Essays on the Effects of Public Policy

by

Mayara Priscila Felix Silva

Submitted to the Department of Economics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2021

© Mayara Priscila Felix Silva, MMXXI. All rights reserved.
The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Author . . . .
Department of Economics
May 14, 2021

Certified by . . . .
Benjamin A. Olken
Jane Berkowitz Carlton and Dennis William Carlton Professor of Microeconomics
Thesis Supervisor

Certified by . . . .
David G. Atkin
Professor of Economics
Thesis Supervisor

Certified by . . . .
Arnaud Costinot
Professor of Economics
Thesis Supervisor

Accepted by . . . .
Amy N. Finkelstein
John & Jennie S. MacDonald Professor of Economics
Chairman, Department Committee on Graduate Theses

Essays on the Effects of Public Policy

by

Mayara Priscila Felix Silva

Submitted to the Department of Economics
on May 14, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This dissertation presents three papers on the effects of public policy on market outcomes. In the first paper, I analyze the effects of trade liberalization on firm labor market power in Brazil. I find that while Brazil's 1990s trade liberalization significantly lowered wages and increased labor market concentration, it did not increase firm labor market power. The negative effects of trade on local wages were therefore likely driven by reductions in the marginal revenue product of labor. In the second paper, in collaboration with M. Chatib Basri, Rema Hanna, and Benjamin A. Olken, I analyze the effects of two corporate taxation reforms in Indonesia: one in tax administration and one in tax rates. We find that the tax administration reform had large effects on tax revenue and reported income, and that the government would have had to raise the corporate income tax marginal tax rate on affected firms by 8 percentage points to match those revenue gains. Finally, the third paper evaluates the impact of a school discipline policy in Massachusetts on student suspensions and test scores at charter schools. I find that the policy reduced charter suspensions by roughly 10 percentage points, but had no impact on charter learning.

Thesis Supervisor: Benjamin A. Olken
Title: Jane Berkowitz Carlton and Dennis William Carlton Professor of Microeconomics

Thesis Supervisor: David G. Atkin
Title: Professor of Economics

Thesis Supervisor: Arnaud Costinot
Title: Professor of Economics

Acknowledgments

This dissertation would not have been written without the help of many people. Some contributed directly with obtaining data or with feedback: my advisers, other faculty members, and graduate students. Others have contributed indirectly with support before and during the PhD: my advisers, my family, and my friends. While this dissertation could not have been written without the former group, I dedicate it to the latter, and to all those who I do not know but whose lives I hope my research could one day improve.

Specifically, I would first like to thank Josh Angrist and Parag Pathak for helping me be admitted into the MIT Economics PhD program. JA and PP: I'm sure your recommendation letters were fundamental, and I hope I have made you proud. Second, I thank my dissertation advisers David Atkin, Arnaud Costinot, and Ben Olken for their unwavering support and intellectual investment in my projects. Thank you Ben, in particular, for listening to me whenever I needed to vent and for guiding me through the hardest times, especially throughout the current Covid-19 pandemic.

Third, I thank those who have helped me walk the long path from poverty in Northeast Brazil to an Economics PhD at MIT. These folks recognized my talent, were inspired by my drive, and gave me the opportunity to succeed via scholarships, provision of test prep materials, and mentoring. Thank you Edvaldo Amorim, Eduardo Carvalho, Israel Souza Jr., Márcia Mizuno, Rita Moriconi, and Serafim Gomes. Your investment in me made a difference. Thank you also uncle Sebastião and aunt Jaciara for introducing me to the English language to begin with, and for lobbying your friends so that I would get tuition discounts at their English teaching schools.

Finally, I thank my family and friends for never letting go of my hand throughout the ups and downs of graduate student life. My husband Seth was a constant source of strength and my most frequent go-to buddy for economic theory questions. My sister Maíra and brother-in-law Bruno provided me with laughter and sanity. My mother Iara reminded me of how far I've come and how proud I should be no matter what lies ahead. My childhood memories of my father Maviael, blessed be his memory, helped me place every adversity in perspective. My in-laws Jane, Gus, Ira, E.Y., and my niece Sylvia gave me reinvigorating love and care. My vast extended family was a constant inspiration. And last but not least, I was lucky to live the PhD experience with a handful of incredible peers. Thank you Anton, Beata, Ben D., Carolyn, Gabriel, Ivan, Masao, Michael W., Mikel, Patricia, Sophie, and Tamar for your presence in my life and for the many shared problem set solutions.

Thank you all for everything you have done for me. This dissertation is for you.

Contents

List of Figures

List of Tables

Chapter 1

Trade, Labor Market Concentration, and Wages: Evidence from Brazil

Abstract

Growing evidence suggests that import competition has long-lasting negative effects on local wages and employment. Why? I study a potential mechanism: trade-induced increases in firm labor market power. Using employer-employee linked data and tariff reductions from Brazil's 1990s tariff reform, I find that a 10 percent increase in import competition exposure increased local labor market concentration by 6.3 percent relative to trend. This effect is led by exporter expansion and contraction of import-competing firms. I then investigate the implications for wage markdowns, the standard measure of firm labor market power. Using a framework of labor market oligopsony to measure wage markdowns, I estimate that prior to the reform Brazilian workers took home 64 cents of every marginal dollar generated by firms, and that a 10 percent increase in import competition reduced that figure by 0.04 cents. This muted effect is driven by a surprising finding: in response to wage changes, Brazilian workers substitute across firms within markets as weakly as they substitute across markets. On net, these findings suggest that: 1) even sizable changes in labor market concentration can have limited consequences for wages; and 2) the large negative effect of trade on wages was driven instead by changes in the marginal revenue product of labor.[1]

[1]I am indebted to David Atkin, Arnaud Costinot, and Ben Olken for invaluable guidance and support. I'm also deeply thankful to Daron Acemoglu, Rodrigo Adão, Nikhil Agharwal, Josh Angrist, David Autor, Abhijit Banerjee, Jonathan Cohen, Dave Donaldson, Arin Dube, Esther Duflo, Amy Finkelstein, Raymond Fisman, Teresa Fort, Matthew Grant, Fadi Hassan, Simon Jager, Xavier Jaravel, Stephanie Kestelman, Jeremy Maskowitz, Gautam Rao, Tobias Salz, Gariam Sharma, Claudia Stenwender, Frank Schilbach, Joonas Tuhkuri, Sean Wang, Iván Werning, Michael Wong, Sammy Young, Roman Andres Zarate, Roman David Zarate, Owen Zidar, and participants of the MIT Trade Tea and Development and Labor lunches for helpful feedback and discussions.

1.1 Introduction

A growing body of evidence suggests that import competition has long-lasting negative effects on wages and employment in local labor markets heavily reliant on jobs in adversely affected industries. These patterns of regional divergence have been documented in various contexts, including Brazil (Dix-Carneiro and Kovak, 2017; Kovak, 2013), the U.S. (Autor et al., 2013), and India (Topalova, 2010). Why does this regional divergence persist?

While one set of explanations points to the role of slow capital adjustment and imperfect labor mobility in an otherwise perfectly competitive labor market (Dix-Carneiro and Kovak, 2017), a mounting body of evidence suggests that these puzzling effects might also be linked to imperfect competition for workers. The most severely affected labor markets are typically remote (Topalova, 2010), face stricter enforcement of labor regulations (Topalova, 2010; Ulyssea and Ponczek, 2018), and/or exhibit limited mobility responses to import competition (Dix-Carneiro and Kovak, 2017; Topalova, 2010). As remoteness and regulations can induce workers to prefer staying in their geography and/or occupation despite wage losses, it is natural to ask whether this source of firm labor market power could generate persistent negative effects on wages in response to trade.

This paper analyses one potential mechanism underlying the negative effect of trade on wages and employment: trade-induced increases in firm labor market power. Firm labor market power has itself received substantial attention in the recent labor literature, with papers estimating that firms charge sizable wage markdowns, reducing worker's take-home wage. Sizable markdowns have been documented in several countries, whether markdowns were measured via oligopsony frameworks (Berger et al., 2019; Lamadon et al., 2019), monopsonistic competition frameworks (Dube et al., 2020; Card et al., 2018; Tucker, 2017; Ashenfelter, 2010; Manning, 2003), or via production function approaches (Hershbein et al., 2019; Hoang, 2019; Tortarolo and Zarate, 2018). Other papers have also documented that US local labor markets are highly concentrated, that high levels of concentration correlate with negative wages (Azar et al., 2017, 2018), and that concentration increased in local labor markets more exposed to import competition (Benmelech et al., 2018). Yet whether the effects on concentration translate into meaningful effects of trade on wage markdowns and wages is an open question. This paper addresses this knowledge gap.

Specifically, I analyze the effect of trade on local labor market concentration and wage markdowns in the context of Brazil's 1990s unilateral tariff reductions. I begin by presenting a simple model and empirical evidence that unilateral import tariff reductions increase local labor market concentration — specifically, the wage bill Herfindahl of the local labor market.

Theory-wise, I show that a unilateral import tariff reduction unambiguously increases labor market concentration in a two-sector small protected economy whose absolute and comparative disadvantage lies in the protected sector. The mechanism is simple: lower tariffs shift employment to and increase wages in the comparative advantage sector, whose firms — all of which are exporters — are already larger and pay higher wages in the pre-reform equilibrium. The key assumptions underlying my sharp result are: a) absolute and comparative disadvantage in the protected sector; and b) firm differentiation in local labor markets. These are very plausible assumptions.[2] The former gives the result's sign. The latter is required such that firm size (and therefore, the concept of Herfindahl) is pinned down in equilibrium. Overall, my simple model illustrates a cross-sector reallocation mechanism for increases in concentration in response to unilateral tariff reductions.[3]

After presenting the simple model illustrating how trade affects concentration, I present empirical evidence that Brazil's 1990s tariff reform increased local labor market concentration. The first step in this process is to define what constitutes a local labor market. While data restrictions have limited prior work to defining local labor markets as either geography $\times$ sector cells (e.g. Berger et al. (2019); Zarate (2016)), or as geography $\times$ occupation cells (e.g. Azar et al. (2018)) without much empirical support for these boundaries, I use employer-employee links to inform my definition. In particular, I construct worker transition matrices across geographies, sectors, and occupations, and find that Brazilian workers' transitions follow a striking pattern of permanence within geographic regions and occupation groups, but not within sectors. Based on this evidence I define local labor markets as geography $\times$ occupation group cells.

My identification strategy leverages quasi-exogenous variation in import tariff changes from Brazil's 1990s tariff reform to construct an "import competition exposure" shift-share instrument, following a similar methodology as in (Kovak, 2013) and (Dix-Carneiro and Kovak, 2017). The key idea is that the tariff reform induced differential exposure to import

[2] The rationale behind protection policy is to shield less productive national industries from cheaper foreign goods, and firms are typically differentiated either by productivity or by the amenities they can offer to employees.

[3] Another well-known mechanism whereby trade can increase labor market concentration is illustrated by Melitz (2003)'s one-sector model of symmetric (as opposed to unilateral) trade liberalization. In that model, employment is also reallocated from import-competing firms towards exporters, but the reallocation is driven by increased foreign demand for domestic goods, a consequence of the symmetric feature of the liberalization. In contrast, in my simple model only the comparative advantage sector exports, and the reallocation towards exporters is a cross-sector response to unilateral trade liberalization, which further pushes the country towards specializing in its comparative advantage. In practice, since even protected sectors have a few exporters, my empirical estimates of the effect of trade liberalization on the size of exporters vs. import-competing firms may capture both within-sector and cross-sector reallocation mechanisms.

competition across local labor markets depending on each market's pre-reform industrial composition. The "shift" is the set of tariff reductions experienced by the industries employing in a local labor market,, whereas the "share" is the set of pre-reform wage bill shares of those industries. I then use this shift-share instrument as the tariff reform treatment intensity in a difference-in-differences empirical strategy. I present visual checks for parallel trends prior to the reform, and report treatment effects relative to trend for the outcomes where pre-trends are in fact present.

To ensure that the key identifying variation in my shift-share instrument is the "shift" rather than the "share", I follow advice from the recent literature on shift-share instruments (i.e. Adao et al. (2019); Borusyak et al. (2018)). Specifically, I create a new dataset of finer industry level tariff shocks, which increases the number of "shifts" used to construct the shift-share instrument to 285 from the 53 used by the current literature.[4] I then show that these shifts are indeed the identifying variation in my treatment effects using a placebo exercise in which I replace the real tariff shocks in the shift-share instrument with randomly drawn tariff shocks, and run 1,000 placebo regressions of local labor market outcomes on the placebo shift-share instrument. Consistent with the real tariff shocks being the key identifying variation, I show that the distribution of treatment effects from the placebo regressions is centered at zero and have the expected rejection mass (i.e. roughly 5% of t-statistics beyond +/- 1.96).

Using my shift-share instrument as the measure of change in import competition exposure, I find that a 10 percent increase in import competition exposure increased the wage bill Herfindahl of local labor markets by roughly 6.3 percent relative to a declining pre-reform trend. Consistent with the mechanism illustrated in my simple model, this effect is accompanied by a mild expansion of employment at exporting firms and a substantial contraction of employment at import-competing firms. Since rising wage bill Herfindahls may imply greater labor market power in standard oligopsony models, could these effects have substantially changed wage markdowns and thus explain the divergence in wages documented in the literature?

To understand whether increases in concentration had meaningful effects on wage markdowns, I then outline a simple labor market oligopsony framework that allows me to measure firms' wage markdowns. On my framework's supply side, workers have discrete choice over firms, choosing local labor markets first and firms last. On the demand side, firms compete

[4]Specifically, I map product-level import tariffs to the finest industry code, "CNAE95", available in my employer-employee linked dataset at the time of the reform. Previous papers used averaged tariffs provided by Kume et al. (2003), which are either at "Nível 50" (total of 20 industries) or at "Nível 80" (total of 53 industries) . See Appendix A.2 for details.

for workers à la Cournot, making hiring decisions while taking other firms' hiring decisions as given. Combined, these assumptions imply two things. First, the wage markdown of a firm z in local labor market m is $1+\varepsilon_{zm}^{-1}$, where ε_{zm}^{-1} is the firm's inverse elasticity of residual labor supply. Second, $\varepsilon_{zm}^{-1} = \frac{1}{\eta}(1 - s_{zm}) + \frac{1}{\theta}s_{zm}$ is the weighted average of two key labor supply elasticities: the cross-firm within-market inverse elasticity of substitution $\frac{1}{\eta}$, and the cross-market inverse elasticity of substitution $\frac{1}{\theta}$, with weights given by the firm z's wage bill share s_{zm} in local labor market m, as in standard models of local labor market oligopsony (e.g. Berger et al. (2019)).

The oligopsony framework provides a simple link between firm size (and thus, Herfindahls) and wage markdowns. If workers substitute more easily across firms within a market than across markets (i.e. $\frac{1}{\theta} > \frac{1}{\eta}$), then larger firms set larger wage markdowns. If, however, workers substitute just as easily across firms than across markets (i.e. $\frac{1}{\theta} = \frac{1}{\eta}$), then firm size has no consequence for wage markdowns. As derived in Berger et al. (2019), the structural relationship between a market's average Herfindahl and its average wage markdown carries a similar dependence on the elasticities of substitution. If $\frac{1}{\theta} > \frac{1}{\eta}$, then increased Herfindahls raise the average wage markdown. If, however, $\frac{1}{\theta} = \frac{1}{\eta}$, then increased Herfindahls do not affect wages. In fact, $\frac{1}{\eta} = \frac{1}{\theta}$ is the monopsonistic competition (that is, constant wage markdowns) limiting case of the oligopsony framework.

I then proceed to use rich employer-employee linked data combined with reform-induced import tariff reductions to estimate the key elasticities of labor supply comprising the wage markdown. To estimate the cross-firm within-market elasticity of substitution, I bring my framework's firm-level wage equation to the data, instrumenting its firm employment component with the change in tariffs faced by the firm (i.e. a firm-level labor demand shifter). I introduce a new method for ensuring that the tariff shocks (measured at the sector level) are firm-specific, which is to estimate the within-market cross-firm elasticity using a sample of firms that are the unique producers of their sector in the local labor market in which they operate. I document that failing to focus on unique producers would substantially bias my estimate of $\frac{1}{\eta}$ in the direction that suggests more labor market power (i.e. larger $\frac{1}{\eta}$). To estimate the cross-market elasticity of substitution, I bring my framework's market-level wage index equation to the data, instrumenting its employment index component with the change in market-level import competition exposure (i.e. a market-level demand shifter).

I estimate $\frac{1}{\eta} = 0.542$ (with a s.d. of 0.207) and $\frac{1}{\theta} = 0.627$ (with a s.d. of 0.231). In other words, firms would have to increase wages by 0.542 percent if they wished to employ one percent more workers, and wages would have to rise by 0.627 at the market-level in order to attract one percent more workers from other markets. Whether these inverse elasticities are small and similar because workers find it hard to move anywhere or because they have strong preferences not to move anywhere depends on how one interprets the idiosyncratic shocks faced by workers when choosing a workplace.

This is in sharp contrast with corresponding estimates by Berger et al. (2019) the US, based on local labor markets defined at the geography $\times$ sector level. Berger et al. (2019) estimate that $\frac{1}{\theta} = \frac{1}{0.66} = 1.51$, which is nearly eight times larger than their estimate of $\frac{1}{\eta} = \frac{1}{5.38} = 0.19$, suggesting US workers substitute a lot more swiftly across firms within markets than across markets.[5] In contrast, I cannot reject that $\frac{1}{\eta} = \frac{1}{\theta}$, and the close magnitude of the two estimates is consequential for the effect of trade on wage markdowns.

Combined with measures of firm wage bill shares, my estimates of $\frac{1}{\eta}$ and $\frac{1}{\theta}$ imply that workers took home 64% of their marginal revenue product. This estimate is in line with corresponding estimates in the literature (e.g. 65% for US manufacturing by Hershbein et al. (2019)'s, between 60 and 64% to for US tradable sector firms by Berger et al. (2019),[6] and 53% by Hoang (2019) for Chinese manufacturing).

Once I have estimated wage markdowns for all firms, I quantify the change in local labor markets' average wage markdowns induced by trade. I find that a 10% increase in market-level import competition exposure reduced workers' average take-home wage share of 64% by 0.04 percentage points. This effect is largest among exporter firms (0.05 percentage point reduction) and is five times smaller and statistically insignificant for non-tradable sector firms. Since I find that the wage bill Herfindahl (sum of squared wage bill shares) increased by 6.3% in response to a 10% increase in import competition, the small effect on firms' wage markdowns is driven primarily by the similarity between the elasticities of $\frac{1}{\eta}$ and $\frac{1}{\theta}$. That is, primarily by the fact that the market structure of Brazilian local labor markets is much closer to monopsonistic competition (i.e. constant markdowns) than to oligopsony (i.e. variable markdowns).

My findings contribute to growing literatures on labor market oligopsony and on the regional incidence of trade. While recent papers have examined the impact of import competition on local labor market concentration (Benmelech et al., 2018), the literature still

[5]These estimates also suggest that local labor markets are less substitutable to US workers than to Brazilian workers (since $1.51 > 0.627$), and that firms are more substitutable to US workers than to Brazilian workers (since $0.19 < 0.542$). In other words, US workers are more "picky" than Brazilian workers when it comes to choosing local labor markets, while Brazilian workers are more "picky" than US workers when it comes to choosing firms.

[6]Calculated as $1/\left(1 + \hat{\epsilon}^{-1}\right)$ using the authors' reported $\hat{\epsilon}^{Data} = 1.80$ (for firms with less than 5% wage bill shares), and $\hat{\epsilon}^{Data} = 1.52$ (for remaining firms with shares below 10%).

lacks empirical support for a trade-related mechanism underlying changes in concentration. I provide evidence of the presence of a simple employment reallocation mechanism, whereby employment gets mildly shifted towards exporters while import-competing firms experience significant contraction. In addition, while other papers have explored the impact of trade on both labor market concentration and on similar measures of wage markdowns (Hoang, 2019), no one paper has been able to study the impact of trade on both outcomes using employer-employee links (as opposed to firm-level data) and the universe of formal sector employment (as opposed to manufacturing only). But employer-employee links and full coverage of firm employment is essential in a study of firm labor market power because wage data can be confounded by worker characteristics and because workers' full set of employment options is not limited to manufacturing.

My employer-employee linked data and quasi-exogenous tariff shocks allow me to make several empirical contributions to the existing literature. First, I measure local labor market concentration more accurately because I observe the universe of formal employment (as opposed to manufacturing employment only). Second, I define local labor market boundaries that are consistent with workers' job-to-job transition patterns. The transition matrices I compute are the first documented job-to-job transition matrices for a developing country, adding to a growing literature on worker network mobility (e.g. Schubert et al. (2019); Nimczik (2017); Schmutte (2014)).[7] Third, my estimates of firm-specific wage markdowns are not confounded by worker characteristics because I observe the latter and can partial out their effects from wages. Finally, my large dataset allows me to introduce a new method for finding firm-specific tariff shocks, described above, strengthening the identifying variation for the cross-firm elasticity of substitution.

Combined, my data and identification strategy improvements are consequential to my finding that trade had a negligible effect on wage markdowns. My null effect differs from the only other available estimate of the effect of trade on labor market distortions, by Hoang (2019), which is based on plant-level data from Chinese manufacturing firms. Hoang (2019) finds that import tariff reductions on firms' outputs had a positive but statistically insignificant effect on labor market distortions, but import tariff reductions on firms' inputs substantially reduced labor market distortions.[8] While there are various reasons why our

[7]Import competition had a similar effects on wages, employment, and concentration at these more finely defined markets. See Appendix Table 1.1 and Appendix Figures A.10 and 4.

[8]Specifically, Hoang (2019) reports that a 10 percent reduction in a firm's input tariffs driven by China's entry into the WTO reduced wage distortions among Chinese manufacturing firms by 5.8 percent. In contrast, my estimates are insensitive to whether import tariffs on firms' inputs are incorporated into the tariff reduction shocks or not.

estimates may differ,[9] my identification strategy offers the advantage of its transparency in documenting the lack of pre-trends in the effect of import competition exposure on average wage markdowns.

Overall, this is the first paper to show that even sizable changes in local labor market concentration can have limited consequences for wage markdowns. The key reason for the limited effect is the surprising finding that the key elasticities comprising markdowns are too close in magnitude. In addition, while other papers have shown that in theory labor market concentration is relevant for wage markdowns under oligopsony market structures (e.g. Berger et al. (2019); MacKenzie (2018)), my paper is the first to adopt an oligopsony framework and still find that trade had large effects on local labor market concentration but little effect on wage markdowns. On net, my findings suggest that rather than being driven by changed wage markdowns, the long-lasting negative effects of trade on wages were instead primarily driven by changes in the marginal revenue product of labor.

The remainder of the paper proceeds as follows. Section 1.2 presents motivating theoretical and empirical evidence that import tariff reductions increase local labor market concentration. Section 1.3 outlines the labor market oligopsony framework I use to measure markdowns, which I take to the data and setting described in Section 1.4 using the empirical strategy outlined in Section 1.5. Section 1.6 presents my findings, and Section 1.7 concludes.

