	1990	2000
Total population	81.25	97.48
Economically active pop. age > 14	31.23	40.16
Remunerated workers	25.96	32.01
Remunerated workers, private sector	21.27	27.20
Workers registered in IMSS	10.76	15.24
Workers registered in IMSS, permanent	9.53	13.53

Notes: Numbers in millions. Figures drawn from *Anuario Estadístico de los Estados Unidos Mexicanos* [Statistical Yearbook of Mexico], 2005, which draws in turn from the following: decennial population censuses (total population), 1991 *Encuesta Nacional de Empleo* [National Employment Survey] (economically active population age > 14), and INEGI Banco de Información Económica (remunerated employees), *IMSS Memoria Estadística*.

Table C.2.2. Summary statistics, IMSS individual-level data, before limiting to largest connected sets

year	# individuals	# establishments	avg. age	fraction male	avg. daily wage (raw, 2002 pesos)		avg. daily wage (winsorized, 2002 pesos)	
					mean	std. dev.	mean	std. dev.
1988	5,257,200	426,570	31.76	0.72	146.81	702.83	115.61	62.21
1989	5,993,961	469,018	31.29	0.70	151.86	618.46	125.50	73.70
1990	6,869,806	538,274	31.02	0.69	144.82	417.67	123.23	78.16
1991	7,546,628	596,124	31.01	0.68	153.03	235.67	134.68	86.85
1992	7,756,268	621,246	31.10	0.68	161.18	264.05	142.49	94.67
1993	7,659,363	615,684	31.41	0.68	180.30	249.97	152.09	105.07
1994	7,843,005	619,991	31.58	0.67	190.37	259.41	155.34	109.18
1995	7,413,728	600,015	32.01	0.67	152.28	202.47	122.94	87.77
1996	7,998,174	617,721	31.95	0.67	139.70	187.39	111.67	80.87
1997	8,592,365	640,381	31.92	0.67	140.61	194.91	112.75	83.02
1998	9,001,372	653,151	32.03	0.67	142.99	177.21	115.34	85.46
1999	9,578,857	674,710	32.17	0.66	145.03	176.33	117.56	86.67
2000	10,203,195	711,176	32.32	0.65	153.09	181.65	125.09	91.72
2001	10,103,668	736,849	32.85	0.65	160.87	187.25	132.29	97.03
2002	10,151,601	748,620	33.20	0.65	163.62	189.03	134.94	98.13

Notes: Sample is from IMSS employer-employee records after cleaning steps 1-6 in Appendix C.1 (before restricting to largest connected sets). Winsorization is at 10th and 90th percentiles. Wages are reported both in "raw" (i.e. pre-winsorized) and winsorized form. See Section 3.3 and Appendix C.1 for further details. Average 2002 exchange rate: 9.60 pesos/US\$1.

Table C.2.3. Summary statistics, EIA panel, 1993

	non-exporters (1)	exporters (2)	all plants (3)
Total revenue	151.88 (7.96)	417.78 (47.51)	226.55 (14.65)
Employment	184.56 (5.66)	370.31 (21.78)	236.78 (7.48)
K/L	146.63 (4.82)	194.80 (8.96)	160.17 (4.30)
Export share of sales		0.15 (0.01)	0.04 (0.00)
Avg. hourly wage (EIA)	43.64 (0.80)	60.98 (1.16)	48.51 (0.67)
N	2537	992	3529

Notes: Table reports statistics using 1993 data from EIA panel (before linking to IMSS data). Standard errors of means in parentheses. Exporter defined as export sales > 0. Export share is fraction of total sales derived from exports. Sales are measured in millions of 2002 Mexican pesos, capital-labor ratio in thousands of 2002 pesos, and average daily wage in 2002 pesos. Average 2002 exchange rate: 9.60 pesos/US\$1. For further details, refer to Section 3.3 and Appendix C.2.

Table C.2.4. Number of EIA plants linked to IMSS data, by connected set status

	EIA panel plants	EIA panel plants linked to IMSS		EIA panel plants not linked to IMSS	EIA-IMSS panel plants
		Total	Not connected		
Period 1 (1992-1994)	3,529	2,769	2,746	23	2,625
Period 2 (1996-1998)	3,529	2,903	2,868	35	2,625
Period 3 (2000-2002)	3,529	2,872	2,812	60	2,625

Notes: Data from IMSS employer-employee records and EIA plant panel as described in Section 3.3. “Connected” means contained in the largest connected set, as described in Section 3.2.