1.2 Motivating theory and evidence

I start by outlining a simple model that illustrates how import competition can increase local labor market concentration (specifically, the wage bill Herfindahl). I then present evidence that Brazil's 1990-1994 trade reform increased labor market concentration in local labor markets more exposed to import competition, and that this effect appears to be driven by the model's illustrated mechanism, namely that firms in the non-exporting tradable sector (smaller, lower-wage) lose labor market share relative to firms in the exporting tradable sector (larger, higher wage) in response to import tariff reductions.

[9] Specifically, the difference in estimates could be accounted for by several reasons, including different theoretical frameworks (oligopsony approach vs. production function estimation), contexts (Brazil vs. China), local labor market boundaries (geography $\times$ occupation vs. geography), sample (all firms vs. manufacturing only), and data structure (employer-employee-links vs. plant-level survey).

1.2.1 A simple model of trade and labor market concentration

My model describes what happens to labor market concentration in a simple small protected economy (that is, open to trade but with import tariffs) that implements a unilateral import tariff reduction. The key assumptions underlying the results are that: a) the country has absolute and comparative disadvantage in producing the protected good relative to the world; b) firms are differentiated (e.g. either via productivity or via worker preferences) such that firm size is pinned down in equilibrium. The former assumption is what gives the results' sign: a unilateral reduction in tariffs increases labor market concentration. The latter assumption is needed such that firm size (and thus, concentration) is well defined. In what follows I make these two key assumptions, along with others that simplify exposition of the mechanism of interest.[10]

To keep things simple, I assume labor is the only input of production, and that there is no firm entry or exit. The economy consists of two goods: b ("bananas") and c ("cell phones"), each produced by a continuum of firms at constant returns to scale but varying levels of productivity. A representative household demands these two goods with elasticity of substitution $\sigma > 1$, and supplies labor across firms with elasticity of substitution $\varepsilon > 0$. In terms of market structure, firms are price takers in product markets and wage setters in labor markets. For simplicity I assume that firms mark worker wages down by constant than an variable markdown (i.e. monopsonistic competition instead of oligopsony). Finally, the country has both absolute and comparative advantage in producing bananas, such that it exports bananas (at world prices) and imports cell phones (at world prices times an import tariff $1 + \tau$). Appendix A.3 presents the detailed model and its solution.

The result that tariff reduction unambiguously increases concentration in this economy is illustrated in Figure 1, which plots the sectoral wage bill share of cell phones on the y-axis against the sectoral wage bill share of bananas on the x-axis.[11] Prior to the reform, the economy is at point A. What is important about point A is that it is located below the 45-degree line, a result I show in Appendix A.3 and that follows from the assumptions that

[10]In particular, I assume: a) that preferences over goods are CES with $\sigma > 1$ such that in equilibrium both goods will still be consumed by households; b) that preferences over labor are CES with $\varepsilon > 0$ such that in equilibrium all firms hire (e.g. no entry or exit); c) that firms are differentiated both by productivity and by worker preferences to make the setup more general than under identical productivities, even if differentiation via worker preferences is sufficient to derive the model's prediction on concentration.

[11]The figure is drawn under the simplifying assumption that within sectors all firms have identical productivity. Firms are however still differentiated workplaces from workers' perspective, so that in equilibrium firm size is pinned down by the each firm's residual labor supply. See Appendix A.3 for details and for the solution under more general within-sector firm productivity heterogeneity.

the country has both an absolute and a comparative disadvantage in producing cell phones. Note that point A is pinned down by the intersection of two objects: a "budget" (in this case, the identity that the sectoral wage bill shares must sum to one) and a relative "price" (in this case, the ratio of sectoral wage bill shares).

Now we consider the effect of a reduction of tariffs on cell phones. Lower tariffs imply that the relative price of cell phones to bananas faced by the country's consumers fall, and as a result so do the equilibrium wages in the cell-phone sector. Lower wages in cell phones reallocates labor towards banana production, thus reducing the ratio of the cell phone sector's wage bill to the banana sector's wage bill.[12] On net, the effect of a tariff reduction is depicted by the dark black arrow in Figure 1: cheaper cell phone imports lower the relative "price" that pins down equilibrium, bringing the economy to the post-reform equilibrium point B.

The consequence of a move from point A to point B to the wage bill Herfindahl is depicted in Figure 1 as the move from the blue arc to the red arc. These arcs are Herfindahl level sets: the bigger the radius of the arc, the higher the Herfindahl.[13] Since the economy starts below the 45-degree line, and since a tariff reduction reduces the relative price of cell phones to bananas, the effect of a unilateral import tariff reduction in this economy is to move the economy to a higher Herfindahl level set; that is, to a more concentrated labor market.

Importantly, the key assumptions driving the result in my simple model are very plausible. First, the rationale behind protection policy is to shield relatively less productive (i.e. comparative disadvantaged) national industries from cheaper foreign goods. Consistent with this idea, Appendix Figure A.1 shows that prior to the reform most of Brazilian exports came from the least protected industries. Second, firms are not identical: they are typically differentiated either by productivity or by the amenities they can offer to employees.

Overall, this simple model illustrates how a unilateral import tariff reduction increases labor market concentration by triggering the cross-sector reallocation of employment towards the exporting sector. Since conditional on switching firms most Brazilian workers switch sectors,[14] cross-sector reallocation is likely an important mechanism driving effects on local

[12]Of course, consumers gain by lower prices on cell phones, consuming imported instead of domestic phones. See Appendix A.3 for the complete proof.

[13]Specifically, in Appendix A.3 I show that

$$HHI = \left(\frac{w_b l_b}{WL}\right)^2 + \left(\frac{w_c l_c}{WL}\right)^2$$

In other words, Herfindahl level sets can be visualized as circles on the $\frac{w_b l_b}{WL} \times \frac{w_c l_c}{WL}$ plane centered at $(0, 0)$ and with radius $\sqrt{HHI}$.

[14]See Appendix Table A.2. Conditional on switching employers, only 18% of workers stay within their

labor market concentration.

1.2.2 Evidence on the effect of trade on concentration

The model in Section 1.2.1 predicts that, following a unilateral import tariff reduction, labor market concentration increases as firms in the comparative disadvantage sector (namely, the non-exporting sector) shrink relative to firms in the comparative advantage one (namely, the exporting sector). I next test these hypotheses against the data by leveraging unilateral tariff reductions from Brazil's 1990s trade reform.

Specifically, I test whether in local labor markets more exposed to import competition: a) non-exporting tradables fared worse than exporting tradables; and b) the wage bill Herfindahl increased. To make progress, I begin by defining local labor markets as microregion $\times$ occupation pairs, a measure that is consistent with the recent literature on labor market monopsony and oligopsony (e.g. Azar et al. (2017, 2018); Hoang (2019); Berger et al. (2019)), and finer than the definition adopted by previous papers on the impact of trade in Brazil (i.e. Kovak (2013); Dix-Carneiro and Kovak (2017)). As I later elaborate in Section 1.4, my choice of defining local labor markets at the microregion $\times$ occupation group level is guided by evidence from workers' job-to-job transition patterns that workers' labor market decisions have a significant geographic and occupational component.[15] Conditional on switching employers, 79% of workers stay in the same geography and 53% in the same occupation group, as opposed to 33% in the same sector group (or 18% in the same sector). This definition further allows me to use import tariff reduction shocks – which are sector-specific — to study worker movement *across* occupation $\times$ geography pairs.

My identification strategy follows the shift-share instrument approach adopted by several papers in the literature on the regional impact of trade (i.e. Dix-Carneiro and Kovak (2017); Kovak (2013); Autor et al. (2013); Topalova (2007)). The key idea is that trade liberalization would have had a larger impact in local labor markets whose pre-reform employment composition was highly reliant on industries most affected by the reform, i.e. the industries in which tariff cuts were largest. With that in mind, the identification strategy consists of a difference-in-differences analysis where the shift-share instrument is used as a measure of

sub-industry (the level at which I observe tariff shocks).

[15]While some authors in the monopsony literature define local labor markets at the geography $\times$ sector level, evidence from Brazilian workers' job-to-job transitions shown in Appendix Figures A.4-A.6 and Table A.2 suggests that a geography $\times$ occupation labor market cell is more appropriate. In ongoing work, Felix and Wang (2021) also find preliminary evidence that workers are more likely to switch industries rather than occupations in response to firm-specific mass layoffs.

treatment intensity, allowing me to compare the outcomes of more vs. less affected local labor markets before vs. after the reform. The key identification assumption is that, had the reform not occurred, the outcomes of more affected labor markets would have followed similar trends as the outcomes of less affected markets.

In the context of Brazil's trade liberalization, the shift-share instrument most commonly adopted in the literature was first proposed by Kovak (2013), with its quasi-exogenous properties being extensively discussed in Dix-Carneiro and Kovak (2017). Following a similar methodology, I define the reform-induced change in Import Competition Exposure (ICE) of a local labor market as:

$$\Delta ICE_m \equiv -\sum_i s_{im} \ln\left(\frac{1+\tau_{i,1994}}{1+\tau_{i,1990}}\right) \tag{1.1}$$

where s_{im} is the wage bill share of sector i in local labor market m in 1991,[16] and $\tau_{i,t}$ is the import tariff faced by industry i's output in year t.[17] Note that the key tariff variation exploited is the long-difference in import tariffs between 1994 and 1990, which was the policy-induced change announced at the beginning of the trade reforms in 1990.

Figure 2 displays the variation in ΔICE_m across geography for two example occupations, while Appendix Table A.1 provides the mean and key percentiles of the distribution of ΔICE_m across local labor markets. The mean change in import competition exposure was 5%, ranging from a 10th percentile of no exposure change (i.e. a local market made primarily of non-tradable sector firms) to a 90th percentile of 14% increase.

An important identification concern with shift-share instruments is that the ideal causal variation to be exploited should come from the "shifts" (i.e. the change in import tariffs) as opposed to from the shares, as discussed in recent methodological work on shift-share instruments (i.e. Adao et al. (2019); Borusyak et al. (2018)). I address this concern in two manners. First, I create a new sub-industry level dataset on import tariffs, increasing the number of industries used to construct ΔICE_m to 285 from 53 in the currently available literature.[18]

[16] I follow Dix-Carneiro and Kovak (2017) in choosing 1991 as the base year for all analyses. I choose wage bill shares to be the weights of this shift-share instrument to keep the measure more aligned with the theory of labor market oligopsony, but all results are virtually identical if (the more commonly used) employment shares were used instead.

[17] All results are virtually identical if the effective rate of protection of the industry (that is, the import tariff on the industry output netted out by the import tariff on the industry's inputs) is used instead of output tariffs alone. For simplicity of exposition and for all main specification I refer to τ_i as the import tariff on industry i's output.

[18] Previous papers used tariffs at either Nível 50 (20 industries) or Nível 80 (53 industries) from Kume

Second, I conduct a placebo exercise to check whether the variation driving my treatment effects are driven by the reform-induced shock or by the pre-reform wage bill shares. Specifically, I run 1,000 placebo regressions, separately for log employment and log wage premia.[19] In these regressions, the dependent variable is the 1997-1991 long-difference in the local labor market outcome,[20] and the independent variable is a placebo ΔICE_m, constructed with the same industry wage bill shares but with randomly drawn shocks instead of the real tariff shocks. Appendix Figure A.2 displays histograms of t-statistics associated with each round of placebo regressions. The figure shows that the roughly 5% of placebo regressions fall beyond the 1.96 standard deviations to the left and right of the distributions, as should be expected if the key source of variation driving the estimates is the tariff change.

Having defined my shift-share instrument of local labor market exposure to import competition, I proceed to estimate the effect of import competition on local labor market outcomes via a difference-in-differences strategy. Specifically, I estimate the cumulative effect (as of year k) of import competition on a local labor market's outcome Y_m as β_k from the following regression:

$$\Delta Y_{mt} = \alpha + \sum_{k \neq 1991} \beta_k \left(\Delta ICE_m \times 1_{t=k} \right) + \delta_m + \delta_t + \epsilon_{mt} \tag{1.2}$$

where ΔY_{mt} denotes the long difference in Y_m from year t back to the base year 1991,[21] and δ_m and δ_t are local labor market and year fixed effects. I estimate this regression using years 1986 to 2000, and let standard errors be two-way clustered by microregion and occupation group. Note that because ΔICE_m was defined with a negative sign, a positive β_k indicates that the import tariff reductions had a positive effect on the outcome (i.e. raised wages or expanded employment).

My findings on wage premia and employment are in line with previously documented patterns by Dix-Carneiro and Kovak (2017): import tariff reductions reduced both wages and employment in local labor markets more exposed to import competition. However, at the finer local labor market level it is apparent that both wages and employment, and especially

et al. (2003). See Appendix A.2 for details.

[19]Wage premia are wages conditional on worker observables. See Appendix A.2 for details on wage premia estimation.

[20]For simplicity I focus on the 1997-1991 long-difference as 1997 is the mid-point of the post-reform years.

[21]I follow the same long-differences convention adopted by Dix-Carneiro and Kovak (2017): long differences are taken using 1991 as the base year, and to keep the timing convention (i.e. future minus past) consistent, for the pre-treatment years ΔY_{mt} is the long difference from 1991 back to year t.

employment, were on a strong growth trend prior to the reform years.[22] Therefore, I make the identification assumption that, absent the reform, outcomes in more affected markets would have continued to follow the same pre-reform growth trend relative to least affected markets. I do this by estimating the effect of import competition *relative to trend* as the $\tilde{\beta}$ coefficients using the following regression:

$$\Delta\tilde{Y}_{mt} = \tilde{\alpha} + \sum_{k \neq 1991} \tilde{\beta}_k \left(\Delta ICE_m \times 1_{t=k}\right) + \tilde{\delta}_m + \tilde{\delta}_t + \tilde{\epsilon}_{mt} \tag{1.3}$$

where $\Delta\tilde{Y}_{mt}$ is the predicted outcome from the following de-trending regression, which I estimate using the pre-treatment years 1986-1990 only:

$$\Delta Y_{mt} = \omega + \theta \left(\Delta ICE_m \times t\right) + \kappa_m + \kappa_t + \nu_{mt} \tag{1.4}$$

in which κ_m and κ_t are local labor market and year fixed effects.[23] Standard errors are once again two-way clustered by microregion and occupation group.

The key motivating empirical evidence of the effect of trade on labor market concentration is presented in Figures 3 and 4. These figures present the effect of import competition on the employment of non-exporting vs. exporting tradables and on local labor markets' wage bill Herfindahl. I start by discussing Figure 3, which plots β_k coefficients from equation 1.2. Consistent with the model, Figure 3 shows that as of 1997 (the mid-point of the post-treatment period), the total employment of non-exporting tradable sector firms fell by 75 percent for every 10 percent increase in import competition exposure, whereas the employment of exporters grew by roughly 10 percent.[24]

Figure 4 shows the effects on concentration both on levels (β_k coefficients from equation 1.2, in Panel B) and relative to trend ($\tilde{\beta}_k$ coefficients from equation 1.3, in Panel A). Although noisy on levels, the effect of import competition on concentration is positive and large relative to trend, of an average of roughly 6.3 percent increase for every 10 percent increase in

[22] The pre-trend in wages is also present at the microregion-level only, as documented by Dix-Carneiro and Kovak (2017), but the strong pre-trend in employment is specific to the finer local labor markets defined by occupation $\times$ geography.

[23] Specifically, $\Delta\tilde{Y}_{mt} = \Delta Y_{mt} - \left[\hat{\omega} + \hat{\theta}\left(\Delta ICE_m \times t\right) + \hat{\kappa}_m + \sum_{i=1986}^{1990} \hat{\kappa}_i\right]$, where the hats indicate estimates of the respective parameters from equation 1.4.

[24] In addition, it's important to note that all effects in Figure 3 are for *total* employment by firm type. Evidence from Menezes-Filho and Muendler (2011) suggests that displaced workers are not themselves absorbed by exporters, suggesting that the increase in exporter employment comes either from outside the formal sector or from non-tradable sector firms. Finally, I find a negative but not statistically significant effect on the log employment of non-tradable sector firms.

import competition exposure. Table 1.1 summarizes these effects. As the median labor market concentration in the base year was 0.20,[25] this is equivalent to reducing the median "effective" number of equally-sized employers from $5 = 1/0.20$ to $4.7 = 1/0.23$. Overall, import competition increased labor market concentration, and reduced the number of firms, employment, and wages in more affected markets.

Do the theoretical and empirical findings linking import competition to labor market concentration have implications for the effect of trade on wages? In the simple model outlined in Section 1.2.1, the answer is no. Since firms are assumed to compete monopsonistically for labor, their wage markdowns are constant over time, and therefore cannot be affected by changes in firms' relative size. In the simple model, real wages are driven instead by changes in the price index alone, which falls, thus raising real wages.

But what if firm size did affect wage markdowns? In particular, what if firms set wages as in oligopsony rather than monopolistic competition, such that changes in their relative size were directly translated to changes in wage markdowns? I next investigate these questions empirically with the help of an oligopsony framework that allows me to measure wage markdowns, and hence their change in response to trade.

1.3 Measuring wage markdowns

In this section I outline the simple oligopsony framework I use for measuring wage markdowns. Under this framework wage markdowns can be measured as the weighted average of workers' elasticities of substitution across firms vs. across markets, where the weights are a function of firm wage bill shares. Sections 1.4 and 1.5 then outline how I estimate these key elasticities leveraging detailed data and Brazil's rich institutional setting.

1.3.1 Labor supply

I follow a similar setup as in Berger et al. (2019)'s micro-foundation of a nested CES labor supply system, which I extend to incorporate worker taste shifters for specific markets and for firm-market pairs. Incorporating these taste-shifters allows me to bring the model's wage equation directly to the data to estimate the elasticities of interest, bypassing the need for estimating elasticities via indirect inference as in Berger et al. (2019).

Discrete choice. The economy consists of a continuum of homogenous workers j, a

[25](See Appendix Table A.1 for descriptive statistics of local labor markets)

large but finite number of local labor markets m, and a finite number of firms z within each local labor market. Each worker chooses to which firm-market pair zm she provides h^j_{zm} units of labor by minimizing her indirect disutility of work V_{zm} subject to the constraint of making reservation earnings $y^j \sim F(y)$:

$$\min_{zm} V_{zm} = \left\{ \ln h^j_{zm} + \ln \xi_{zm} + \ln \xi_m - \xi^j_{zm} \right\}$$
$$\text{s.t. } h^j_{zm} w_{zm} \geq y^j$$

where ξ_{zm} and ξ_m are firm-market- and market-specific taste shifters common to all workers, and ξ^j_{zm} is an idiosyncratic worker taste for working at firm z in local market m, and w_{zm} is the wage paid by firm z in local labor market m to identical workers.[26]

The probability that worker j chooses firm z in local labor market m is given by $P^j_{zm} = P\left(\xi^j_{zm} > \ln w_{zm} - \ln y^j - \ln \xi_{zm} - \ln \xi_m\right)$. Assuming ξ^j_{zm} has a generalized extreme value form with the nesting structure shown in Appendix Figure A.3, by the results in McFadden (1978) and McFadden (1981) this probability can be written as

$$P_{zm} = \underbrace{\frac{\left(\frac{w_{zm}}{\xi_{zm}}\right)^{1+\eta}}{\sum_{k\in\Theta_m}\left(\frac{w_{zm}}{\xi_{zm}}\right)^{1+\eta}}}_{P_{z|m}} \times \underbrace{\frac{\xi_m^{-1}\left[\sum_{k\in\Theta_m}\left(\frac{w_{km}}{\xi_{km}}\right)^{1+\eta}\right]^{\frac{1+\theta}{1+\eta}}}{\sum_l \xi_l^{-1}\left[\sum_{k\in\Theta_l}\left(\frac{w_{kl}}{\xi_{kl}}\right)^{1+\eta}\right]^{\frac{1+\theta}{1+\eta}}}}_{P_m} \tag{1.5}$$

where Θ_m is the set of firms operating in local labor market m, P_m is the marginal probability that workers choose local labor market m, and $P_{z|m}$ is the probability that workers would choose firm z conditional on them choosing local labor market m.

Aggregation. Total labor supplied to firm z in market m can be found by integrating probabilities P^j_{zm} (times $h^j_{zm} = y^j/w_{zm}$ supplied by each worker) over the continuum of workers:

$$l_{zm} = \int_0^1 P^j_{zm}\left(\frac{y^j}{w_{zm}}\right) dF(y) = w_{zm}^{-1} P_{z|m} P_m \int_0^1 y^j dF(y) \tag{1.6}$$

where $\int_0^1 y^j dF(y) \equiv Y$ is national labor income. Lastly, to obtain an expression for l_{zm} that is a function of w_{zm}, parameters, and market-level aggregates, I define the following wage

[26]Empirically, firm z's wage premium in local labor market m. See Appendix A.2 for details on estimation of wage premia for firms and local labor markets.

indexes labor supply indices:

$$W_m \equiv \left[\sum_{k\in\Theta_m}\left(\frac{w_{km}}{\xi_{km}}\right)^{1+\eta}\right]^{\frac{1}{1+\eta}} \qquad W \equiv \left[\sum_m\left(\frac{W_m}{\xi_m}\right)^{1+\theta}\right]^{\frac{1}{1+\theta}}$$

$$L_m \equiv \left[\sum_{k\in\Theta_m}\left(\xi_{zm}l_{km}\right)^{\frac{1+\eta}{\eta}}\right]^{\frac{\eta}{1+\eta}} \qquad L \equiv \left[\sum_m\left(\xi_m L_m\right)^{\frac{1+\theta}{\theta}}\right]^{\frac{\theta}{1+\theta}}$$

which imply $Y = WL$ and

$$P^j_{zm} = P_{z|m}P_m = \frac{w_{zm}}{\xi_m\xi_{zm}}\left(\frac{w_{zm}}{W_m}\right)^{\eta} \times \left(\frac{W_m}{W}\right)^{\theta} W \tag{1.7}$$

Plugging this expression back into 1.6 gives the firm-specific labor supply

$$l_{zm} = \xi_m^{-1}\xi_{zm}^{-1}\left(\frac{w_{zm}}{W_m}\right)^{\eta} \times \left(\frac{W_m}{W}\right)^{\theta} L \tag{1.8}$$

Interpretation. The parameters η and θ have an important economic interpretation. Atkeson and Burstein (2008) show that the nested discrete choice setup can be mapped into a representative worker problem where the representative worker has nested CES preferences over firms and markets. As a result, through the lens of this model η corresponds to workers' wage elasticity of substitution across firms within markets, whereas θ is workers' wage elasticity of substitution across markets.

1.3.2 Imperfect Competition in Labor Markets

On the labor demand side, firms compete a la Cournot, choosing their labor demand in each market to maximize their profits while taking as given the labor demand of other firms. Firm profits are given by

$$\Pi_z = R_z\left(\varphi_z, l_{zm}, l_{-zm}\right) - \sum_m w_{zm}\left(l_{zm}, l_{-zm}\right) l_{zm}$$

where R_z is the firm's revenue function, and $w_{zm}\left(l_{zm}, l_{-zm}\right)$ is the inverse of the firm's labor supply described in equation 1.8, i.e. the wage that firm z would need to pay to obtain l_{zm} units of labor in local labor market m, conditional on other firms obtaining l_{-zm}. The first-order condition with respect to l_{zm} implies the equality of marginal revenue and marginal cost:

$$w_{zm} = \frac{\overbrace{\partial R_z / \partial l_{zm}}^{\text{MRPL}}}{\underbrace{1 + \varepsilon_{zm}^{-1}}_{\text{Wage markdown}}} \tag{1.9}$$

where $\partial R_z / \partial l_{zm}$ is the marginal revenue product of labor (MRPL), and $\varepsilon_{zm}^{-1} \equiv \frac{\partial \ln w_{zm}}{\partial \ln l_{zm}}$ is the elasticity of firm z's inverse labor supply curve in market m.

1.3.3 Wage markdown

According to equation 1.9, the wage markdown is a function of the firm-specific inverse elasticity of residual labor supply $\varepsilon_{zm}^{-1} \equiv \frac{\partial \ln w_{zm}}{\partial \ln l_{zm}}$. I now derive an explicit expression for ε_{zm}^{-1} as a function of wage bill shares and model parameters only. I start by inverting equation 1.8 (omitting the time subscript for simplicity), which gives:

$$w_{zm} = \xi_{zm} \xi_m \left(\frac{l_{zm}}{W_{mt}^{-1}} \right)^{\frac{1}{\eta}} \left(\frac{W_m^{-1}}{W^{-1}} \right)^{\frac{1}{\theta}} \tag{1.10}$$

where $W_m^{-1} = L_m, L^{-1} = W$. I then take logs of equation 1.10 and differentiate it with respect to l_{zm} holding l_{-zm} constant. This gives:

$$\varepsilon_{zm}^{-1} = \frac{1}{\eta} (1 - s_{zm}) + \frac{1}{\theta} s_{zm} \tag{1.11}$$

where

$$s_{zm} \equiv \frac{w_{zm} l_{zm}}{\sum_k (w_{km} l_{km})} = \frac{\partial \ln L_m}{\partial \ln l_{zm}} \tag{1.12}$$

is firm z's wage bill share in market m.[27]

Equation 1.11 is the framework's key expression linking wage bill shares to wage markdowns. Specifically, note that equation 1.11 implies that the markdown is increasing in the firm's market share when $\theta > \eta$,[28] which is the typical assumption in oligopsony models.

[27] To see why equation 1.12 holds, compute $\partial \ln L_m / \partial \ln l_{zm}$ to obtain $\partial \ln L_m / \partial \ln l_{zm} = (\xi_{km} l_{km})^{\frac{1+\eta}{\eta}} / \sum_{j=1}^{N_m} (\xi_{jm} l_{jm})^{\frac{1+\eta}{\eta}}$. Now depart from $s_{zm} \equiv w_{zm} l_{zm} / \sum_k (w_{km} l_{km})$ and plug in equation 1.10 to similarly obtain $s_{zm} = (\xi_{km} l_{km})^{\frac{1+\eta}{\eta}} / \sum_{j=1}^{N_m} (\xi_{jm} l_{jm})^{\frac{1+\eta}{\eta}}$. Therefore, $s_{zm} = \partial \ln L_m / \partial \ln l_{zm}$.

[28] Since $\frac{\partial \left(1 + \varepsilon_{zm}^{-1}\right)}{\partial s_{zm}} = \left(\frac{1}{\eta} - \frac{1}{\theta} \right) \varepsilon_{zm}^{-2}$.

Limiting cases are perfect competition ($\eta, \theta \to 0$) (implying zero markdowns) and monopsonistic competition ($\eta = \theta$) (implying constant markdowns).

Finally, note that because inverse elasticity of firm-specific residual labor supply ε_{zm}^{-1} ranges from zero to infinity, the larger the markdown, the smaller the worker's take-home wage. Therefore, for ease of interpretation I define workers' wage take-home share as $\mu_{zm} \equiv [1+\varepsilon_{zm}^{-1}]^{-1}$. Substituting ε_{zm}^{-1}:

$$\mu_{zm} = \left[1 + \frac{1}{\eta}\left(1 - s_{zm}\right) + \frac{1}{\theta} s_{zm}\right]^{-1} \tag{1.13}$$

The wage take-home share is easier to interpret because $\mu_{zm} \in [0,1]$ is decreasing in the inverse elasticity of residual labor supply ε_{zm}^{-1}. When $\varepsilon_{zm}^{-1} = 0$, $\mu_{zm} = 1$ and the worker receives its full marginal revenue product. When $\varepsilon_{zm}^{-1} \to \infty$, $\mu_{zm} \to 0$ and the worker receives zero wages. I will therefore report empirical results in terms of both wage markdowns and wage take-home shares.

1.4 Data and Setting

I use three main data sources for workers, tariffs, and exporting activity. Appendix A.2 describes these datasets in detail.

1.4.1 Data

First, rich labor market data come from Brazil's administrative employer-employee linked database Relações Anuais de Informações Sociais (RAIS), spanning years 1986-2000. RAIS covers the universe of Brazilian formal sector workers. I focus on the sample of private sector workers aged 18 to 65, or roughly 15 million private sector workers per year.

While my results are internally valid for the formal sector, a key concern on external validity is that employment might be underreported at RAIS, particularly by small firms, as documented in Ulyssea and Ponczek (2018). To gage the importance of this concern, I check whether local labor market concentration (measured using RAIS) correlates with microregion-level rates of informality (measured using Census data). Appendix Figure A.16 presents these correlations. While concentration and informality appear to be positively correlated in each census year, the correlation is weak, with high levels of informality observed at various levels of regional concentration. This suggests my findings are unlikely to be driven by under-reporting of employment at RAIS.

Second, data on tariffs come from UNCTAD TRAINS, which I map to RAIS via the 5-digit economic activity code CNAE95.[29] Finally, exporting activity is mapped to RAIS using firms' unique identifier CNPJ. What I observe in terms of exporting activity is the list of exporting firms for years 1990-1994, which were provided by the (extinct as of 2019) Ministry of Development, Industry, and Foreign Trade, currently a part of the Ministry of the Economy.

1.4.2 Setting: local labor market boundaries

All components of the wage markdown depend on how local labor markets are defined. I identify local labor market boundaries by adopting the notion that boundaries should be consistent with workers' job-to-job transitions. Empirically, I explore this by leveraging the employer-employee links in RAIS to construct job transition matrices that are conditional on workers switching employers. I produce these matrices along three potentially key dimensions of workers' job search process: geography, occupation, and sector. Figures A.4 and A.5 plot the 1990-1991 transitions are plotted for the top 50 "microregions" and top 50 occupation groups.[30]

Brazilian workers' transitions have clear geographic and occupational components: conditional on switching employers, most workers stay within either the same microregion or the same occupation group. Appendix Table A.2 summarizes this information for all—not only top 50—microregions, occupations, and sectors. It shows that 79% of workers stay in the same microregion conditional on switching employers, and 53% of workers stay in the same occupation group. In contrast, only 32% of workers stay within the same broad industry group, and 18% stay within the same industry. Overall, based on these patterns I follow others in the recent literature on labor market power (e.g. Azar et al. (2018)) and define local labor markets at the microregion $\times$ occupation group cell. Appendix Table A.1 provides descriptive statistics of the local labor markets.

[29] See Appendix A.2 for details on mapping procedures.

[30] Microregions are areas of contiguous municipalities that are economically integrated as defined by Brazil's statistical agency IBGE. I use the 2019 microregion borders, for a total of 558 microregions. See Appendix A.2 for details on mapping microregion borders to RAIS via municipality codes. Occupation groups are 68 groups defined based on 2-digit CBO94 codes, which are grouped by Brazil's statistical agency IBGE according to a combination of worker education and task content. See Appendix A.2 for details.

1.4.3 Setting: Brazil's 1990s import tariff reform

The key policy-induced variation I leverage throughout my analyses stem from Brazil's 1990s unilateral import tariff reductions. Dix-Carneiro and Kovak (2017) provide an in-depth discussion of Brazil's 1990s import tariff reform. Tariffs were reduced from a pre-reform average of 33% to a post-reform average of 13%.[31] with some sectors experiencing larger reductions than others because they were previously more protected, as shown in Appendix Figure A.11.

Kovak (2013) argues that the striking correlation between pre-reform tariff levels and reform-induced tariff cuts, as documented in Appendix Figure A.11, is precisely the biggest support for exogeneity of the tariff cuts. In particular, because the pre-reform levels of protection were set decades earlier (Kume et al., 2003), it is unlikely that the 1990s tariff cuts were correlated with counterfactual industry performance at the time. Instead, the reductions were motivated by the broader national goal to reduce all tariffs towards a much lower and much more equalized level of protection across all industries.

Given their plausible exogeneity, I exploit the cross-sector heterogeneity in tariff reductions to estimate the effect of import competition on firm-level outcomes (e.g. firm employment, wage premia, $\frac{1}{\eta}$, etc.), and the cross-market heterogeneity in tariff reduction exposure (i.e. ΔICE_m from equation 1.1) to estimate the effect of import competition on market-level outcomes (e.g. market concentration, $\frac{1}{\theta}$, etc.).

Finally, despite the plausible exogeneity in tariff cuts, one might be concerned that the decades-long level of protection enjoyed by the industries experiencing the largest tariff cuts might induce differential trends in industry outcomes. To address this concern, I estimate year-specific regression coefficients for all outcomes of interest to check for parallel trends. Finally, for the market-level regressions where pre-trends are indeed present, I estimate effects relative to trend based on equations 1.4 and 1.3.[32]

1.5 Empirical Strategy

I next describe how I bring the framework equations to the data to estimate the key components of wage markdowns (namely, the elasticities of substitution $\frac{1}{\eta}$ and $\frac{1}{\theta}$), and how I estimate the effect of import competition on markdowns and corresponding wage take-home

[31] Simple 1990 averages of nominal tariffs at CNAE95 level. See Appendix A.2 for details.

[32] As shown in Table 1.1, the presence of pre-trends tends to attenuate the treatment effect of trade towards zero, suggesting that more protected industries were thriving, not shrinking, prior to the reform.

shares.

1.5.1 Estimation of framework parameters

Within-market cross-firm inverse elasticity of substitution $\frac{1}{\eta}$

I derive the regression equation for estimating $\frac{1}{\eta}$ by taking logs of a time-specific version of equation 1.10, which gives:

$$\ln w_{zmt} = \underbrace{-\frac{1}{\theta}\ln W_t^{-1}}_{\text{Constant}} + \frac{1}{\eta}\ln l_{zmt} + \underbrace{\left(\frac{1}{\theta}-\frac{1}{\eta}\right)\ln W_{mt}^{-1} + \ln \xi_{mt}}_{\text{Market x Year FE}} + \ln \xi_{zmt} \tag{1.14}$$

The empirical equation for estimation $\frac{1}{\eta}$ then becomes:

$$\ln w_{zmt} = \alpha + \frac{1}{\eta}\ln l_{zmt} + \delta_{mt} + \ln \xi_{zmt} \tag{1.15}$$

An unbiased estimate of $\frac{1}{\eta}$ requires $E\left[\ln \xi_{zmt} | \ln l_{zmt}\right] = 0$. The key potential violation to this assumption is simultaneity of supply and demand, which makes demand an omitted variable in equation 1.15.

I address this omitted variable problem by leveraging Brazil's policy-induced import tariff reductions as labor demand shifters. The effect of import tariffs on $\ln x_{zmt}$ (log wages or employment) can estimated as:

$$\ln x_{zmt} = \beta \ln\left(1+\tau_{zt}\right) + \delta_{mt} + \epsilon_{zmt} \tag{1.16}$$

Since the plausibly exogenous tariff reduction in Brazil's reform was the 1990-1994 long-change in tariffs rather than the year-to-year variation in tariff changes, I estimate equation 1.15 in long-differences as follows:

$$\Delta \ln \tilde{w}_{zm} = \frac{1}{\eta}\Delta \ln l_{zm} + \Delta\delta_m + \Delta \ln \xi_{zm} \tag{1.17}$$

$$\Delta \ln l_{zm} = \beta \Delta \ln\left(1+\tau_z\right) + \Delta\delta_m + \Delta \ln \epsilon_{zm} \tag{1.18}$$

where $\tilde{w}_{zm}$ is the wage premium of firm z in local labor market m,[33] Δx is the long difference

[33] Wage premia are wages conditional on worker observables. See Appendix A.2 for details in wage premia estimation. Using wage premia instead of wages ensures that the variation in wages used is not due to heterogeneity in worker characteristics across firms, a feature outside of my labor supply framework.

between each respective variable from the post-reform mid-year of 1997 back to the base year 1991,[34]and

$$\Delta \ln(1+\tau_z) \equiv -\ln\left(\frac{1+\tau_{i(z),1994}}{1+\tau_{i(z),1990}}\right)$$

is the firm-level change in import competition exposure, which is based on the industry i to which firm z belongs. Therefore, the identifying variation in equation 1.18 comes from firms of different industries operating in the same local labor market (i.e. hiring in the same microregion $\times$ occupation group pair). Finally, to focus on the tariff change variation facing the firms where most workers are located, I weight all regressions by the firm's base year employment.

A potential violation to this strategy is the correlation of shocks across firms within the same local labor market. While this does not introduce an identification concern through the lens of my nested CES model,[35] as the market-level fixed effects absorb any confounding effects due to other firms in the same market also being shocked, in more general models (i.e. where there are additional nests such that firms within the same industries are closer substitutes), this introduces the possibility that $\frac{1}{\eta}$ reflects industry-level responses to tariff changes rather than firm-level responses, which would trace the industry-level labor supply curve. Since the industry-level labor supply curve is likely upward-sloping, failing to ensure that there is plenty of variation of tariff shocks across firms within a market would bias my estimates of $\frac{1}{\eta}$ upwards, making workers seem more inelastic in substituting across firms than they really are.

I address this issue two ways. First, to increase the variation in tariff changes within a market, I create a new dataset of industry-level tariff shocks at the CNAE95 level, which increases the number of sector-specific tariffs to 285 from 53 in the currently available literature.[36] Appendix Figure 5 shows that there is substantial variation in tariff changes across sectors (Panel A), and that plenty of variation remains once market-level fixed effects are removed (Panel B). Second, to guarantee that the industry-level shocks I use induce firm-level responses, I estimate equations 1.17 and 1.18 in a sample of firms that are the

Appendix Table A.5 presents robustness checks to using wage levels rather than premia.

[34]The only exception is the long-difference in tariffs, which is the reform-induced difference $\Delta \ln(1+\tau_z) = -\left[\ln(1+\tau_{z1994}) - \ln(1+\tau_{z1990})\right]$. Appendix Table A.5 further presents robustness checks to using effective rates of protection (rather than import tariff levels) as the tariff shocks.

[35]That is, through the lens of the model so long as there are at least two firms from different sectors operating in a local labor market, $\frac{1}{\eta}$ would be identified.

[36]That is, tariffs at Nível 80 from Kume et al. (2003). See Appendix A.2 for details.

unique producers of their industry in the local labor market in which they operate (plus all non-tradable sector firms), such that their shocks are firm-specific. Finally, as treatment is assigned at the industry level, I cluster standard errors at the CNAE95 industry level (i.e. 285 tradable industries, 329 non-tradable industries).[37]

Cross-market inverse elasticity of substitution $\frac{1}{\theta}$

I follow Costinot et al. (2016) in my method to derive the regression equation for the upper nest elasticity of substitution $\frac{1}{\theta}$. The key idea is that, once regression equation 1.17 is estimated, η and the regression residual $\Delta\xi_{zro}$ provide the inputs needed to estimate the model's remaining upper-nest elasticity. To see that, note that long-differencing equation 1.14 gives:

$$\Delta \ln w_{zm} = \frac{1}{\eta}\Delta \ln l_{zm} + \underbrace{\left(\frac{1}{\theta} - \frac{1}{\eta}\right)\Delta \ln W_m^{-1} - \frac{1}{\theta}\Delta \ln W^{-1} + \Delta \ln \xi_m}_{\Delta\delta_m} + \Delta \ln \xi_{zm} \tag{1.19}$$

From which it follows that estimates of the fixed effects $\Delta\hat{\delta}_m$ in equation 1.5.1 can serve as left-hand-side variables for the estimation of $\frac{1}{\theta}$ as follows:

$$\Delta\hat{\delta}_m = \left(\frac{1}{\theta} - \frac{1}{\eta}\right)\Delta \ln W_m^{-1} + \underbrace{\frac{1}{\theta}\Delta \ln \left(\frac{1}{W^{-1}}\right)}_{\text{Constant}} + \Delta \ln \xi_m \tag{1.20}$$

where — given an estimate of η and corresponding regression residuals,[38] the inverse wage index $\Delta \ln W_m^{-1} = \Delta \ln L_m$ can be computed as

$$\Delta \ln W_m^{-1} = \frac{\eta}{\eta+1}\left[\ln\left(\sum_{z\in\Theta_m} \xi_{zm,1997} l_{zm,1997}^{\frac{\eta+1}{\eta}}\right) - \ln\left(\sum_{z\in\Theta_m} \xi_{zm,1991} l_{zm,1991}^{\frac{\eta+1}{\eta}}\right)\right] \tag{1.21}$$

To keep the sample consistent with the estimation of $\frac{1}{\eta}$, I compute the employment index 1.21 using the same sample as the one used to estimate $\frac{1}{\eta}$. To focus on the import tariff reduction exposure variation facing the local labor markets where most workers are located,

[37]Appendix Table A.3 presents estimates of $\frac{1}{\eta}$ using all firms instead, while Appendix Table A.4 presents robustness to alternative clustering schemes. As I discuss in Section 1.6, the choice of sample matters for the estimate of $\frac{1}{\eta}$, while clustering schemes does not.

[38]Note that the regression residuals needed to compute $\Delta \ln W_m^{-1}$ are year-specific, one for the base year 1991 and one for the outcome year 1997. I recover the these year-specific regression residuals from the levels-equivalent regression of the long-differenced regression 1.5.1.

I weight all regressions by the market's base year employment. Finally, I cluster standard errors at the broader microregion (rather than local labor market) level.[39]

1.5.2 Effect of trade on wage markdowns

With measures of wage bill shares for each firm and estimates of $\frac{1}{\eta}$, $\frac{1}{\theta}$ at hand, I can proceed to quantify the effect of trade on markets' average wage markdown and on the (more easily interpretable) market average take take-home share. I do this by first computing the average wage markdown and average take-home share of each market and then estimating the effect of trade on them using the difference-in-differences regression equations 1.2 and 1.3.[40]

1.6 Findings

1.6.1 Elasticities of substitution

Table 1.2 presents my estimate of $\frac{1}{\eta}$ based on equations 1.17-1.18. The first stage in Panel A shows that a 1 percent increase in firm-level import competition exposure reduced employment by 0.611 percent. Panel B shows that the effect on firms' wage premia is roughly half of the employment effect, at a 0.331 percent reduction. Combined, these effects imply an inverse elasticity of substitution of 0.542.

This means that if a firm wished to increase its size by 1 percent, it would have to increase the wage premium offered to workers in its respective local labor market by 0.542 percent.

[39] Appendix Table A.6 presents robustness checks to using all firms and to alternative clustering schemes, while Appendix Table A.7 presents robustness to alternative clustering schemes.

[40] An alternative approach would be to proceed via the mathematical relationship between the average wage markdown and the wage bill Herfindahl. For the simple case where the number of firms N_m in a market is constant, the relationship between the effect of trade on the average markdown and the effect of trade on the market Herfindahl is:

$$\frac{\partial\left[\frac{1}{N_m}\sum_{z=1}^{N_m}\varepsilon_{zm}^{-1}\right]}{\partial ICE_m} = \frac{1}{N_m}\left(\frac{1}{\theta}-\frac{1}{\eta}\right)\frac{1}{2}\frac{\partial HHI_m}{\partial ICE_m}$$

which is driven by the difference in magnitude between $\frac{1}{\theta}$ and $\frac{1}{\eta}$. But the expression becomes more complicated if N_m is not constant—which is empirically the case as shown in Table 1.1. Therefore, for simplicity and for the added bonus of pre-trends check, I proceed empirically via difference-in-differences rather than analytically via the formula-driven approach. Either way, since I find very similar magnitudes for $\frac{1}{\theta}$ and $\frac{1}{\eta}$ (and in fact cannot reject the null that are the same), the formula-driven approach also suggests that the change in import competition exposure did not meaningfully change wage markdowns, despite having a large effect on the wage bill Herfindahl.

This is a large estimate, about three times larger than Berger et al. (2019)'s corresponding estimate for the US of 0.18,[41] suggesting that Brazilian workers substitute a lot less swiftly across firms in response to wage changes than US workers do. In addition, a $\frac{1}{\eta}$ estimate of 0.542 places a lower bound of 1.542 on markdowns, or an upper bound of $1/1.542 \approx 0.65$ on the wage take-home shares. In other words, the slow change in firm choice in response to wage changes imply that Brazilian workers are paid at most 65% of their marginal revenue product.

Next, Table 1.3 presents my estimate of $\frac{1}{\theta}$ based on equation 1.20. The first stage in Panel A shows that a 1 percent increase in firm-level import competition exposure reduced the local labor market labor supply index by 0.582 percent. Panel B shows that the same effect on the market-level wage change $\Delta\hat{\delta}_m$ is roughly ten times smaller, at a statistically insignificant 0.049 percent reduction. These effects imply a statistically insignificant difference between $\frac{1}{\theta}$ and $\frac{1}{\eta}$ of 0.084, which mean that $\frac{1}{\theta}$ is estimated at 0.627. This estimate is more than twice smaller than Berger et al. (2019)'s corresponding estimate of 1.51 for the US, suggesting that Brazilian workers switch *more* swiftly across local labor markets than US workers do.[42] A $\frac{1}{\theta}$ estimate of 0.627 places an upper bound on markdowns of 1.627, or a lower bound of $1/1.627 \approx 0.61$ on the wage take-home share.

1.6.2 Pre-trends and robustness

Within-market cross-firm elasticity of substitution $\frac{1}{\eta}$

A key concern with the long-differenced estimate of $\frac{1}{\eta}$ is that it might reflect pre-reform trends, since sectors experiencing larger tariff reductions were also previously more protected. To check whether that is the case, Appendix Figure A.13 plots the year-by-year estimates underlying the estimation of $\frac{1}{\eta}$. Each connected point on the graph is the effect of the firm-level import competition exposure on the long-difference of the outcome between each respective year and 1991, such that the point estimate for year 1997 corresponds to the first stage and reduced form regressions displayed in Table 1.2.[43] While mild pre-trends are present at the firm-level, they are small relative to the post-reform effect of tariff changes,

[41]Berger et al. (2019) reports an η of 5.38, whose inverse is 0.18. Their estimate is estimated via indirect inference using local labor markets defined at commuting zone $\times$ sector level and variation in corporate income tax rates.

[42]Berger et al. (2019) reports an θ of 0.66, whose inverse is 1.51.

[43]I estimate these effects separately for each year (rather than as a single stacked regression, as in the case for equations 1.2 and 1.3, which are estimated on a balanced sample of local labor markets) to avoid conditioning pre-reform estimates on firm survival.

and in the opposite direction. This is reassuring evidence that the identifying variation underlying $\frac{1}{\eta}$ is policy-induced.

Another important set of checks on the estimate of $\frac{1}{\eta}$ concerns its estimation sample. As discussed in Section 1.5.1, I estimate $\frac{1}{\eta}$ in a sample where industry-level tariff shocks are firm-specific because all tradable sector firms included in the regression are the unique producers of their sector within their local labor markets. This restriction ensures that the variation identifying $\frac{1}{\eta}$ is a firm-level labor demand shock in the local labor market, rather than an industry-level labor demand shock, which would potentially bias the estimates of $\frac{1}{\eta}$ upwards by tracing the industry-level labor supply curve (a concept outside of the labor supply framework I use, but a real force potentially biasing the estimates nonetheless).