Table C.2.5. OLS estimates, exports and plant-level outcomes

	(1)	(2)	(3)	(4)	(5)
	$\Delta \log$ K/L	$\Delta \log$ avg. hourly wage (EIA)	Δ avg. log daily wage (IMSS)	Δ plant component	Δ avg. person component
A. Not including initial value of capability proxy					
Δ export share	0.043 (0.082)	0.042 (0.039)	0.034 (0.023)	0.058** (0.026)	-0.024 (0.024)
B. Including initial value of capability proxy					
Δ export share	0.001 (0.082)	0.016 (0.038)	0.019 (0.023)	0.046* (0.026)	-0.028 (0.024)
initial log employ.	0.072*** (0.009)	0.045*** (0.004)	0.026*** (0.003)	0.021*** (0.004)	0.005 (0.003)
6-digit industry \times period effects	Y	Y	Y	Y	Y
region (state) \times period effects	Y	Y	Y	Y	Y
N (plants)	2625	2625	2625	2625	2625
N (obs)	5250	5250	5250	5250	5250

Notes: Table reports OLS regressions corresponding to equation (3.4) in the main text. Panel A omits the initial value of the capability proxy, $\hat{\lambda}_{jp-1}$; Panel B includes it. Export share is fraction of total sales derived from exports. Robust standard errors in parentheses. *10% level, **5% level, ***1% level.

Table C.2.6. Construction of alternative proxies, TFP and predicted export share

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	log VA ind. 31	log VA ind. 32	log VA ind. 33	log VA ind. 34	log VA ind. 35	log VA ind. 36	log VA ind. 37	log VA ind. 38	log VA ind. 39	export share
log w.c. empl.	0.176*** (0.024)	0.122*** (0.036)	0.126*** (0.038)	0.365*** (0.034)	0.311*** (0.031)	0.073** (0.028)	0.223*** (0.061)	0.156*** (0.018)	0.291** (0.135)	
log b.c. empl.	0.259*** (0.030)	0.295*** (0.047)	0.510*** (0.102)	-0.041 (0.028)	0.038 (0.036)	0.261*** (0.053)	0.141 (0.090)	0.157*** (0.032)	0.014 (0.247)	
log capital	0.341*** (0.055)	0.054 (0.050)	0.128 (0.109)	0.145** (0.059)	0.299*** (0.052)	0.514** (0.227)	0.165 (0.100)	0.128* (0.068)	0.247 (0.185)	
log empl.										0.053*** (0.003)
log sales										0.040*** (0.003)
log K/L										0.027*** (0.002)
N (plants)	611	575	123	312	743	226	81	814	40	3529
N (obs)	6519	6117	1317	3364	8003	2407	870	8769	429	38819

Notes: Columns 1-9 report coefficients from Levinsohn and Petrin (2003) TFP estimation, with log value-added as outcome, log employment (white-collar and blue-collar separately) and log capital as covariates, and log materials and log electricity as proxies, separately by 2-digit industry. The industries (indicated at top) are: food, beverages, tobacco (31); textiles, apparel, leather goods (32); wood products, including wood furniture (33); paper, papers, products, publishing (34); chemical products (35); non-metallic mineral products (36); basic metal products (37); metal products, machinery, equipment (38); other manufacturing (39). Column 10 reports the coefficients of a tobit model of the export share using log employment, log sales, and log capital-labor ratio as covariates, and including 4-digit sector fixed effects. Observations with negative or zero value-added omitted in Columns 1-9. For all columns, if capital or employment variable has value zero, log is set to zero. Standard errors in parentheses. *10% level, **5% level, ***1% level.

Table C.2.7. Correlations between proxies, EIA-IMSS panel, pooling periods

	log empl. (hours)	log dom. sales	log TFP (L-P)	pred. exp. share
log employment (hours)	1.0000			
log domestic sales	0.8142	1.0000		
log TFP (Levinsohn-Petrin)	0.5578	0.6979	1.0000	
predicted export share	0.8893	0.9306	0.6247	1.0000

Notes: Table reports bilateral correlation coefficients using the EIA-IMSS panel, pooling periods 1 (1992-1994), 2 (1996-1998), and 3 (2000-2002).