I therefore check whether the estimate of $\frac{1}{\eta}$ is sensitive to this sample restriction. The answer is a resounding yes: failure to focus on the unique producers in a labor market would more than double my estimates of $\frac{1}{\eta}$, placing the upper bound on worker's take-home wage at 46% (rather than 65%) of their marginal revenue product. Restricting to unique producers is therefore my preferred specification.

A final set of robustness checks on $\frac{1}{\eta}$ concern the level of clustering and the measurements used for both firm wage and for the sector tariff shock. Appendix Tables A.4 and A.5 present estimates of $\frac{1}{\eta}$ based on alternative measures of clustering and alternative measures of wages and tariff shocks, respectively. Clustering at alternative levels (e.g. local labor market or microregion) reduces standard errors—indicating error terms of different sectors in the same market are negatively correlated—while using alternative measures of either wages or tariff shocks provide slightly larger but very similar estimates of $\frac{1}{\eta}$. Interestingly, the robustness of $\frac{1}{\eta}$ to the wage measure used suggests that, ex-post, the availability of employer-employee links for estimating wage premia (as opposed to firm-level average wages) was not crucial for identifying $\frac{1}{\eta}$. This is of course unknowable ex-ante.

Cross-market elasticity of substitution $\frac{1}{\theta}$

Of special concern with my estimate of $\frac{1}{\theta}$ is that local labor markets more exposed to import competition did follow different wage and especially employment trends prior to the reform, as shown in Appendix Figures A.9 and A.10.[44] While I am not able to directly

[44] Note that the outcomes actually used in the estimates of $\frac{1}{\theta}$ are not the local labor market wage premium and log employment displayed in Figures A.9 and A.10, but rather the local labor market fixed effect $\Delta\hat{\delta}_m$ from the estimation of $\frac{1}{\eta}$ and the change local labor market employment index $\Delta \ln L$ from equation 1.21. However, the patterns in Figures A.9 and A.10 are indicative of the trends $\Delta\hat{\delta}_m$ and $\Delta \ln L$ might follow as these are functions of the underlying wage premia and employment, respectively.

test whether the estimates are affected by these pre-trends, the patterns in Figures A.9 and A.10 strongly suggest that the policy induced a trend break in the growth rates of wages and employment, further suggesting that the post-reform variation used to estimate $\frac{1}{\theta}$ is also policy-induced.

Another set of robustness checks concern the set of firms used in the estimation of $\frac{1}{\theta}$ and the appropriate level of clustering. In terms of the set of firms used, to be consistent with the lower nest elasticity estimation $\frac{1}{\theta}$ is also estimated using data only from unique producers (plus non-tradables). A natural question is therefore how sensitive this upper nest elasticity is to using data from all firms to compute the right-hand-side variable, the local labor market's change in employment index.[45] Appendix Table A.6 presents the results. Using all firms, instead of the unique producers only, has nearly no impact on the estimate of $\frac{1}{\theta}$, but does reduced the first stage F by more than half, suggesting the main specification using a consistent set of firms is more appropriate.

Finally, Appendix Table A.7 presents robustness checks to clustering. The main specification uses two-way clustering by microregion and occupation group. Columns (2) and (3) present alternative clustering schemes, by microregion only and by occupation group only, respectively. Clustering by occupation group only does not meaningfully change either standard errors or the first stage F-statistic, but clustering by microregion only substantially shrinks standard errors, more than doubling the first-stage F-statistic. These results suggest that the standard errors of $\frac{1}{\theta}$ are quite sensitive to robustness choice, and that the choice of two-way clustering is quite conservative.

1.6.3 Measured markdowns and implied take-home shares

Combining measures of firm wage bill shares with the elasticity estimates for $\frac{1}{\eta}$ and $\frac{1}{\theta}$ from Section 1.6.1 place the range of markdowns between 1.542 and 1.627, or wage take-home shares between $0.65 = 1/1.542$ (for any firm with $s_{zm} = 1$) and $0.61 = 1/1.627$ (for any firm with $s_{zm} = 0$). In other words, workers take home on average between 61% and 65% of their marginal revenue product of labor.

Note that this tight range of estimated wage take-home shares is entirely driven by the similarity between the within-market cross-firm elasticity of substitution $\frac{1}{\eta}$ and the cross-market elasticity of substitution $\frac{1}{\theta}$. Whether these inverse elasticities are small and similar because workers find it hard to move anywhere (across firms or markets) or because they have

[45]In this robustness check, the change in the wage premium index is not changed, as it is the local labor market fixed effect from the the estimation of $\frac{1}{\eta}$.

strong preferences not to move anywhere depends on how one interprets the idiosyncratic shocks faced by workers in their discrete choice problem of choosing a workplace (i.e. ξ^j_{zm} in the discrete choice model of Section 1.3.1).

My estimates of wage take-home shares are in line with corresponding estimates in the literature (e.g. 65% for US manufacturing by Hershbein et al. (2019)'s, between 60 and 64% to for US tradable sector firms by Berger et al. (2019),[46] and 53% by Hoang (2019) for Chinese manufacturing). Figure 6 displays the distribution of markdowns for 1991 vs 1997 given firms' actual wage bill shares. Two patterns are clear: most firms have markdowns closer to the infinitesimal firm size benchmark, around 0.64, and the cross-section of markdowns did not change much between 1991 and 1997.

1.6.4 Effect of trade on wage markdowns

With estimates of $\frac{1}{\eta}$ and $\frac{1}{\theta}$ at hand, I now proceed to estimate how market-level average markdowns and take-home shares respond to import competition. Figure 7 plots coefficients of the effect of trade on local labor market's mean wage take-home share, estimated based on equation 1.2. It shows no pre-trends in wage markdowns leading up to the reform in 1991, but after the reform the take-home shares in local labor markets more exposed to import competition decrease steadily.

Table 1.4 summarizes these effects and includes the corresponding regressions on wage markdowns. It shows that a 10 percent increase in import competition exposure increased wage markdowns by 0.07 (0.03). This is nearly *sixty* times smaller than the only other comparable empirical estimate of the effect of trade on labor market distortions, by Hoang (2019), who estimates that a 10 percent increase in trade exposure driven by China's entry into the WTO increased wage distortions among Chinese manufacturing firms by 5.8 percent.[47]

The effect on the average take-home share is of similar magnitude: a 10 percent increase in import competition reduced the average wage take-home share by 0.04 (0.02). Interestingly, Table 1.4 shows that the effect of import tariff reductions was not uniform across firms, being larger at exporters (i.e. reduction of 0.05 (0.02)) than non-exporting tradables (i.e. reduction of 0.02 (0.01)), with a statistically insignificant effect on the markdowns for non-tradable

[46] Calculated as $1/\left(1+\hat{\epsilon}^{-1}\right)$ using the authors' reported $\hat{\epsilon}^{Data} = 1.80$ (for firms with less than 5% wage bill shares), and $\hat{\epsilon}^{Data} = 1.52$ (for remaining firms with shares below 10%).

[47] This difference in estimates could be accounted for by several reasons, including different methodologies (oligopsony approach vs. production function estimation), contexts (Brazil vs. China), sample (all firms vs. manufacturing only), and data (employer-employee-links vs plant-level survey).

sector firms.[48]

Overall, these very small effects suggest that — while trade had a sizable effect on local labor market concentration following Brazil's import tariff reductions — it had no meaningful effect on firm local labor market power. At its core, the reason for this small effect is that the two key elasticities that comprise the wage markdown (i.e. $\frac{1}{\eta}$ and $\frac{1}{\theta}$) are very similar in magnitude, and in fact not even statistically distinguishable. This suggests that essentially workers are choosing over a large set of jobs, looking beyond their microregion × occupation cells, such that firm labor market power is pinned down by how easily (or how hard) workers switch across any jobs, and is unaffected by how large the firm is in the worker's local labor market. From a theoretical perspective, what is interesting about the proximity of these elasticities is that, while the general framework allowed firms to be oligopsonists, the empirical estimates suggest the elasticities of substitution are a lot more consistent with a market structure of monopsonistic competition, where wage markdowns are constant.

1.6.5 Implications for market average wage

What is the implication of trade impacting wage take-home shares to our interpretation of reduced-form estimates of the effect of trade on average wages? The answer is simple because from equation 1.9 it follows log wages are linear in log take-home shares. Specifically:

$$\ln w_{zmt} = \ln (R_{zt}) + \ln \mu_{zmt}$$

I can therefore compute $\ln (R_{zt})$, the marginal revenue product of labor, for each firm and year as the difference between log wages and log take-home shares, and estimate the effect of trade on the $\ln (R_{zt})$. This would give us an estimate of the *counterfactual* effect of trade on wages but for its effect on firm labor market power.

Figure 8 displays the results: the counterfactual effect of trade on wages would have been nearly identical to the actual effect of trade on wages but for its effect on firm labor market power. Algebraically this implies that the effect of trade on wages is entirely driven by changes in the marginal revenue product of labor (i.e. productivity, prices, or product market markups).

[48]Note that I am able to estimate effects on the average wage markdowns of non-tradable sector firms because the shock is measured at the local labor market level. The effects on the markdowns of non-tradable sector firms are therefore entirely driven by spillovers from tradable sector firms operating in the same local labor market.

1.7 Conclusion

This paper studied a potential mechanism via which import competition could have long-lasting negative effects on wages and employment: trade-induced increases in firm labor market power. I presented theoretical and empirical motivating evidence suggesting trade increases local labor market concentration, a key diagnostic statistic of firm labor market power in models of labor market oligopsony. Theory-wise, I show that a unilateral import tariff reduction unambiguously increases labor market concentration in a small protected economy if the economy's absolute and comparative disadvantage lies in the protected sector. Empirically, I show that trade increased local labor market concentration relative to trend in local labor markets more exposed to import competition following Brazil's 1990s unilateral tariff reductions.

To understand whether changes in concentration had any meaningful effect on wage markdowns — the standard measure of firm labor market power — I then built a framework of labor market oligopsony whereby workers have discrete choice over firms and firms compete à la Cournot for workers. According to the framework, a firm's wage markdown is a weighted average of two key elasticities of labor supply: the within-market cross-firm elasticity of substitution, and the cross-market elasticity of substitution, where weights are a function of the firm's wage bill share in the local labor market, as in standard models of labor market oligopsony (e.g. Berger et al. (2019)).

I then proceeded to leverage employer-employee linked data and tariff reductions shocks from Brazil's 1990s tariff reform to estimate the key elasticities of substitution. The elasticities of substitution I estimated imply that Brazilian workers received roughly 64% of their marginal revenue product prior to the reform. Importantly, the estimated elasticities are of similar magnitude, and I cannot reject that they are the same. Consequently, regressing the average markdown at a market on the market import competition exposure measure I found that trade increased markdowns by very little, leaving the 64% figure virtually unchanged.

Overall, this is the first paper to show empirically that, even in a labor market oligopsony framework, changes in local labor market concentration need not imply changes in wage markdowns if the key elasticities comprising markdowns are too close in magnitude. On net, the small effects of trade on wage markdowns suggest that the long-lasting negative effects of trade on wages are instead entirely driven by changes in the marginal revenue product of labor (e.g. reduced markups, product demand, or productivity) rather than by increased firm labor market power.

Figure 1: Effect of a unilateral tariff reduction in a small protected economy

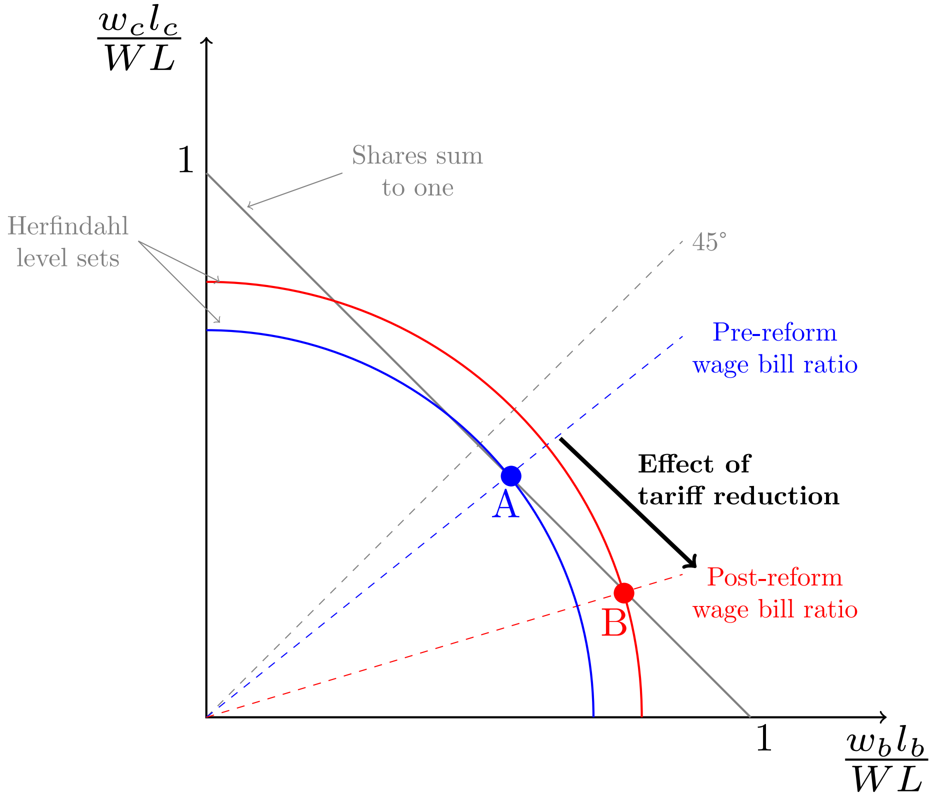

Notes: See Section 1.2.

Figure 2: Variation in Import Competition Exposure across local labor markets

Panel A: Office administration workers

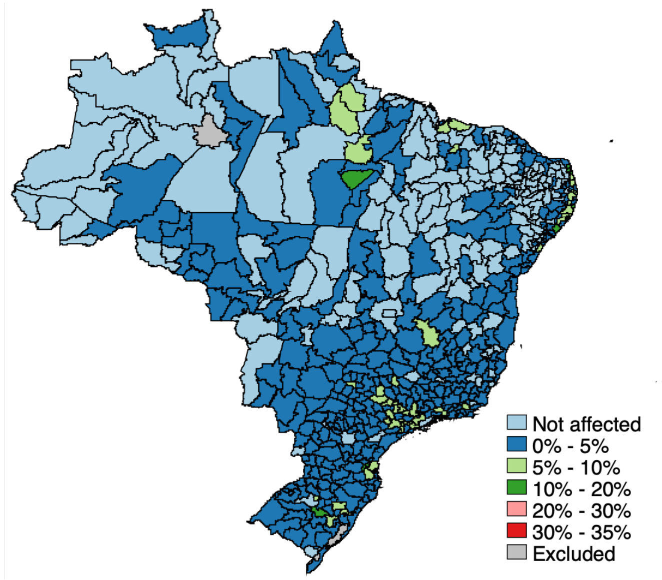

Panel B: Managers and supervisors of industrial workers

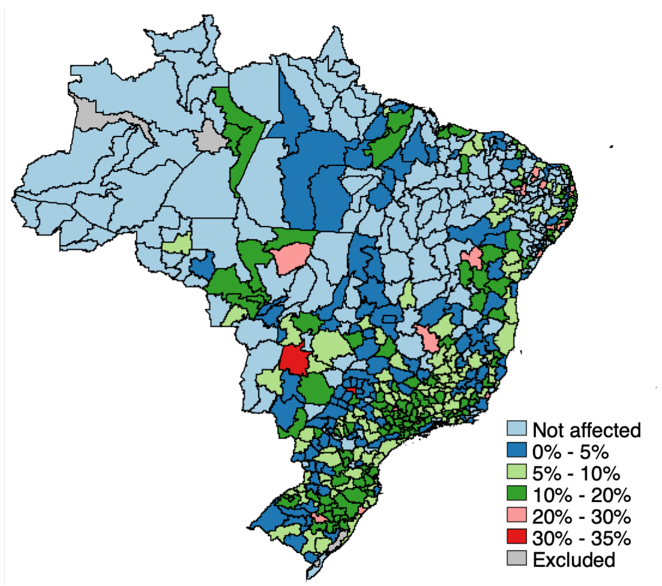

Note: This figure displays variation in ΔICE_m: the change in import competition exposure across local labor markets for two examples of occupation groups. Borders are drawn for Brazil's microregion using IBGE's 2019 border definitions, plotted using Stata's maptile program written by Michael Stepner with borders drawn by Stephanie Kestelman. The microregion of Manaus is excluded from analyses as it is a free trade zone since the 1950s. Microregions with less than one worker are also excluded.

Figure 3: Effect of import competition on exporter vs. non-exporter size

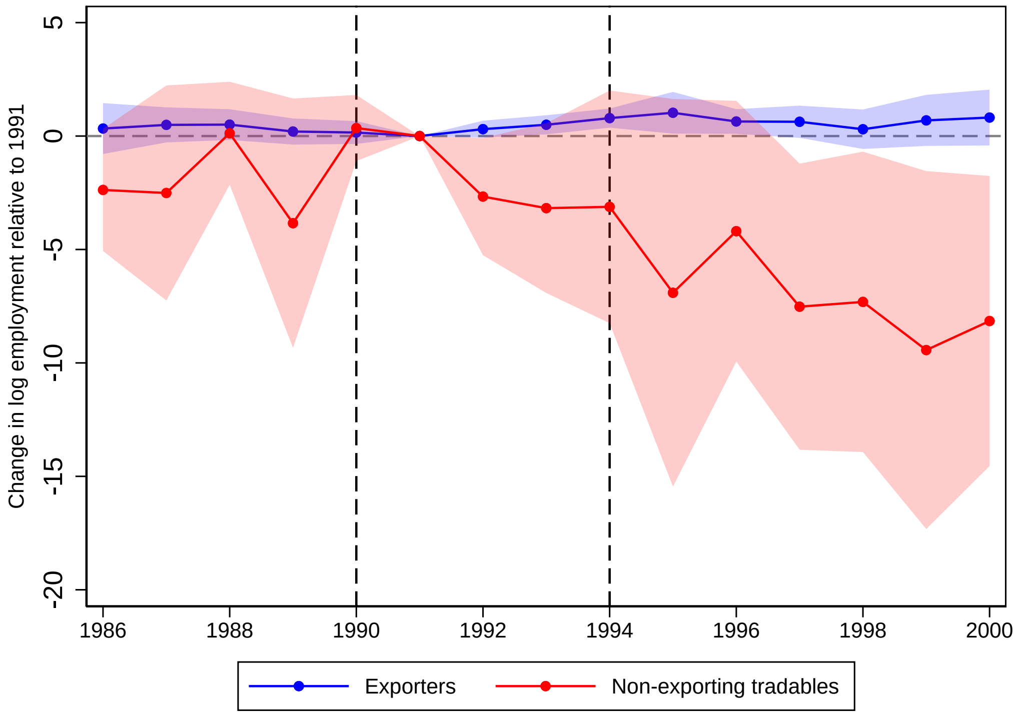

Note: This figure plots coefficients of two regressions about the cumulative effect of the change in import competition exposure: on changes in log employment of exporters; and on changes in log employment of non-exporting tradables. Each point is a β_k coefficient from equation 1.2. Dotted lines indicate the beginning and end of the tariff reductions reform. So that all differences reflect a change from a future year to a past year, for the pre-reform years the outcome is the 1991 log employment minus each respective year's log employment, whereas for the post-reform years the outcome is each respective year's log employment minus the 1991 log employment. Log employment is measured using inverse hyperbolic sine to account for markets where exporter employment is zero. All regressions are weighted by 1991 employment. Standard errors are two-way clustered by microregion and occupation group.

Figure 4: Effect of import competition on local labor market wage bill Herfindahl

Panel A: Relative to pre-trend

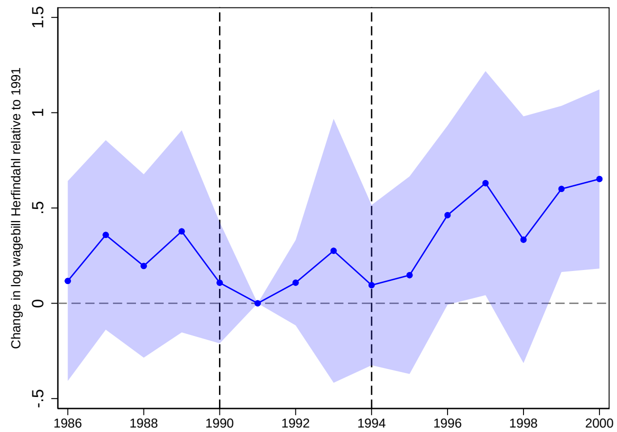

Panel B: Relative to base year

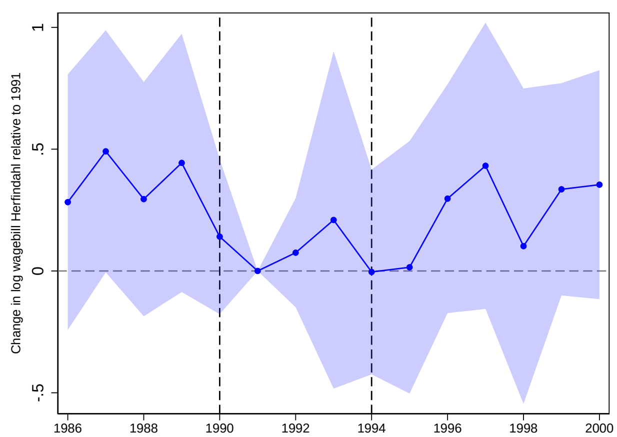

Note: This figure plots the cumulative effect of the change in import-competition exposure on the change in a local labor market's log wage bill Herfindahl, where the wage is based on firm's wage premium. Each point in Panel A is a $\tilde{\beta}_k$ coefficient from equation 1.3. Each point in Panel B is a β_k coefficient from equation 1.2. Dotted lines indicate the beginning and end of the tariff reductions reform. The 'wage bill Herfindahl of a local labor market is computed by first squaring the wage bill shares of each firm in the local labor market, then summing them. Log Herfindahl is calculated using the inverse hyperbolic sine function because its values are between 0 and 1. So that all differences reflect a change from a future year to a past year, for the pre-reform years the outcome is the 1991 log wage bill Herfindahl minus each respective year's log wage bill Herfindahl, whereas for the post-reform years the outcome is each respective year's log wage bill Herfindahl minus the 1991 log employment. All regressions are weighted by 1991 employment. Standard errors are two-way clustered by microregion and occupation group.

Figure 5: Variation in import tariff reductions across firms

Panel A: Cross-sector tariff change variation

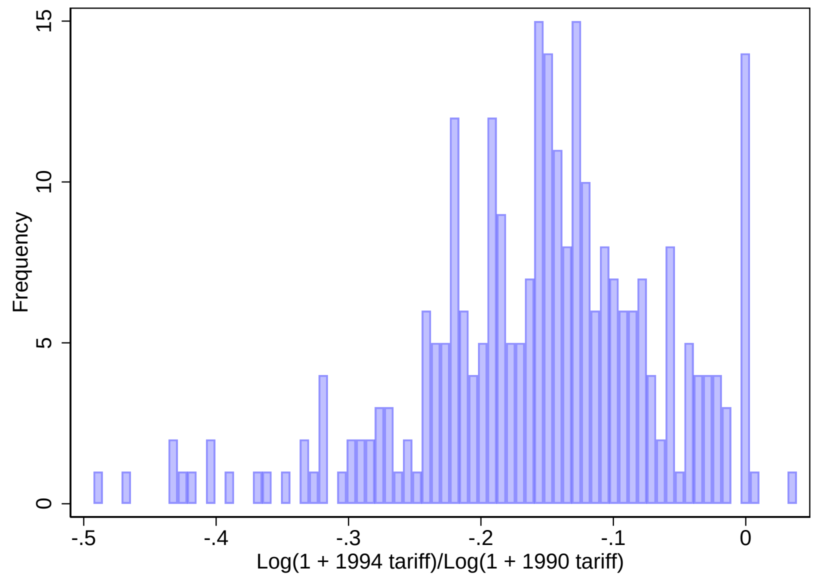

Panel B: Residual cross-firm tariff change variation

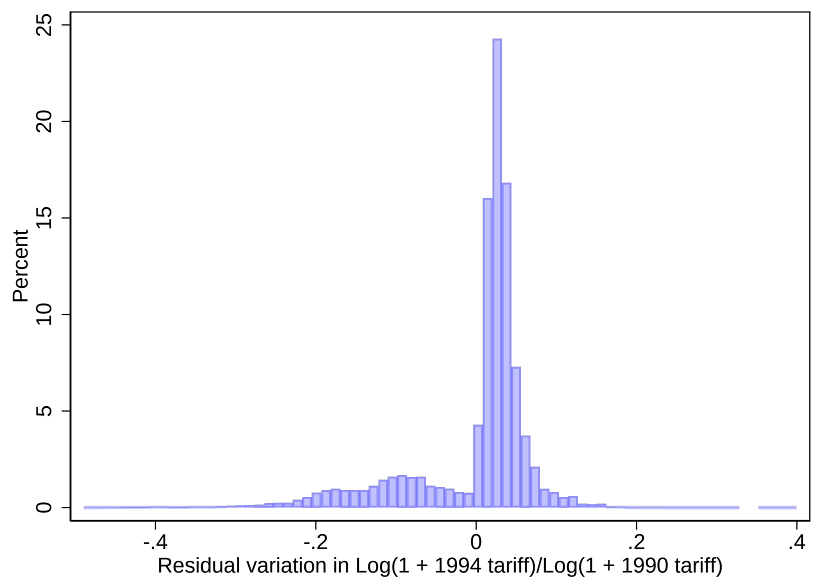

Note: This figure shows the variation in tariff changes at the CNAE95 level (285 tradable sector sectors) induced by Brazil's 1990s import import tariff reform. Panel A displays the raw data, while Panel B displays the residualized changes from a regression of tariff changes for all firms (included non-tradables, for whom the tariff change is zero) on market fixed effects.

Figure 6: Wage take-home shares at baseline (1991) and after (1997) import tariff reform

Note: This figure plots the 1991 and 1997 distributions of wage markdowns across local labor markets. See Section 1.3.3 for the definition of wage markdowns.

Figure 7: Effect of trade on local labor market mean wage take-home share

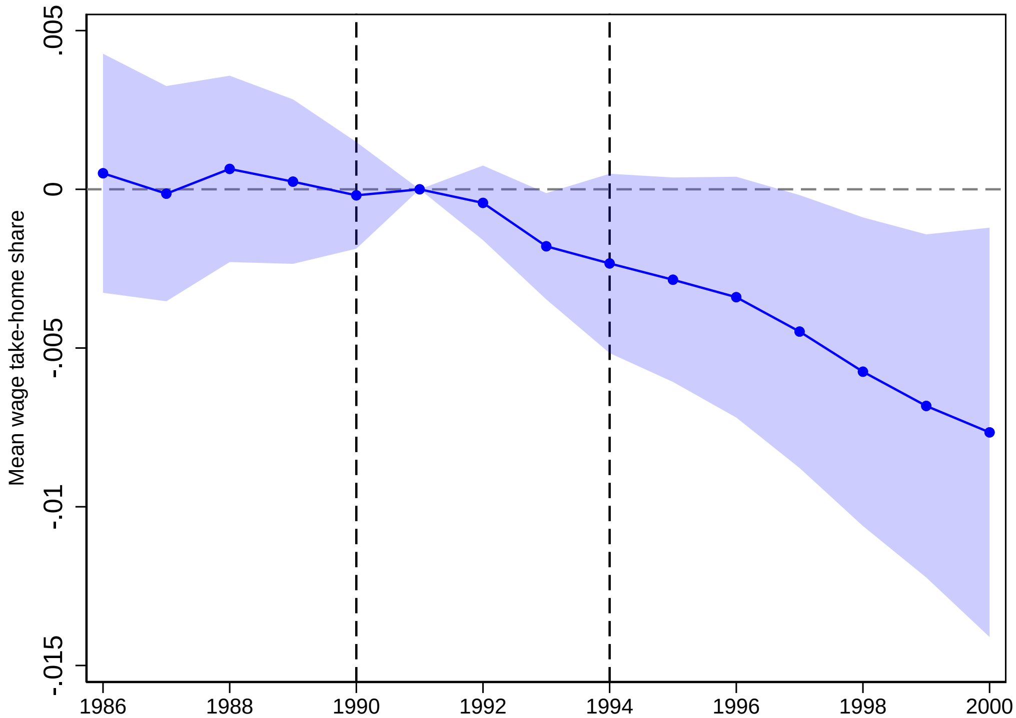

Note: This figure plots the cumulative effect of the change in import-competition exposure on the change in a local labor market's mean wage take-home share, defined in equation 1.13. Each point is a β_k coefficient from equation 1.2. Dotted lines indicate the beginning and end of the tariff reductions reform. So that all differences reflect a change from a future year to a past year, for the pre-reform years the outcome is the 1991 mean wage take-home share minus each respective year's mean wage take-home share, whereas for the post-reform years the outcome is each respective year's mean wage take-home share minus the 1991 mean wage take-home share. All regressions are weighted by 1991 employment. Standard errors are two-way clustered by microregion and occupation group.

Figure 8: Effect of trade on wages but for effect on wage take-home shares

Panel A: Relative to pre-trend

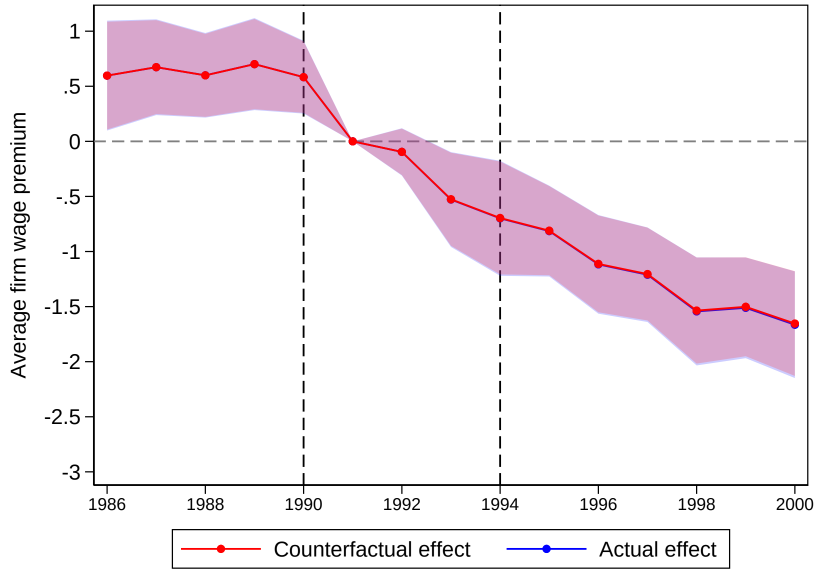

Panel B: Relative to base year

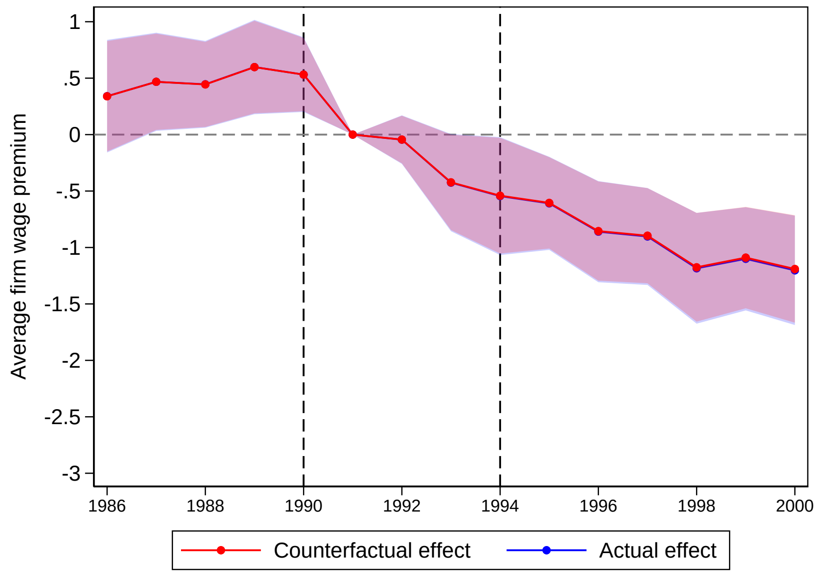

Note: See notes to Figure 4 and Section 1.6.4.

Table 1.1: Effect of import competition exposure on local labor market outcomes

	Δ Import Competition	
	Level effect	De-trended effect
	(1)	(2)
Panel A: Labor market concentration		
Δ Log Wage Bill Herfindahl (based on wage premia)	0.432	0.630
	(0.294)	(0.294)
Δ Log Wage Bill Herfindahl	0.490	0.666
	(0.292)	(0.292)
Δ Log Employment Herfindahl	0.067	0.340
	(0.193)	(0.193)
Panel B: Log number of firms and log employment		
Δ Log number of firms	-0.216	-0.834
	(0.418)	(0.418)
Δ Log total employment	-0.939	-2.177
	(0.445)	(0.445)
Δ Exporter log employment	0.629	1.016
	(0.356)	(0.356)
Δ Non-exporting tradables log employment	-7.522	-9.997
	(3.159)	(3.159)
Δ Non-tradables log employment	-0.152	-0.793
	(0.545)	(0.545)
Panel C: Log wage premium		
Δ Log wage premium	-0.660	-0.889
	(0.175)	(0.175)
Local Labor Market (LLM) FE	Yes	Yes
Year FE	Yes	Yes
Outcome is residual from regression on 1986-1990 trend	No	Yes
Observations	337,650	337,650

Note: This table presents estimates of the cumulative effect of import competition exposure on local labor market outcomes as of 1997, the mid-point of the post-reform period, estimated as β_{1997} from equation 1.2 (Column 1) and $\tilde{\beta}_k$ from equation 1.3 (Column 2). Log Herfindahl is calculated using the inverse hyperbolic sine function because its values are between 0 and 1. Log employment is measured using inverse hyperbolic sine to account for markets where exporter employment is zero. All regressions are weighted by 1991 employment. Standard errors are two-way clustered at the microregion and occupation group levels.

Table 1.2: Labor supply cross-firm within-market inverse elasticity of substitution $\frac{1}{\eta}$

	Δ in Log Import Tariff faced by firm (1)
Panel A: First stage	
Δ Firm log employment in LLM	-0.611
	(0.146)
First stage F	17.568
Panel B: Reduced form	
Δ Firm wage premium in LLM	-0.331
	(0.091)
Panel C: 2SLS	
Labor supply within-market cross-firm inverse elasticity of substitution	0.542
	(0.207)
Implied upper bound on wage take-home share	65%
Local labor market (LLM) FE	Yes
Observations	759,612
Firms	328,462
Local labor markets	14,918

Note: This table presents estimates of the coefficients from equations 1.17-1.18 and their corresponding reduced form equation. Standard errors are clustered at the sector level. An outcome's Δ is the outcome's value in 1997 (the mid-point of the post-reform period) minus its 1991 value. All regressions are weighted by 1991 employment. Standard errors are clustered at the sector level.

Table 1.3: Labor supply cross-market inverse elasticity of substitution $\frac{1}{\theta}$

	Δ Import Competition Exposure (1)
Panel A: First stage	
Δ LLM employment index	-0.582
	(0.256)
First stage F	5.161
Panel B: Reduced form	
Δ LLM wage premium index	-0.049
	(0.175)
Panel C: 2SLS	
$\frac{1}{\theta} - \frac{1}{\eta}$	0.084
	(0.308)
Panel D: Cross-market inverse elasticity of substitution	
$\frac{1}{\theta}$	0.627
	(0.228)
Implied lower bound on wage take-home share	61%
Observations (Local labor markets)	13,965

Note: This table presents estimates of the coefficients from equation 1.20 and its corresponding first stage and reduced form equations. An outcome's Δ is the outcome's value in 1997 (the mid-point of the post-reform period) minus its 1991 value. All regressions are weighted by 1991 employment. Standard errors are two-way clustered by microregion and occupation.

Table 1.4: Estimates of effect of trade on LLM average log wage markdown

	Δ Import Competition Exposure	
	Level effect	De-trended effect
	(1)	(2)
Panel A: LLM-level outcomes		
Δ Average log markdown	0.007	0.006
	(0.003)	(0.003)
Δ Average MRPL take-home share	-0.004	-0.004
	(0.002)	(0.002)
Panel B: Effect on Δ in average MRLP take-home share by LLM subgroups		
Exporters	-0.005	-0.005
	(0.002)	(0.002)
Non-exporting tradables	-0.002	-0.002
	(0.001)	(0.001)
Non-tradables	0.001	0.001
	(0.002)	(0.002)
Local Labor Market (LLM) FE	Yes	Yes
Year FE	Yes	Yes
Outcome is residual from regression on 1986-1990 trend	No	Yes
Observations (Local labor markets)	261,165	261,165

Note: This table presents estimates of the effect of local labor market import competition on the market's average log wage markdown and wage take-home shares. Outcomes are aggregated to the market level using firm total employment weights. An outcome's Δ is the outcome's value in 1997 (the mid-point of the post-reform period) minus its 1991 value. All regressions are weighted by each market's 1991 total employment. Standard errors are two-way clustered by microregion and occupation.

Chapter 2

Tax Administration vs. Tax Rates: Evidence from Corporate Taxation in Indonesia

JOINT WITH M. CHATIB BASRI, REMA HANNA, AND BENJAMIN A. OLKEN[1]

Abstract

We compare two approaches to increasing tax revenue: tax administration and tax rates. We show that when Indonesia moved top regional firms into "Medium Taxpayer Offices," with high staff-to-taxpayer ratios, tax revenue more than doubled. Examining non-linear changes to corporate income tax rates, we estimate an elasticity of taxable income of 0.59. Combining these estimates, improved tax administration is equivalent to raising rates on affected firms by 8 percentage points. Improved tax administration flattened the relationship between firm size and enforcement, reducing the "enforcement tax" on large firms. On net, improved tax administration can have significant returns for developing countries.

[1]We thank Pierre Bachas, Michael Best, Lucie Gadenne, Louis Kaplow, Henrik Kleven, Dina Pomeranz, James Poterba, Danny Yagan, and Owen Zidar for helpful comments and discussion. We thank the Indonesian Ministry of Finance and the Directorate General of Taxation, in particular Sri Mulyani Indrawati, Bambang Brodjonegoro, Robert Pakpahan, Luky Alfirman, Suahasil Nazara, Puspita Wulandari, Arfan, Peni Hirjanto, Yon Arsal, Harry Gumelar, Romadhaniah, Riana Budiyanti, Aprinto Berlianto, Nur Wahyudi, Wijayanti Kemala, Raden Yulius Wahyo Riyanto, Rudy Yuliyanto Kurniawan, Tibrizi, Harsugi, as well as many other DGT staff for their support and assistance, and thank Aaron Berman, Jonathan Cohen, Amri Ilmma, Aqila Putri, Nurul Wakhidah, and Poppy Widyasari for outstanding research assistance. Disclosures: Financial assistance from the Australian Government Department of Foreign Affairs and Trade, the J-PAL Government Partnership Initiative, and the National Science Foundation is gratefully acknowledged. Chatib Basri was Minister of Finance of Indonesia, which includes oversight of the Directorate General of Taxation, from 2013-2014, subsequent to the period of reforms that we study in this paper. The views expressed here are those of the authors and do not necessarily reflect those of the many individuals or organizations acknowledged here.

2.1 Introduction

Low tax revenue is a central challenge in many developing countries. While high-income countries typically collect around 40 percent of their GDP in tax revenue, low and middle-income countries typically collect between 10 and 20 percent. Many features of developing countries' economies, such as the informality of employment relationships, small firms, limited banking systems, and so-on, combine to limit governments' ability to tax more (Gordon and Li, 2009; Kleven et al., 2016; Jensen, 2019).

For such countries, a key question is whether they are constrained in raising revenue primarily by the elasticity of taxable income with respect to the tax rate, or whether increases in tax administration may have direct, first-order effects on the amount of tax collected. The idea is that enhanced tax administration may make evasion and avoidance more difficult, enabling governments to not only collect more, even at current tax rates; moreover, if better tax administration reduces the elasticity of taxable income, governments may be able to raise rates as well (e.g., Besley and Persson, 2014). While an emerging literature focuses on particular pieces of the tax administration puzzle in developing countries (e.g., Pomeranz, 2015; Khan et al., 2016; Naritomi, 2019; see Slemrod, 2019 for a review), there are relatively few studies that examine these types of large-scale tax administrative investments comprehensively, and can contrast these administrative reforms with more conventional attempts to raise revenue by raising tax rates (Keen and Slemrod, 2017; Slemrod, 2019).

In this paper, we study these questions in the context of corporate taxation in Indonesia. We study the introduction of a large corporate tax administration reform in Indonesia, the creation of 'Medium Size Taxpayer Offices' (henceforth, MTOs) throughout the country. These offices can be thought of as more 'intensive' tax administration, as they more than triple the staff-to-taxpayer ratio for firms. The aim is to increase both enforcement and customer service, while holding the *de jure* tax regime and the administrative structure of the tax office constant. We first study how this intensified tax administration affected actual tax filings and payments using a 9-year firm-level panel of administrative tax data. We then compare this with a differential change in the statutory marginal corporate income tax rates enacted several years later, which applied regardless of whether taxpayers were in these special tax offices.

To compare these two approaches – tax administration and tax rates – we build a model of corporate taxation in which firms can chose to keep certain parts of their business 'off the books' – i.e., hidden from the tax authority. We then adapt the framework of Keen and Slemrod (2017) to corporate taxation, showing that the two tax reforms that we consider

– improvements in tax administration and changes in tax rates – can both be analyzed by comparing their effects on net government revenue. Importantly, the fact that we consider both reforms in the same context – corporate taxpayers in Indonesia, analyzed using the same administrative tax records, and even zooming in on the effect of both types of reforms on corporate income tax payments – allows us to compare the marginal returns to both types of policies on an equal footing.

We begin by analyzing improvements in tax administration through the creation of the MTOs. In virtually all countries, corporate income tax revenues are heavily skewed, with a small number of large taxpayers comprising a considerable share of revenues. As such, many countries have created special large taxpayer offices to focus on the largest firms in the country; these are present in at least 62 countries (Lemgruber et al., 2015; Almunia and Lopez-Rodriguez, 2018). Despite being a common policy, there is relatively little evidence on whether these reforms have been effective in the developing world, and if so, on the magnitude of the gains relative to the costs of this increased supervision.[2] We study the introduction of such a reform, introduced at an unusually large scale: In the mid-2000s, Indonesia moved the largest several hundred corporate taxpayers in each of its 19 main tax regions to a special MTO in each region that focused exclusively on them. The MTOs had similar administrative structure to the regular taxpayer offices, but focused exclusively on large corporate taxpayers, and had 4-5 times as many tax staff per corporate taxpayer as regular taxpayer offices.

To identify the MTO's impact, we use a matched difference-in-differences design to examine what happens when firms are moved into the MTO. While we know which firms are in the MTO in which years, the original Excel files used to select firms were not archived, and so we cannot recreate the assignment scores and processes. Instead, we match the set of taxpayers included in the MTOs with similar taxpayers based on the taxpayers' region and based on the two variables that were primarily used for MTO assignment: the level of their pre-period tax payments and gross revenue. We use the value of these variables in 2005, the last unaffected tax year, for matching. Our preferred specification uses the entropy-balancing method of Hainmueller (2012) to create matched treatment and control samples balanced on these covariates, although other matching approaches produce similar results. We show that

[2]There is relatively little evidence on these types of reforms even for developed countries. An important exception is Almunia and Lopez-Rodriguez's (2018) study of Spain. Exploiting the fact that large firms in Spain are monitored by a national large tax office, they show that firms bunch beneath the threshold of inclusion into the LTO, and that those above the threshold report a 20 percent higher valued added tax base than those below.

the treatment and matched control group of taxpayers are on very similar trends prior to the MTOs' establishment, and then identify the impact of being moved into an MTO using a matched difference-in-differences design.

The introduction of enhanced tax administration via the MTOs dramatically increased tax revenue, at a very low cost. Real total taxes paid increased by 128 percent for affected firms; that is, moving firms to the MTOs more than *doubled* average tax collections from these firms over the subsequent six years. The government's increased costs of administering taxes through the MTOs were minuscule – about 1.5 percent of the additional revenue collected – so the net increase in government revenue is almost identical to the gross increase.[3] Put another way, the reform increased net tax revenues by Rp. 517 million per year for an annual investment of Rp. 8 million, or a 64-1 net return. All types of taxes paid by these firms rose dramatically: corporate income tax (CIT) payments rose by 111 percent, VAT payments rose by 137 percent, and other tax payments (primarily withholding taxes remitted by firms on behalf of employees) rose by 113 percent.[4] Examining MTO and non-MTO firms separately, the results suggest that this effect appears to be driven by dramatic increases in revenue from MTO firms, rather than declines for non-MTO firms, whose tax revenues remain on a similar trajectory to what they were in the pre-period.

The estimated net revenue increase from enhanced tax administration, which covered just 4 percent of all firms, amounts to a lower-bound total effect of IDR 40 trillion (USD 4.0 billion at the 2007 exchange rate).[5] Importantly, the MTO effects grow over time: the effects of the MTO on taxes paid and on reported gross incomes 6 years after firms were transferred into the MTO were between 1.5 and 2.5 times larger than they were 2 years after being moved to the MTO, despite the fact that staffing levels and enforcement actions from the MTO (as well as from Primary Tax Offices, or PTOs) remained essentially constant.

[3]This very high ratio of revenues received to additional administration costs makes the MTO tax administration reform distinctively cost-effective compared to tax administration interventions studied elsewhere. For example, Gadenne (2017) finds a near one-to-one cost/benefit ratio from a tax administration program implemented by the Brazilian Development Bank in 1998, which provided municipalities with subsidized loans for investments in items such as improved taxpayer registry systems, streamlining of audit processes, and simplifying taxpayer interactions with authorities.

[4]Throughout the period studied (2003-2011), third-party cross-checking of VAT payments in Indonesia was a manual process conducted by tax office staff, limiting the self-enforcing aspect of VAT, and thus increasing the scope for large VAT effects once taxpayers were moved to a higher enforcement regime.

[5]The large impact of improved tax administration that we find is not mechanical – the fact that the level of tax collection may be low in a developing country like Indonesia does not necessarily imply, a priori, that the *derivative* of tax collections with respect to improved administration would be high. This is in contrast to, for example, the comparison between *de jure* changes in the tax base and *de jure* tax rates, where, as suggested by Suárez Serrato and Zidar (2018), there is a mechanical interaction between tax base and tax rate changes.

One question raised by our model is whether the impacts come from previously hidden transactions being brought 'on the books,' or instead whether the impacts come from a greater scrutiny of deductions or better tax collection of tax arrears. We find that the creation of the MTOs also led to an increase in reported revenues, reported costs, the reported number of permanent employees, and a higher reported wage bill. The findings that reported costs, revenues, and taxable income all increase at roughly similar rates, with no impacts on reported profit margins or collections as a share of taxes due, suggest that the MTO may have led to more of the business being reported to the tax authority.

While the MTO can affect enforcement in many ways, our model suggests one mechanism in particular that we can investigate in the data: a reduction in size-dependent enforcement.[6] We investigate this using detailed data on a few types of enforcement activities tracked consistently by the government – formal audits and letters sent to taxpayers regarding late VAT payments and underpayment. In the standard (i.e., non-MTO) tax administration, with low staff-to-taxpayer ratios, we show that tax staff prioritize their efforts by focusing on the largest taxpayers. Given this, firms may want to avoid growing too large and drawing the attention of the tax authorities. By contrast, we document that the tax offices with more tax staff (i.e., MTOs) pay attention to taxpayers more uniformly, regardless of firm size. Thus, while the effective tax rate may increase for smaller firms who are moved to the MTO – since they face higher enforcement overall – the better tax administration eliminates the additional enforcement tax on firm growth.

We next compare the tax administration reform to a second reform that changed the *de jure* corporate income tax rate schedule. In 2009, Indonesia changed from a system with progressive corporate income tax rates (i.e., a system with three marginal rates, ranging from 15 to 30 percent, with the marginal rate based on a firm's taxable profits) to a flat 28 percent corporate income tax rate, with discounts given as a nonlinear function of a firm's gross revenues. The flat 28 percent corporate income tax rate was then lowered in 2010 to 25 percent, with a proportionate adjustment to the revenue-based discounting scheme. This differential tax change, in which the marginal tax rate moved from being a function of net *profits* to being a function of gross *revenues*, meant that firms faced different marginal tax rate changes as a nonlinear function of the combination of both their gross and net revenues.[7]

[6] These ideas are related to Bigio and Zilberman (2011), who show in a more general setup that this type of size-dependent enforcement can be optimal, even if it leads to distortions.

[7] Corporate tax schedules based on firm revenues, rather than taxable income, are currently used by several other developing countries, including Costa Rica, India, Thailand, and Vietnam (Bachas and Soto, 2018).

We exploit these changes to estimate the elasticity of taxable income (ETI) with respect to the net of tax rate. Following Gruber and Saez (2002) and others, we instrument for the change in a firm's marginal tax rate by applying the new tax formula to gross and net revenue reported by the firm in the pre-period. This approach isolates the variation in changes in marginal tax rates stemming only from the tax schedule change, and has strong predictive power, with a first-stage F-statistic of over 3,000.

We estimate an elasticity of taxable income of 0.59. This implies that, perhaps surprisingly, corporate income taxes for relatively large firms in a developing country setting are not vastly more elastic than in developed countries.[8] We also investigate whether the ETI differs depending on whether firms have been moved to the MTOs or not. While our point estimates suggest that the ETI is lower for firms that are in the MTO than for those that are not, we cannot reject that the ETI under the two different enforcement regimes is the same. The results suggest that the effects of the MTO documented above do not come primarily through a reduction in the ETI.

Finally, we can put our estimates together to compare raising revenue through improvements in tax administration and increases in statutory tax rates. Specifically, we can compute, using our estimated ETI, how much marginal corporate income tax rates would have had to be increased to raise the same amount of revenue that the government obtained from the same corporate income tax by improving tax administration. The answer is substantial: to obtain the increases in corporate income taxes paid by MTO taxpayers alone, top marginal corporate income tax rates on all firms would have had to be raised by 8 percentage points (i.e., from 30 percent to 38 percent).[9]

To compare the welfare impacts of tax administration improvements and tax rate changes, one needs an additional component – namely, the change in firms' administrative costs for complying with the new regime. While this change is unobserved, we adapt the framework of Keen and Slemrod (2017) to characterize the conditions under which the welfare gains

[8]The estimated ETI for Indonesia is larger than the estimate from Gruber and Rauh (2007) using Compustat data in the United States (0.2), but close to Dwenger and Steiner's (2012) estimate using a pseudo-panel of German corporate taxpayers' average tax rates (0.6). Our estimate is, however, smaller than Bachas and Soto's (2018) estimates from Costa Rica (3-5), though the firms in their sample are very small, with revenues of only approximately USD100,000 - USD200,000, making them between 8-17 times smaller than the medium-sized firms we consider here.

[9]Achieving the total increase in revenue from improved tax administration, including the higher VAT and withholding payments received, would not have been feasible by changing the corporate income tax rates alone (i.e., that would have required raising rates well above the revenue-maximizing rate); likewise, it would not be possible to raise the amount of corporate income tax revenue generated by the MTO from the MTO firms themselves simply by raising the marginal corporate income tax rate.

from raising revenue through improved tax administration exceed those from increased rates. Our results suggest that these conditions are likely to hold unless the additional compliance costs associated with the MTO are extremely high. Since the MTO actually appears to have made compliance for firms easier– firms report higher customer satisfaction when dealing with MTOs than when dealing with PTOs – the conditions seem likely to be satisfied.

In short, our findings suggest that developing country governments may have substantial room to raise revenue through both administrative improvements and raising rates, but that at least in the case of medium-sized firms, the dramatic returns from improved tax administration suggest it is likely to be a particularly important policy tool.

This paper builds on a number of literatures. First, we build on the growing new literature documenting the importance of tax administration in developing countries. Important recent work in the developing world has focused on improvements to third-party reporting (Pomeranz, 2015; Carrillo et al., 2017; Almunia et al., 2017; Naritomi, 2019; Brockmeyer et al., 2019), computerization (Fan et al., 2018), and performance pay (Khan et al., 2016).[10] The reform that we study, coupled with an extensive panel of administrative tax data, allows us to contribute to this literature by understanding the impacts of a change in the overall level of tax administration and by understanding how this sustained increase in tax administration over many years affects firms after they are able to adjust to a new paradigm.

Second, we build on the recent literature understanding the *de jure* impacts of corporate income taxes. While most recent work in the United States and Europe, such as Suárez Serrato and Zidar (2016) and Fuest et al. (2018), focuses on the impact of corporate income tax changes on investment and wages, our paper follows instead in the tradition of Gruber and Rauh (2007) and Kawano and Slemrod (2016) in estimating the elasticity of taxable income for corporate income tax. Recent papers in this literature that use administrative tax data, notably Devereux et al. (2014) and Boonzaaier et al. (2019), use regression kink designs to estimate elasticities based on excess mass at kink points, which requires substantial assumptions restricting heterogeneity in preferences to generate identification of elasticities (Blomquist and Newey, 2017). Our paper, by contrast, uses the large and differential changes in marginal tax rates stemming from Indonesia's tax reform, which generates substantial variation in marginal tax rates and does not require these additional assumptions.

Finally, and perhaps most importantly, this paper bridges these two literatures to highlight the tradeoffs between tax administration and rate changes. Keen and Slemrod (2017),

[10] Other recent work focuses on what to tax, such as Best et al. (2015), who explore whether one should tax profits or revenues in low-information, developing country settings, and the role of liquidity constraints in limiting tax ability (Brockmeyer et al., 2020).

in particular, theoretically show that the key parameter of interest to study the impact of changes in both tax administration and tax rates is their impact on taxable income, and suggest the importance of studying both changes in the same context for comparison. In fact, they specifically point out that "the new wave of empirical literature on the impact of tax enforcement activities has not yet produced estimates of the elasticities our approach shows to be critical." Part of the reason why this has not been done before is that doing so requires clear, credible natural experiments varying both tax rates and administration in the same setting, as well as access to high quality administrative tax data to evaluate the impacts of these changes. Indonesia's reforms, coupled with its rich administrative data, provide a unique opportunity to bring empirical evidence into this broader theoretical debate, particularly in the developing country context.

The rest of this paper is organized as follows. Section 2.2 describes the setting, the two reforms that we study, and the data. Section 2.3 develops a model of corporate tax evasion that guides our empirical approach. Section 2.4 estimates the impact of improved tax administration. Section 2.5 presents the estimated elasticity of taxable income from tax rate reform, and uses this to contrast the tax administration reforms with changes to the tax schedule. Section 2.6 concludes.

2.2 Setting and Data

2.2.1 Corporate taxation reforms in Indonesia

Indonesian taxation is administered by the Directorate General of Taxation (DGT). Overall, in 2005 Indonesia had a tax-to-GDP ratio of 14.9%, which puts it at the 42nd percentile of low and middle-income countries (UNU-WIDER, 2021), and comparable to countries such as Philippines (12.4%), Costa Rica (13.6%), Malaysia (14.3%), Senegal (14.6%), and India (16.4%).

Corporate taxpayers must pay both corporate income tax and value-added taxes, as well as file withholding taxes on behalf of their employees. As in most countries, corporate income taxes are levied on net income (profits), with standard depreciation schedules for capital assets. In our study period, the tax schedule moved from a progressive corporate income tax rate, with three brackets ranging from 10 to 30 percent, to a flat 25 percent rate, with discounts based on gross income (see Section 2.2.1). Value-added taxes are assessed at a flat 10 percent rate, with rebates for exports. Taxpayers remit payments for both corporate income tax and individual income taxes monthly. Annual corporate tax returns follow a

January - December tax year, and must be filed by the end of April of the following year.

Tax administration reform and the introduction of Medium Tax Offices

Indonesia began comprehensive reforms of its tax administration system in 2002, to improve fiscal balance in the wake of the 1997-1998 Asian Financial Crisis. This was the first year it transitioned to a modern, centralized IT system to handle all tax transactions. It also restructured the organization of its tax offices.

The organizational reform had two main features. First, following typical practice worldwide (Lemgruber et al., 2015), large corporate taxpayers were moved to centralized offices, with higher staff-to-taxpayer ratios to allow for more intensive followup. The largest 200 taxpayers nationwide would be serviced centrally by a Large Taxpayer Office (LTO) based in Jakarta. Analogously, the top several hundred taxpayers in each region would be handled by a special Medium Taxpayer Office (MTO) in their tax regions. All remaining corporate taxpayers, as well as all individual taxpayers, would be handled by the network of about 300 Primary Taxpayer Offices (PTOs).[11] We focus on firms serviced by MTOs and PTOs.[12]

Second, the office structure was also reformed. Prior to the reform, tax offices were organized by tax type, such that taxpayers filed different taxes in different locations, and auditing was conducted by a separate network of audit offices (Brondolo et al., 2008). The reorganization centralized all of each taxpayer's payment obligations and auditing into a single office, and put a single contact person, known as an account representative, in charge of each taxpayer. This new centralized organizational structure was identical at LTOs, MTOs, and PTOs.

We study the impact on firms of being assigned to an MTO, as opposed to a PTO. Tax liabilities and procedures are identical for MTO and PTO firms.[13] Instead, the primary difference was that the MTOs had higher staff-to-taxpayer ratios.[14] We focus on the two

[11] Eight "special" tax offices were also created to handle foreign corporate taxpayers, publicly traded companies, and oil and gas firms.

[12] Since LTO firms and firms in the special tax offices are large and easily identifiable, their data could not be shared in a way that would assure anonymity in accordance with Indonesian regulations.

[13] Firms in Indonesia can have multiple branches. Excluding headquarters, the MTO firms in our sample have on average 0.25 branches, while PTO firms have 0.06 branches. The only difference between MTO and PTO treatment of VAT is that PTO firms can file VAT either branch-by-branch or in aggregate; MTO firms report a single aggregated VAT; corporate income taxes are always filed centrally in both PTO and MTO firms. We combine all branches of a given firm to a single observation per firm per year using the common company identifier, so that firms with multiple branches are always treated identically in our analysis.

[14] On net, the MTOs we study have a total taxpayer-to-staff of about 6 as of 2011 (see Appendix Table B.1). These staffing ratios are broadly comparable to other high-intensity tax administration settings in similar countries. For example, large tax offices (LTOs) in other upper-middle countries have a corporate

main types of tax staff who deal with taxpayers: account representatives (ARs), who are the main tax staff responsible for interactions with taxpayers and routine enforcement (including sending letters asking for clarification, calling in taxpayers for meetings, and visiting taxpayers to confirm that firm activities appear commensurate with tax reports); and auditors, who conduct in-depth formal financial audits. Importantly, the MTOs feature a low taxpayer-to-staff ratio: approximately one AR and one auditor for each 17-26 corporate taxpayers. By contrast, at PTOs, each AR and auditor handled between 56 and 125 corporate taxpayers – in addition to hundreds or, in many cases, thousands of individual taxpayers (see Appendix Table B.1). The staff-to-corporate taxpayer ratio was therefore about 4-5 times higher in MTOs compared to PTOs.[15]

Although staff-to-taxpayer ratios were higher, the MTO staff were broadly similar in terms of experience (e.g., account representatives at MTOs had 8.3 years of experience at DGT in 2008, compared with 7.9 at PTOs; see Appendix Table B.1) and had similar scores at baseline on the subjective performance assessments that are explicitly used for promotions (see Appendix Table B.2).

The higher staff-to-taxpayer ratios in the MTO can affect tax revenues in many ways. For example, de facto enforcement levels can increase if ARs handling fewer firms per person in MTOs can spend more time developing detailed firm profiles to help spot evasion. ARs can call in taxpayers for discussions or send letters asking for clarification (both of which are key enforcement activities and are not counted as formal "audits"), and they can do more of these activities per firm in the MTO since they handle fewer firms. The increased ratio of auditors to taxpayers also means that formal audit probabilities may increase at the MTOs, and when audits are conducted, auditors may be able to conduct more detailed audits.

The MTO may also reduce compliance costs, since ARs have more time to answer each firm's questions. In fact, anecdotal evidence suggests that this was the case: a survey of corporate taxpayers in the Jakarta and Banten regions conducted by ACNielsen showed 5 percentage points higher "satisfaction" with tax office interactions at MTOs compared to PTOs.[16] The MTO effects that we estimate should, therefore, be interpreted to include both increased enforcement (through higher ratios of account representatives and formal

taxpayer-to-FTE ratio of 8-8.5 (Crandall et al. (2019)).

[15] As noted, PTO staff also had to handle individual taxpayers, whereas MTO staff could focus exclusively on corporate taxpayers. While we do not know the precise allocation of effort in PTOs between corporate and individual taxpayers, if one assumes that roughly half of PTO staff time was spent on individual taxpayers, then the taxpayer-to-staff ratio is about 8-10 times higher in the MTO.

[16] Summary statistics from the ACNielsen survey were obtained from an internal DGT presentation dated January 2016; the original microdata have not been retained.

auditors), as well as potentially easier compliance.

We focus primarily on the wave of MTOs created in 2007, which covered the vast majority (13 out of 19) of tax regions. Prior to this, in 2004-2006, the new organizational structure was piloted in 6 regions, but the primary tax offices were not yet changed to have the same structure as MTOs (i.e., all taxpayer processes centralized into one office and a modern IT system). Hence, in these pilot districts, the MTOs differed from PTOs on a number of different characteristics (see Appendix Table B.3 for a list of these pilot districts). In 2007, two changes occurred. First, MTOs were created in all remaining 13 regions, with the lists of firms assigned to MTOs developed in late 2006 and officially published in January 2007. Second, the PTOs were reorganized in all regions, so that the PTOs and MTOs would have the same responsibilities, IT, and structure, but now the key difference would be that MTOs would have high staff-to-taxpayer ratios. Therefore, we focus on the 13 regions where MTOs were created in 2007, in order to examine the more intensive staff-to-taxpayer ratios that taxpayers were subject to, holding the overall administrative and organizational structure fixed between the MTO and the PTO, though results are strikingly similar using the full set of MTOs (see Section 2.4.2).

Within each region, taxpayers were assigned to the MTO based primarily on a formula involving pre-period taxpayer size. While neither the exact formula nor the Excel spreadsheets used to assign taxpayers were retained, interviews with tax officials shed light on its inputs. Our understanding is that the formula combined gross income and total taxes paid for the prior three tax years into a score, and the several hundred largest taxpayers in each region were generally included in each MTO. At the time the MTOs were created, the formula was not published, nor were explicit criteria announced as to how the lists would be revised in the future. As of December 2006, when the MTO assignment was conducted, the latest data available to DGT were for tax years 2003-2005, filed in April-May of 2004-2006. On average, about 4 percent of the taxpayers per region – about 330 taxpayers – were initially assigned to each MTO.

Descriptive statistics comparing the taxpayers that are assigned to MTOs vs. remaining in the PTOs are shown in Appendix Tables B.17 through B.20. MTO taxpayers are, as expected, substantially larger than non-MTO firms on almost all dimensions. As such, they account for a large share of taxes even though they are a small number of firms. They are widely represented across sectors and geographies. The manufacturing and mining sectors appear disproportionately likely to be in the MTOs compared to other sectors, likely reflecting the fact that these sectors are more likely to have large firms than other sectors.

The 2009 corporate income tax rate reform

In September 2008, Indonesia passed a new law outlining a restructuring of the corporate income tax rate schedule beginning in tax year 2009. This had two main components: a) corporate tax rates would now be determined according to gross income (i.e., revenues) rather than taxable income (i.e., profits); b) the top marginal tax rate of 30 percent would be cut to 28 percent in 2009, and to 25 percent from 2010 onwards. Other than the change in statutory rates, the other features of the corporate income tax code (e.g., depreciation schedules and allowances) were unaffected by this reform.

Prior to this reform, corporate income tax rates followed a three-tiered schedule defined over taxable income (i.e., bottom-line profits): a rate of 10 percent for the first IDR 50 million (USD 5,000) in taxable income; a rate of 15 percent for the next IDR 50 million; and a rate of 30 percent on all taxable income over IDR 100 million (USD 10,000).

Starting in 2009, however, the system shifted to a flat rate, with discounts given based on gross income (i.e., top-line revenues). For firms with gross income above IDR 50 billion (USD 5 million), a 28 percent rate over all taxable income was applied. For firms with gross income below IDR 4.8 billion (USD 480,000), a 50 percent discount was applied, resulting in a 14 percent rate over all taxable income. For firms with gross income between IDR 4.8 billion and IDR 50 billion, a non-linear schedule was implemented, whereby a taxpayer with IDR g billion in gross income was assessed at a rate of 14 percent over the $(\frac{4.8}{g})$ share of its taxable income, and 28 percent over the remaining share, i.e., the tax rate was $14\frac{4.8}{g} + 28(1 - \frac{4.8}{g})$ percent. In 2010, the 28 percent flat rate was reduced to 25 percent, but the discounts were similar, so the final tax rate in this region became $12.5\frac{4.8}{g} + 25(1 - \frac{4.8}{g})$ percent, with a similar notch at IDR 50 billion in gross income. Note that the tax is still levied on a firm's taxable income; however, the tax rate charged depends on the firm's gross income.

Figure 9 illustrates the marginal tax rate under the original regime (Panel A) and the post-reform regime (Panel B). The x-axis, which determines the marginal tax rate, is different in the two regimes – it is based on taxable income (i.e., profits) in Panel A, and on gross income (i.e., revenues) in Panel B. We exploit this change, which meant that taxpayers with different combinations of gross and taxable income faced different changes in their marginal tax rate, in the empirical analysis below.[17]

[17] The formula creates a notch at IDR 50 billion in gross revenue, where the tax rate on all taxable income jumps discontinuously from 26.65 to 28 percent. The data confirm that there is bunching at the notch, with the density of taxpayers falling discontinuously by about 30 taxpayers in each IDR 1 billion bin to the right of it (see Appendix Figure B.17). However, since the notch is on gross income, not taxable income, this may understate the true elasticity, since many margins available to taxpayers to affect taxable income (i.e.,

2.2.2 Data

We obtained anonymized microdata covering all corporate taxpayers registered in the regional tax offices where an MTO was ever created, from 2003 through 2011.[18] These data include detailed information on corporate income reporting (from corporate income tax forms), employment and wage bills (from employee income tax withholding forms), daily payments data from the Treasury (separated for corporate income tax, VAT, and withholding), and administrative information of tax audits and VAT tax assessments, including the dates and types of assessment-related letters sent to taxpayers. We use reported income data from original corporate income tax filings only (that is, excluding correction filings). We aggregate tax payments data from all branches of a given corporate taxpayer to a single observation per company-year. See Appendix B.2 for details.

2.3 Theoretical Framework

We build a simple model of corporate tax evasion to examine how the levers empirically assessed in this paper (tax administration and tax rates) might affect corporate taxpayers' business and evasion decisions. Broadly speaking, firms can evade taxes in two ways. They can evade taxes by hiding pieces of business activity from the government, i.e., keeping certain transactions, certain customers, or certain types of its business 'off-the-books.' In this case, the firm pays an evasion cost to keep this piece of its business secret, and then does not report any revenues, costs, or taxes from that piece of its business. This type of extensive margin evasion is akin to what Pomeranz (2015) refers to as 'Omission.' For this type of evasion, the key point is that all revenues and costs associated with the evaded activity are hidden. A second type of evasion is to misreport costs (or revenues) to reduce tax liability on business activities that the government knows about. This type of intensive-margin evasion is central to many models of tax evasion, such as Best et al. (2015); this is referred to as 'Distortions' in Pomeranz (2015).

We build a model in Section 2.3.1 that focuses on the first type of evasion – omission of complete transactions or even entire business lines – to illustrate key mechanisms. We present a generalized version of the model that includes both types of evasion in Appendix B.3. Section 2.3.2 considers changes in tax enforcement and tax rates in this model. Section 2.3.3

deductions) may not be available for adjusting gross income.

[18]Since the data are anonymized per DGT regulations, we cannot match it to external datasets – such as surveys of manufacturing – to analyze the effect of MTOs on other outcomes. We also do not observe MTO status in those other datasets, so cannot independently use them for analysis.

adapts Keen and Slemrod (2017)'s analysis, which generalizes the arguments of Feldstein (1999), Chetty (2009), Saez et al. (2012), and others to allow for changes in tax enforcement in addition to tax rate, to provide conditions for the welfare effects of tax rate and tax administration changes in the corporate taxation setting. Section 2.3.4 then extends the model to consider what happens when enforcement is not uniform across firms.

2.3.1 Setup

Suppose a firm has a continuum of business lines indexed from $[0, L]$.[19] Each business line has convex costs of production, so that the revenue from line l is y_l and the costs are given by the convex function $c(y_l)$. We assume that all lines are symmetric, and normalize output prices to 1. Pre-tax profits from line l are, therefore, $\pi(y_l) = y_l - c(y_l)$. With no taxes, the firm sets $c'(y_l) = 1$ and produces equally on all business lines.

Following Best et al. (2015), we assume that a proportion μ of costs are deductible from taxes. Setting $\mu = 1$ is therefore a pure, non-distortionary profit tax; setting $\mu = 0$ is a pure output tax. Since we examine firms that pay a mix of VAT (for which labor and many other expenses are not deductible) and corporate income taxes (for which these costs are deductible), we assume $0 < \mu < 1$.[20] Firms pay a tax rate τ on revenues less the deductible component of costs.

For a line on which it pays taxes, the firm therefore solves:

$$\max_{y_l}(1-\tau)y_l - (1-\tau\mu)c(y_l) \tag{2.1}$$

which yields the optimum conditions:

$$c'(y^p) = 1 - \tau\frac{1-\mu}{1-\tau\mu} = 1 - \tau_E \tag{2.2}$$

where $\tau_E = \tau\frac{1-\mu}{1-\tau\mu}$ is the firm's effective tax rate and y^p is the optimal level of production y for firms that pay tax.[21]

We now introduce the possibility that firms can hide activity from certain business lines by paying an evasion cost. If a firm evades on line l, it does not report either revenue y_l

[19] We use business 'lines' as the units of analysis here, but one could imagine these 'lines' also refer to specific customer relationships or even specific transactions, where there is heterogeneity among customers or transactions in terms of the ease of keeping various transactions 'off-the-books.'

[20] See Best et al. (2015) for a detailed discussion of why setting $0 < \mu < 1$ may be optimal.

[21] In this simple model, conditional on paying taxes, the firms report c truthfully. We generalize the model to allow misreporting of c in Appendix B.3.

or costs $c(y_l)$ to the government, and does not pay taxes on this production. Suppose that the cost of hiding line l is given by $\alpha b(y_l)h(l)$, where both $b(y_l)$ and $h(l)$ are increasing and continuous and $b(y_l)$ is convex. The business lines l are implicitly ordered in terms of how difficult they are to evade, from easiest to hardest; this heterogeneity across lines is captured by $h(l)$.[22] We assume that the easiest line can be evaded at cost 0 and that $h'(0) = 0$, so that firms will always evade at least somewhat. The fact that $b(y_l)$ is increasing in output y_l captures the idea that larger business lines are more easily detectable and harder to evade, and more generally, that there may be an interaction between real decisions and evasion costs (Slemrod, 2001). For example, with some probability, each worker in a given business line, or counterparty in a transaction, might reveal information about evasion to the government (as in Kleven et al., 2016). Finally, the parameter α captures the level of enforcement. We assume these evasion costs are real costs, and not transfer costs.[23]

Given this setup, the firm will make its evasion and production decisions as follows. If line l is hidden, the firm sets output y to solve:

$$\max_{y_l} y_l - c(y_l) - \alpha b(y_l)h(l) \tag{2.3}$$

and so sets:

$$c'(y^e) = 1 - \alpha b'(y^e)h(l) \tag{2.4}$$

where $y_l^e(\alpha)$ is the optimal level of output under evasion.

Firms choose which lines to evade and which to pay taxes on. In particular, the firm chooses the point l^* such that the firm is indifferent between evading on line l^* or not, comparing after tax profits with and without evasion. The indifference point l^* is given by the solution to the equation:

$$y_{l^*}^e(\alpha) - c(y_{l^*}^e(\alpha)) - \alpha b(y_{l^*}^e(\alpha))h(l^*) = (1-\tau)y^p - (1-\tau\mu)c(y^p) \tag{2.5}$$

Total taxes collected are therefore given by $\tau \int_{l^*}^L y_l^p - \mu c(y_l^p)$, where $z \equiv \int_{l^*}^L y_l^p - \mu c(y_l^p)$ is the firm's taxable income. The fact that after-tax profits if evasion takes place are strictly

[22]Heterogeneity in $h(l)$ could come from certain customers being more willing to engage in under-the-table transactions, or certain types of businesses being easier to conduct with informal labor, for example.

[23]These evasion costs could take many forms. Grubert and Slemrod (1998), for example, discuss location shifting to lower-tax locations as an example. In this context, it could entail costs to facilitate financial evasion (e.g., using cash instead of banks, or other financial mechanisms); having to pay employees higher wages to compensate them for forgone social security payments; or inefficient production technologies to keep factories from being detected. Fines (which would be transfers, not real costs) are empirically very small in our context, accounting for only 0.08% of tax revenues collected between tax years 2004 and 2011.

decreasing in l gives a unique solution l^*.

2.3.2 Changes in enforcement and tax rates

There are several remarks worth making about the effects of increasing enforcement (α) in this model.[24] Increasing α leads to more business lines being reported, i.e., a lower optimal level l^*. There are two forces, which go in the same direction. First, even holding y_l^e fixed, increasing α has a direct increase in the costs of evasion for a given line l. Second, from equation (2.4), increasing α further reduces y_l^e – and hence profits under evasion – for a given business line l. Real output will therefore decrease for those lines that continue to evade, but firms will evade on fewer lines.

What happens at the margin when a business line switches from being hidden to being reported? First, there is a large and immediate jump in *reported* revenues y, costs c, and taxes paid that comes from the line now being reported to the tax authorities. Note that in this model, reported revenues *and* costs both increase in response to increased enforcement, as all aspects of the new business lines are reported to the government.

The effect on *real* activity of the marginal line l that switches to becoming formal is ambiguous, as there are two offsetting effects. When a business line switches from being hidden to being taxed, the additional 'enforcement tax' – $\alpha b'(y)h(l^*)$ in equation (2.4) – disappears. However, the firm now has to pay a distortionary tax on that line, given by the effective tax rate $\tau_E = \tau\frac{1-\mu}{1-\tau\mu}$ from equation (2.2). Real output on that line will increase if and only if the size-dependent 'enforcement cost' effect is greater than the effective tax rate, i.e.:

$$\alpha b'(y)h(l^*) > \tau\frac{1-\mu}{1-\tau\mu} \tag{2.6}$$

For real activity as a whole to increase with α, equation (2.6) would need to hold, and the increase in real activity from these marginal lines induced to be reported would need to be larger than the decline in output lines that continue to evade. While the results are ambiguous, the point is that real activity could actually increase at the margin as more activity is brought into the tax net. Figure 10 shows an example of an increase in enforcement α in the case where real activity increases (i.e., where the real distortions from the enforcement

[24] While we focus on increased enforcement (α), improved tax administration can also make paying taxes easier. This can be incorporated by modifying the taxpayer's maximization problem in equation (2.1) to be $\max_{y_l}(1-\tau)y_l - (1-\tau\mu)c(y_l) - \kappa\tau(y_l - \mu c(y_l))$, where κ is the administrative cost associated with filing taxes of size $\tau(y_l - \mu c(y_l))$. The effects of reducing κ would be similar to increasing α for lines induced to start paying taxes by the change; the only difference is that for infra-marginal lines, reducing κ would also increase real output among lines that are already paying taxes, rather than those who are evading.

tax on the margin are greater than the real distortions from taxation).

Changing statutory tax rates (i.e., increasing τ) in the model has several effects. First, from equation (2.2), it decreases real activity on all tax-paying business lines as long as $\mu < 1$. Second, because it decreases profits on tax-paying business lines, equation (2.5) shows that evasion will also increase. The model also implies the possibility of complementarity between tax administration and tax rates, as in Besley and Persson (2014). This is because, from equation (2.5), a higher level of enforcement α implies that the elasticity of taxable income with respect to tax rates will be smaller in absolute value (i.e., $\frac{\partial^2 z}{\partial\tau\partial\alpha} > 0$), though whether this is quantitatively important is an open empirical question.

2.3.3 Welfare analysis

Social welfare in this context is given by:

$$W = \underbrace{\int_{l^*}^{L} (y_l^p - c(y_l^p)) - \tau z}_{\text{firm post-tax profits from taxed business lines}} + \underbrace{\int_0^{l^*} y_l^e(\alpha) - c(y_l^e(\alpha)) - \alpha b(y_l^e(\alpha))h(l)}_{\text{firm post-tax profits from evaded business lines}} + \underbrace{v\left(\tau z - a(\alpha)\right)}_{\text{social value of public funds}} \tag{2.7}$$

where $v \geq 1$ is the marginal value of government funds and $a(\alpha)$ are administration costs.

We can use this expression to calculate the welfare effects of changes in both enforcement levels and tax rates. We define private compliance costs $\gamma = \int_0^{l^*} \alpha b(y_l^e(\alpha))h(l)$ to simplify notation.

To calculate the effect of changing enforcement levels on welfare, we take the derivative of (2.7) with respect to tax enforcement α and apply the envelope theorem, which yields:

$$W_\alpha = -\frac{d\gamma}{d\alpha} + v\left(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}\right) \tag{2.8}$$

where $\frac{d\gamma}{d\alpha}$ is the change in private compliance costs.

This change in private compliance costs is unobserved. Instead, we estimate the change in net government revenue with respect to improved tax administration (i.e., $\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}$); see Section 2.4. This allows us to bound how large the change in private compliance costs would have to be for the change in administration to be welfare-improving.

We can do a similar calculation for the welfare effect of a change in tax rates. Taking the derivative of (2.7) with respect to τ and applying the envelope theorem yields:

$$W_\tau = -z + v\left(z + \tau\frac{dz}{d\tau}\right) = -z + vz\left(1 - \frac{\tau}{1-\tau}\varepsilon_{1-\tau}\right) \tag{2.9}$$

where $\varepsilon_{1-\tau}$ is the elasticity of taxable income with respect to the net of tax rate.

This simple framework also allows us to ask whether, if the government is seeking to raise an additional dollar of revenue, it is better to do so through improvements in tax administration or increases in tax rates. We begin by calculating the tax change such that government revenue is the same after a marginal change in tax administration (i.e., a change in α). Given that net government revenues $R = \tau z - a(\alpha)$, we can write:

$$\frac{dR}{d\tau} = \tau\frac{dz}{d\tau} + z = z\left(1 - \frac{\tau}{1-\tau}\varepsilon_{1-\tau}\right) \tag{2.10}$$

$$\frac{dR}{d\alpha} = \tau\frac{dz}{d\alpha} - \frac{da}{d\alpha} \tag{2.11}$$

This implies that:

$$\frac{d\tau}{d\alpha}|_R = -\frac{\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}}{z\left(1 - \frac{\tau}{1-\tau}\varepsilon_{1-\tau}\right)} \tag{2.12}$$

Thus, armed with the elasticity of taxable income, we can ask how large a change in tax rates one would need to get the equivalent revenue change from improved tax administration, and vice versa. After estimating the elasticity of taxable income with respect to the net of tax rate in Section 2.5.2, we compute this ratio (i.e., $\frac{d\tau}{d\alpha}|_R$) in Section 2.5.3.

We can use the rate of substitution between tax administration and tax rates in equation (2.12) to ask whether, if the government seeks to raise more revenue, should it do so via improved tax administration or by changing tax rates? Since we are considering marginal changes, this is equivalent to asking whether a revenue-neutral increase in administration and corresponding cut in rates would be welfare improving or welfare decreasing; that is, by evaluating:

$$dW = W_\tau\frac{d\tau}{d\alpha}|_R + W_\alpha \tag{2.13}$$

Substituting W_τ, W_α and $\frac{d\tau}{d\alpha}|_R$ from equations (2.9), (2.8), and (2.12) above, this is equal to:

$$dW = \left(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}\right)\frac{1}{1 - \frac{\tau}{1-\tau}\varepsilon_{1-\tau}} - \frac{d\gamma}{d\alpha} \tag{2.14}$$

By estimating the change in net tax revenue with respect to administration (i.e., $\left(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}\right)$)

and the change in tax revenue with respect to tax rates (i.e., by estimating $\varepsilon_{1-\tau}$), we observe all of the parameters in equation (2.14) except the change in private compliance costs $\frac{d\gamma}{d\alpha}$. Nevertheless, equation (2.14) is useful in several respects. First, holding $\frac{d\gamma}{d\alpha}$ fixed, improving tax administration is likely to be a good idea when both $(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha})$ is large – i.e., gains from improved tax administration are large – and when $\varepsilon_{1-\tau}$ is large – i.e., the behavioral elasticity with respect to tax rates is large. Both will turn out to be true in our empirical context. Second, and more precisely, we can use equation (2.14) to bound how large $\frac{d\gamma}{d\alpha}$ has to be for a change in tax administration to be welfare-improving relative to an equivalent change in tax rates (see Section 2.5.3).

2.3.4 Size-dependent enforcement

The government can affect not just the level of enforcement α, but the degree to which enforcement is size-dependent, i.e., the degree to which the government places higher enforcement costs on larger firms.

Suppose the government conditions its enforcement effort on reported income z, i.e., it spends more effort investigating the unreported business lines of firms that appear larger based on their reported income. For example, the government may choose to allocate the effort of tax collection staff to firms that it observes to be larger based on the tax data it collects. In this case, we can write evasion costs as $\alpha m(z)b(y)h(l)$ with $m' > 0$, where z is the total reported taxable income defined above. We write m as a function of taxable income z to simplify notation, but in principle in this model similar logic applies as long as m is a function of any other reported outcomes of the firm (i.e., total reported revenue, total reported employees, etc).[25]

With this new evasion cost that is a function of total reported income z, the indifference

[25]The government can also potentially alter the slope of the $b(y)$ function, i.e., the degree to which evasion costs increase with the size of unobservable business lines. If these actions increased enforcement activity while making it less size dependent – i.e., increasing α but decreasing $b'(y)$ – the analysis above (e.g., equations (2.4) and (2.5)) shows that one can both increase tax payments while reducing distortions on untaxed business lines at the same time.

condition in equation (2.5) for the marginal line to evade (l^*) then has an additional term

$$\underbrace{y^e_{l^*}(\alpha) - c(y^e_{l^*}(\alpha)) - \alpha m(z) b(y^e_{l^*}(\alpha)) h(l^*)}_{\text{profit from marginal line evading}} = \underbrace{(1-\tau)y^p - (1-\tau\mu)c(y^p)}_{\text{profit from marginal line not evading}} - \underbrace{m'(z)\int_0^{l^*} \alpha b(y^e_l(\alpha))h(l)}_{\text{loss from having higher evasion costs on evaded lines}} \tag{2.15}$$

We can use equation (2.15) to consider what happens when the government changes m'. A flattening of the evasion cost (i.e., holding the level of $\alpha m(z)$ fixed, but reducing m') decreases the benefit from evading, and so will lead the marginal firm to evade less than an equivalent amount of enforcement with a flatter m'. Note also that, by the arguments above, this can also lead to a further increase in real activity. This suggests that one may be interested not just in the level of distortion, but also in the degree to which it is dependent on firm size, as increasing enforcement in a way that makes it less size-dependent will be more effective than increasing enforcement in a way that makes it more so.[26] We explore these issues empirically in Section 2.4.3.

2.4 The Impact of Improved Tax Administration

2.4.1 Empirical strategy

We begin by estimating the impact of being assigned to more intensive tax administration in the MTOs. As described in Section 2.2.1, taxpayers were assigned to MTOs in 2007 based on an increasing function of pre-assignment gross income and total taxes paid (see Appendix Figure B.1 which plots the probability of MTO assignment separately by gross income and total taxes paid, and Appendix Figure B.2, which shows this jointly).[27] This implies that the assigned taxpayers were inherently different from other ones: they were larger and paid more taxes. Therefore, we cannot simply compare the two types of taxpayers.

Instead, we compute taxpayer-level balancing weights that match taxpayers assigned to the MTO with other unassigned taxpayers based on their 2005 gross income, total taxes

[26]It important to note that just because size-dependent enforcement creates distortions does not imply that it is not optimal; in more general models, such size-dependent enforcement may be optimal, even accounting for these additional distortions (see, e.g., Bigio and Zilberman, 2011).

[27]We do not know the precise assignment formula, so we cannot use a regression discontinuity design. While the probability of MTO assignment is strongly increasing in these two variables, we also do not observe a sharp discontinuity. See Appendix Figure B.1.

paid, and region. This step brings the pre-assignment outcome levels of the two groups close together via weighting. We then exploit the panel structure of the data to estimate the effect of MTO assignment using a taxpayer-level weighted difference-in-differences design (WDD), with firm fixed effects.

To compute balancing weights, we follow the "entropy-balancing" methodology proposed by Hainmueller (2012). This method computes exact weights (for the untreated group) such that a set of desired pre-treatment characteristics of the untreated group match those of the treated group, and chooses the set of weights that achieves balance that minimally deviates from uniform weights. This methodology is particularly appropriate in a situation where the true functional form of the propensity score is unknown because it does not impose a rigid functional form on the propensity weights, and in this case, this approach provides better pre-treatment balance than standard inverse propensity-score methods (Hainmueller, 2012; see also the related discussion in Athey and Imbens, 2017 and Athey et al., 2018).[28]

As is standard in the matching literature, we impose a common support restriction on the variables used to match. These distributions are shown in Appendix Figure B.3. In our main specification, we drop firms that fall within the top or bottom 2.5 percent of either the control or treatment distribution of the key matching variables; this implies that we exclude very large firms within the MTO and very small firms not in the MTO. Appendix Table B.6 shows robustness to more or less restrictive common support restrictions.

Since the latest corporate income tax filings available to DGT at the time of the MTO assignment (December 2006) were for tax year 2005, we compute balancing weights by matching on 2005 gross income and total taxes paid.[29] We define treated firms as those who were selected in the initial assignment in 2007. In constructing the variables used for matching, we use corporate income tax filing dates and tax payment dates to discard any data that was neither filed nor paid by December 2006. Columns (1) and (2) of Table 2.1, as well as columns (1) and (2) of Appendix Tables B.10 and B.11, show that the resulting weights produce weighted samples that are broadly balanced not only on the targeted variables (2005 gross income and total taxes paid), but on other variables as well.

[28]We replicate all main findings using inverse probability weights (Abadie and Cattaneo, 2018). Results are qualitatively similar and, if anything, generally slightly larger (Appendix Table B.5).

[29]While we believe that data for three baseline tax years (e.g., 2003-2005) were considered to assign taxpayers to MTOs, neither the formula used nor the procedure for handling missing data (e.g., data not yet filed as of December 2006) are available. Matching on the 2005 level, rather than using all three years, allows us to check whether both sets of matched taxpayers are on similar pre-treatment trends. Matching on all three years (2003-2005) instead of just 2005, which also allows us to match on growth rates in addition to levels, produces similar estimates (Appendix Table B.5).

We then estimate the effect of MTO assignment using weighted difference-in-differences. We define a variable M_{iFC} as a dummy for firm i being in the first cohort of MTO assignment.[30] We then estimate the reduced form effect of MTO assignment in 2007 as follows, where each taxpayer is weighted by its respective balancing weight:

$$Y_{it} = \alpha + \beta^{RF} \left(M_{iFC} \times 1_{t>2005} \right) + \delta_t + \delta_i + \epsilon_{it} \tag{2.16}$$

where Y_{it} is the outcome of interest of taxpayer in year t, δ_i is a taxpayer fixed effect, and δ_t is a year fixed effect. Because corporate income taxes for year 2006 are only filed in April-May 2007, four to five months after taxpayers began being serviced by the MTO, we consider 2005 as the last pre-treatment year, so that any taxes for tax years 2006 or later could have been treated. We estimate equation (2.16) for taxpayers from the 13 regions whose MTOs were created in 2007, using data from tax years 2003-2011.[31] Standard errors are clustered by taxpayer.[32] We also estimate an event study version of equation (2.16) where we estimate separate β^{RF} coefficients by year, which allows us to assess whether these firms were on similar trends in the pre-period, and to assess changes in the MTO's impact over time.

To account for the fact that some firms in the control group were moved to the MTO starting in 2009, we also estimate an instrumental-variables version of equation (2.16), i.e.,

$$Y_{it} = \alpha + \beta^{IV} M_{it} + \delta_t + \delta_i + \epsilon_{it} \tag{2.17}$$

where we instrument for M_{it}, the actual MTO status of firm i at time t, using $(M_{iFC} \times 1_{t>2005})$. This is just a re-scaling of equation (2.16), but may provide a more accurate magnitude for the treatment effect of treated firms being moved to the MTO. The first-stage of this equation is quite strong, with an F-statistic over 6,000 – see Appendix Table B.8. The first stage is shown year-by-year in Appendix Figure B.5.

[30] During the first year of the MTO, firms' taxpayer ID codes were gradually converted to reflect the MTO status. We therefore define M_{iFC} as 1 if the firm's corporate income taxes were filed with an MTO code in 2007 or 2008, i.e., prior to the next wave of MTO expansions in 2009. The first tax year affected for this cohort was 2006, for which final tax returns were filed during calendar year 2007.

[31] We end our analysis in 2011 as there were substantial expansions in the number of firms assigned to the LTO in 2012 (which could create attrition), as well as changes in which firms were in MTOs.

[32] Appendix Table B.7 presents robustness to clustering standard errors at the taxpayer's origin tax office level and at the region. Results are very similar.

2.4.2 Results

Impacts on tax collection

As discussed in Section 2.3, the key parameter needed to estimate the impact of a reform in tax administration is the effect on government revenue. Figure 11 begins by showing the impact of the MTO on total tax payments year-by-year. The left-hand side variable is taxes paid in 2007 billions of Rupiah (IDR 1 billion = USD 100,000 at 2007 exchange rates), where we use the Indonesian GDP deflator to deflate all nominal values to their 2007 equivalents.[33]

Panel A presents the time series for each of the two groups of taxpayers (those assigned to the MTO 2007 group, and those not assigned), where firms are weighted using the balancing weights. Panel B shows the full estimates using equation (2.16). In both panels, the year variable is the tax year, and includes payments by all branches of the same firm for that tax year made up to six months following the end of the tax year.[34] Recall that the MTOs were established by a January 2007 decree and took effect a few months thereafter, before the filing date for 2006 tax year tax returns. We therefore consider 2005 as the final pre-period year, 2006 as partially affected, and 2007 as the first full MTO year.

Examining the pre-period – 2003-2005 – shows that the two sets of firms have similar pre-trends. The two groups of firms match almost exactly in Panel A; the regression version in Panel B shows that the pre-period is flat, indicating no differential pre-trends. This is not mechanical, as we only matched on the 2005 data, rather than on the full 2003-2005 period (i.e., the trends).

The MTO had a large impact. There is a large initial effect of the MTO: for firms assigned to the MTO, tax payments increased in 2006 (the first year that could be somewhat affected by the MTO), and tax payments increased by approximately IDR 400 million per firm by 2007, the first year the MTO was fully in effect. The estimated treatment-on-treated

[33]Note two facts: a) the outcome variable is in levels (billions of Rupiah), not logs, and b) the weights from the entropy weighting match the weights in the treatment group mean. Combined, these two facts imply that our results capture the average effect of the MTO on treated firms within the common support sample. To the extent there is treatment effect heterogeneity among firms in terms of percent increases, we will nevertheless capture the true "average effect" on revenue that the government captures. However, these estimates may underestimate the total extent of revenue increases: if the larger firms that we exclude due to our common support restriction had similar percent increases as the firms in our sample, they will have larger impacts in levels than we estimate here. This will not, however, affect the comparison to tax rate changes in Section 2.5 below, since the samples for both are identical.

[34]Taxpayers typically pay VAT and estimated corporate income taxes monthly, and then are required to file a corporate income tax return by April of the following year. We include all tax payments for a given tax year made during that tax year, and in the six months thereafter; that is, 2007 tax payments include all payments made for tax year 2007 and remitted on or before June 30, 2008. We impose this time limit to focus on payment of each year's taxes due, rather than retrospective payments of delinquent taxes.

effect for the MTO in 2007 represents an increase of 64 percent (over the treated group's counterfactual mean in 2007) for affected firms. The impact continues to grow over time: by 2011, the impact of the MTO increased further, to IDR 605 million (an increase of 129 percent over control firms in the same year). The difference between the effect in 2007 and 2011 is statistically significant (p-value of 0.055). Importantly, the MTO effect is entirely driven by firms actually assigned to an MTO, as tax payments for the control firms remain relatively flat following MTO creation.

Panel A of Table 2.1 shows the results in regression form, based on estimating equations (2.16) and (2.17). For each variable, columns (1) and (2) show the weighted pre-treatment (i.e., 2005) means for the treatment and control group, showing that taxpayers appear balanced not just on the variables that we explicitly match on (total tax payments and gross income), but also on various sub-components of taxable income as well.

We show the reduced form and IV estimates, respectively, in columns (5) and (6). On average, total tax payments increased by IDR 525 million (USD 52,500).[35] About two-thirds of the increase comes from higher VAT collections; and the remaining third comes from higher corporate income tax and other income tax (e.g., withholding) payments. Appendix Table B.11 further disaggregates these tax payments.

As a benchmark of magnitude, we compute the counterfactual control complier means by subtracting the estimated treatment effect from the post-period levels in the treatment group (Katz et al., 2001), shown in column (4) for each variable. We then express the estimated impact of the MTOs as a share of the control complier mean (column 7).

The estimated impacts are substantial. We estimate that the MTOs increased annual tax revenues for affected firms by 128 percent. The increases are seen on all types of taxes: 137 percent for VAT, 111 percent for CIT, and 113 percent for other income taxes.[36]

An important question is whether the impact comes from higher revenues on the part of the treated MTO firms, or a reduction in the non-treated PTO firms, who may have increased

[35]We focus on the IV estimates in the text. The IV estimates adjust for imperfect compliance with the original 2007 MTO list; in particular, some firms were moved to the MTO starting in 2009. By contrast, very few firms were moved out of the MTO during this period: only 44 of the 4,094 firms originally assigned to the MTO were moved to PTOs in 2008-2011 (13 in 2008, 12 in 2009, 11 in 2010, and 8 in 2011). A first stage regression of M_{it} on M_{2007} on our weighted sample (where weights are, as always, determined using 2005 values) yields a first stage coefficient of 0.65 (standard error 0.008; F-statistic is 6,412).

[36]These impacts are not driven by the changes in statutory marginal tax rates: VAT rates are uniform, and Appendix Figure B.15 shows that statutory marginal tax rates (which are a function of firm size, and which change in 2009, as discussed in Section 2.2.1) decrease by only a percentage point or two at most among MTO firms compared with PTO firms, so this cannot explain a 111 percent increase in income tax revenue.

evasion once they learned they would not be in the MTO. Figure 11, which shows dramatic increases in revenues among MTO firms, but flat revenues for non-MTO firms, suggests that the effects are primarily driven by increases for firms being moved to the MTO, rather than decreases for non-MTO firms.[37] To further investigate the possibility of disincentive effects for PTO firms, Appendix Figure B.6 subdivides the PTO firms into larger firms, who could plausibly have been on the margin for inclusion in the MTO, and smaller firms, who were further away from the MTO margin. We find that both sets of control firms appear on similar trajectories, suggesting that the effects are not being driven by changes among those firms who learned they would not be assigned to the MTO.

To estimate the total effect of the MTOs, we need to extrapolate to the full set of firms served by the MTOs, not just those in the common support set. Since the firms excluded from the analysis set tend to be larger than the firms in the estimation sample, different approaches to extrapolation could yield different results. A reasonable lower bound is to assume that all firms experience the same gains, in rupiah terms, as the treatment firms (since the excluded firms are substantially larger). By contrast, a reasonable upper bound is to assume that all firms experience the same percent increase in tax revenues shown in Table 2.1. These are not formal bounds, as we only know the LATE on the estimation set, but they seem reasonable for what to expect.[38] Using this approach, we estimate that the MTOs increased total tax revenues by at least USD 4.0 billion over its first 6 years.

While Table 2.1 presents the effects on gross government revenue, as discussed in Section 2.3, the relevant parameter for welfare is the effect on *net* government revenue; that is, the effect on tax revenue after subtracting off the additional enforcement costs. These additional costs, however, are small. We obtained budget data, as well as the number of corporate taxpayers, for all MTOs and PTOs in Indonesia from 2016 (the earliest available year with complete data for all regional tax offices). We convert the costs to 2007 rupiah using the Indonesian GDP deflator. Since PTOs also handle individual taxpayers, we assume that half of the PTO costs are associated with corporations. (This assumption is inconsequential; results are similar even if we assign all PTO costs to corporate taxation.) Appendix Table B.9 shows that the difference in government enforcement expenditures, per taxpayer, between an MTO and PTO is about IDR 8 million (US $800) per year. These enforcement costs are

[37]Unlike in Almunia and Lopez-Rodriguez (2018)'s study of Spain, where firms strategically bunch below a cutoff to avoid being placed into the Large Taxpayers' Unit, here there is no clear cutoff, and as shown in Appendix Figure B.4, we find no bunching, either in the pre- or post-period.

[38]We can also estimate heterogeneous effects of the MTO within our treated sample. The results, shown in Appendix Table B.12, suggest larger MTO impacts (in rupiah terms) on tax revenue for firms with larger baseline tax revenue. This suggests that the proposed bounds might be reasonable.

thus *almost two orders of magnitude* smaller than the estimated revenue gains (Table 2.1). That is, given an effect on gross taxes paid of IDR 525 million per taxpayer per year, the effect on net government revenues is IDR 517 million per taxpayer per year. Put another way, the government gained a net return of Rp. 517 million for an investment of Rp. 8 million, or a 64-1 return.

Mechanisms: increases in reported business activity, scrutiny of deductions, or increases in collections?

As outlined in Section 2.3, better tax administration could increase tax liabilities in several ways. Taxes due could go up if improved administration results in previously hidden business activities being brought onto the tax rolls, or by increasing the scrutiny of deductions. Tax revenue could also go up if improved administration increases collections (i.e., the share of tax due collected). To investigate these mechanisms, we focus on corporate income tax, for which we observe line-item by line-item reports on each taxpayer's annual tax returns, as well as actual tax payments from the tax authority's treasury system.

The results are shown in Panel B of Table 2.1, and graphically in Figure 12. We present results on several key line items – gross incomes, taxable incomes, corporate income tax due, and the profit margin in Table 2.1. Appendix Table B.10 shows the impact on all major line items of the corporate income tax return in detail, allowing us to decompose how changes in these various line items add up, on net, to a change in taxable income; graphs for many of these additional outcomes, including the costs of sales and other firm expenses, are shown in Appendix Figure B.7.

Several results are worth noting. First, the estimated impact of the MTO on reported corporate income tax due – IDR 0.065 billion – is very similar to the actual increase in corporate income tax payments shown in Panel A – IDR 0.074 billion. This implies that most of the increase in observed corporate income tax payments comes from an increase in reported corporate income due, rather than an increase in collections. In Panel C of Table 2.1, we explicitly report results where the dependent variable is the recovery rate (corporate income tax paid divided by corporate income tax due), and find no impact of the MTO.

Second, the increase in corporate tax due comes from an increase in gross revenues reported. Costs rise at about the same rate, so profit margins remain roughly unchanged. In particular, reported gross income (i.e., revenues) increase by IDR 9.1 billion (US $910,000), or about 76 percent, so firms report more sales once they move to the MTO. Costs of sales (defined as operating expenses, including both material and labor inputs) also increase by

IDR 7.6 billion, or about 82 percent, suggesting that this reflects new business being reported to the government. Other expenses increase as well, at a slightly slower rate, so that on net total reported expenses (costs of sales + other expenses) increase by 77 percent. Since both revenues and total costs increase at about the same rate as revenues, reported profit margins (i.e., net income divided by gross income) remain unchanged. This suggests that the main mechanism through which improved tax administration led to increased revenue is through capturing more top-line business activity on the tax books, as in the theory in Section 2.3, rather than more scrutiny on deductions or increases in collection rates.[39]

Third, the pattern of growth in Figure 12 shows that the MTO firms continue to report growth – in both gross income and taxable income – at substantially higher rates than comparable firms that were not assigned to it. Three years after the MTO introduction, these firms had 41 percent higher gross income than comparable firms; this had increased to 120 percent higher six years after the introduction. This difference is statistically meaningful (p-value 0.007). This implies that the large increases in reported tax revenue from MTO firms over time come not from increased effectiveness of the MTO at collecting taxes due, or from increased scrutiny of deductions, but rather that MTO firms reported substantially higher revenues to the government over time. One possibility, consistent with the model, is that once new business lines become formalized, they no longer need to pay the evasion tax $\alpha b'(y)h(l)$, and that output y increases over time.

Changes in reported employment

We also observe each firm's number of reported employees, which comes from the firms' employee income tax withholding reports. Firms are required to report not just their total wage bill, but also the number of temporary and permanent workers.

In Table 2.2, we examine the effect of the MTO on reported firm employment.[40] We find that the number of permanent employees increases by about 21 percent – an increase of 10 permanent employees per firm (p-value 0.085). These numbers reflect tax withholding payments which are double-reported to workers, so these may be harder for firms to ma-

[39]An alternative view, if firms can manipulate costs directly (as suggested by Carrillo et al., 2017, and as discussed in B.3), is that some costs are not reported if firms are already able to report zero taxable income for other reasons, and so firms report these costs once they are forced to report more revenues. If so, one might expect larger effects on reported costs for these firms with zero taxable income at baseline. To investigate this, Appendix Figure B.8 examines the MTO effects separately for firms with zero and positive baseline taxable income. Although the results are noisy, we find similar effects on reported costs for both sets of firms, with a more rapid response for those firms with positive taxable income at baseline.

[40]Year-by-year figures for employment are shown in Appendix Figure B.10.

nipulate directly (Kleven et al., 2011). The point estimates suggest that the total number of employees increased by the same amount, but the standard errors increase once we include temporary employees, who have much higher variance. This may reflect either true new additional hiring, or increased formalization of temporary workers (since permanent workers receive more employment protections than temporary ones, firms often try to avoid categorizing workers as permanent).

The wage bill for both permanent and temporary employees increases at a similar rate – about 21 percent for permanent employees, and about 24 percent overall. Average yearly wages (computed as wage bill divided by number of employees) increase by about 16 percent for permanent employees, with no meaningful change for temporary employees. This implies that the increases in taxes paid are not coming at the expense of worker wages.

Robustness of MTO effects

We consider robustness checks along multiple dimensions, which indicate that our results are robust to specification choices and are not driven by differential trends among firms that are more likely to be assigned to the MTO. First, Appendix Table B.5 shows that the results are qualitatively robust to alternative weighting strategies. We reproduce our baseline Hainmueller (2012) entropy-balancing weights, and then show results with no weights, using the same matching variables but using a propensity score (estimated both via a logit, in columns 3 and 5, and via a random forest classifier, in column 6) and inverse-propensity score weights (IPW) (see Abadie and Cattaneo, 2018), and using additional years of data for matching, to allow for the possibility that the tax office selected based on growth rates, not just levels. Second, Appendix Table B.6 shows that the common support sample restrictions do not substantively change our qualitative conclusions, though the magnitudes differ somewhat since different samples focus the weights on taxpayers of different sizes, which can matter since all results are in levels. Third, Appendix Table B.7 shows that the main results are robust to the level at which standard errors are clustered.

Fourth, we consider results that include all MTOs started before 2007.[41] As discussed in Section 2.2.1, we focus on the regions where the MTOs started in 2007 in the main specifications, since the PTOs were also reorganized to follow the same administrative structure (albeit with fewer staff per taxpayer) at the same time. We re-estimate equation (2.17), but instead of using $(M_{iFC} \times 1_{t>2005})$ as an instrument, we allow for the fact that MTOs in

[41] We only exclude Central Jakarta's MTO, created in 2004 and thus with no pre-data for matching.

different regions started in different years.[42] The results are presented in column (4) of Table B.6; year-by-year reduced form event-study graphs for total taxes paid and firm reported gross income are also shown in Appendix Figure B.11. The results are qualitatively very similar to the main results, showing quantitatively large and statistically significant increases in tax payments, reported gross incomes, and permanent employees.

Finally, we conduct a placebo analysis among control firms that confirms that our results are not driven by differential trends among firms with characteristics that make them more likely to be assigned to the MTO. We assign placebo firms to mimic the feature that the MTO treatment was assigned as an increasing function of 2005 log gross income and 2005 log total taxes paid.[43] We then reproduce our analysis procedure from Section 2.4.1 on this 'placebo' assignment. Appendix Figure B.13 shows no treatment effects for placebo firms, suggesting that our empirical strategy properly accounts for any differential trends correlated with observable characteristics that predict MTO assignment.

2.4.3 Understanding the MTO's enforcement impacts

The theory in Section 2.3.4 suggests that to understand the impact of improved tax administration, it is important to understand both whether the improved tax administration (the MTOs) increased the *level* of scrutiny of firms, and also how it changed the relationship between firm size and enforcement. In particular, tax administration reforms may be particularly effective to the extent to which they make enforcement less *size-dependent*.

Therefore, we examine both whether the MTOs led to greater enforcement, and how it changed the relationship between firm size and enforcement actions.[44] We have detailed

[42]Specifically, for each region r, we define a variable M_{ir} which is a dummy for whether firm i was in the MTO in region r in the first year it was fully operational. For each region r, we define $\tilde{t_r}$ to be the last year unaffected by the MTO. For example, for the MTOs which opened in 2007, which could have affected 2006 tax returns, we define $\tilde{t_r}$, the last unaffected year as 2005. We use data as of year $\tilde{t_r}$ to do the matching in each region, and we construct our instrument for MTO presence in year t as $\left(M_{ir} \times 1_{t>\tilde{t}_r}\right)$. This notation simply generalizes our estimating equations from Section (2.4.1) to allow for the fact that MTOs started at different times in different regions.

[43]We construct the placebo treatment assignment in three steps. First, we predict the probability $\hat{p_i}$ of MTO treatment for each non-MTO taxpayer i using a logit regression with splines in 2005 log gross income, 2005 log taxes paid, and regional tax office dummies as predictors. We scale these probabilities to match the share of all taxpayers in the analysis sample assigned to MTO (4,181 / 37,629). We then randomly assign non-MTO firms a placebo treatment status according to these scaled probabilities. The resulting probabilities of assignment as a function of baseline taxpayer revenue and taxable income are shown in Appendix Figure B.12, and are similar to the real assignment probabilities shown in Appendix Figure B.1.

[44]In this section, we focus on the fact that larger firms bring higher scrutiny, i.e., $m'(z)$ from the model. In addition, Indonesia's post-2008 CIT regime also has an additional tax on firm size, which comes from the fact that the CIT rates are higher for larger firms, and is applied to all firms regardless of MTO/PTO

data for three types of enforcement actions: formal audits, VAT collection letters, and VAT underpayment letters. These formal actions account for only a small portion of firm interactions with the tax office: an account representative can summon a taxpayer to explain something on their tax form, they can send them a letter for some other purpose, etc., all of which are unfortunately not tracked in the data department's administrative data. However, we focus on these three actions because they are a) relatively serious followup actions and b) systematically logged in the tax department's IT systems in the same way for both MTOs and PTOs. We also have data on corrections to corporate income tax returns filed by taxpayers, though we note that this variable may be harder to interpret if taxpayers file returns that are more accurate to begin with, they would have less reason to correct the returns.

We first document whether the MTO led to greater levels of enforcement. Table 2.3 re-estimates equation (2.17) for corrections to tax returns (Panel A) and VAT underpayment letters (Panel B).[45] We find that being assigned to an MTO leads firms to revise their corporate tax returns. In particular, we find an increase in corporate income tax revisions for *previous* years: that is, once firms enter an MTO, they revise their previous returns (i.e., returns from years prior to the MTO). For tax years in which the original return was filed after the shift, MTO firms are actually less likely to file an amendment, suggesting that original returns filed in the MTO are likely to be more accurate. We find no change in the average level of VAT assessment letters (Panel B).

We then turn to estimating the relationship between the enforcement actions we observe and firm size – the empirical $m(z)$ function – which we measure both in terms of total taxes paid and the number of permanent employees reported by the firm. Figure 13 presents this non-parametrically. We plot these relationships with locally weighted linear regressions separately for MTO firms (in blue) and PTO firms (in red), using the same weights that we have used throughout, so that we are comparing ex-ante comparable firms.

The results tell a consistent story. In virtually all cases, the *level* of enforcement actions is higher at the MTO than for comparably-sized firms serviced by the PTO. However, the *slope* of enforcement with respect to firm-size – i.e., $m'(z)$ – is substantially flatter at the MTO. Thus, the MTO increased enforcement levels, but made enforcement less size-dependent. Following the logic of Section 2.3.4, this raises the possibility that the MTO could have

status. While this type of statutory firm-size could also reduce firm size, since it is not differential based on MTO/ PTO status, we do not focus on it here.

[45]We cannot examine audits here, because we do not have audit data prior to 2008. Audits are tracked by DGT using a separate database that began in 2008.

reduced the size-dependent "enforcement tax" – i.e., firms no longer have to worry that they will face heavier scrutiny when they grow, since they already face high scrutiny.

We test for a change of slope in the $m(z)$ function by estimating the following regressions. We begin with a cross-sectional regression, using the same weights we used in Section 2.4 so that MTO and non-MTO firms are balanced:

$$Y_{it} = \alpha + \beta_1 M_{iFC} + \beta_2 l_{it} + \beta_3 M_{iFC} \times l_{it} + \delta_y + \epsilon_{it} \tag{2.18}$$

The key coefficient is β_3, which shows how the slope of enforcement with respect to firm size l changes for firms assigned to the MTO. This is the regression analogue of Figure 13.

For data on VAT enforcement letters, we observe data in the years prior to 2008 as well. For these variables, we can estimate a difference-in-differences version of equation (2.18):

$$Y_{it} = \alpha + \gamma_1 l_{it} + \gamma_2 M_{iFC} \times l_{it} + \gamma_3 M_{iFC} \times 1_{t>2005} + \gamma_4 M_{iFC} \times l_{it} \times 1_{t>2005} + \delta_y + \delta_i + \epsilon_{it} \tag{2.19}$$

Here, the key coefficient is γ_4, which investigates how the slope on firm size changes once the firm is moved to the MTO. We continue to use the same weights as above. For each table, we examine three separate measures of firm size l_{it}: total taxes paid, the number of reported permanent employees, and the number of reported total employees.

The results of the cross-sectional version estimated using equation (2.18) are shown in Table 2.4; the difference-in-differences results for the VAT enforcement letters estimated using equation (2.19) are shown in Table 2.5. Both tables show similar results: the coefficients on the interaction of $M_{iFC} \times l_{it}$ in Table 2.4, and the coefficients on the interaction of $M_{iFC} \times l_{it} \times 1_{t>2005}$ in Table 2.5, are negative (and statistically significant) for all three variables considered.

The tables thus reinforce the findings from Figure 13: the MTO increases the level of enforcement (shown by the positive main effects on M_{iFC} in the cross-section and $M_{iFC} \times 1_{t>2005}$ in the difference-in-differences regressions, but also reduces the slope of the $m(z)$ function. Quantitatively, the results in Table 2.4 suggest that the slope of the $m(z)$ function was reduced considerably, by between 62 - 100 percent in the case of audits, and by 28 - 90 percent in the case of the VAT letters. These results suggest a potential explanation for the magnitude of the MTO effects over the 6 years we examined them, and in particular why these effects grew substantially over time: by raising the level of $m(z)$, while subsequently flattening its slope, the MTO may have been able to increase tax compliance while simultaneously reducing the tax-induced barriers to firm growth.

One implication of these results is that the impacts of improved tax administration might be smaller for the very largest firms in the country, such as those served by the large taxpayer office (LTO) (which are outside our sample). For such firms, it is possible that the derivative of enforcement with respect to firm size may already be low, and so greater enforcement would increase the level of enforcement without necessarily flattening the slope.

2.4.4 Summing up

The transition to improved tax administration – characterized by higher staff-to-taxpayer ratios – led to substantially higher tax revenues. This came in the form of higher top-line revenues being reported by firms, rather than decreased deductions or changes in the degree to which taxes due were collected, consistent with the ideas laid out in Section 2.3. The increases in tax revenues for the government were more than two orders of magnitude larger than the increases in administrative costs associated with the increased enforcement. Surprisingly, the increased tax enforcement did not slow the rate of firm growth; if anything, the results suggest substantially higher revenue growth in the period after being switched to the MTO than that experienced by similar firms that did not move. We document that one reason why the MTOs may have been particularly successful is that they may have reduced the degree to which enforcement is size-dependent, at least for these firms, which may be an important finding for other countries considering such a tax regime shift.

2.5 Changes in Statutory Tax Rates

2.5.1 Empirical strategy

The second policy reform we study is the changes in Indonesia's corporate statutory tax rates in 2009 and 2010. We begin by using the differential tax change described in Section 2.2.1 to estimate the elasticity of taxable income (ETI) with respect to the net of tax rate. We then use this estimate to benchmark the impact of improved tax administration against more conventional changes in the statutory tax rate.

We follow the approach in Gruber and Saez (2002), Saez et al. (2012), and others. Specifically, since the marginal tax rate is a function of potentially endogenous variables (gross income, taxable income), we instrument for the change in a firm's marginal tax rate by taking the firm's characteristics (gross income, taxable income) from the tax year before the schedule change, and apply the new statutory tax schedule to these pre-period values.

Our estimating equation follows the standard panel-level specification discussed in Saez et al. (2012) in general, and in Gruber and Rauh (2007) in the corporate tax context, with the ETI estimated as the ε coefficient in:

$$\ln\left(\frac{z_{it+1}}{z_{it}}\right) = \alpha + \varepsilon \ln\left(\frac{1-\tau_{it+1}}{1-\tau_{it}}\right) + \varphi_1 \ln z_{it} + \varphi_2 \ln g_{it} + \delta_t + \delta_i + \nu_{it} \tag{2.20}$$

where z_{it} is taxpayer i's reported taxable income for tax year t, g_{it} is taxpayer i's reported gross income for tax year t, τ_{it} is taxpayer i's statutory marginal tax rate for tax year t, and ν_{it} is an error term. The ETI estimates are therefore with respect to the net of tax rate $1-\tau$ (the share of reported taxable income that the taxpayer gets to keep). Importantly, there were two tax changes (2009 and 2010), allowing the inclusion of taxpayer fixed effects (δ_i) in a regression specification that is already estimated in first-differences; we report robustness exercises that drop taxpayer fixed effects and/or only use a single tax change.

We instrument for the change in tax rates, $\ln\left(\frac{1-\tau_{it+1}}{1-\tau_{it}}\right)$, by computing the statutory marginal tax rate τ_{it} for taxpayer i in year t according to the statutory marginal tax rate schedules before and after the reform (described in Section 2.2.1 above), using taxpayer characteristics from the year prior to the reform.[46] We denote by τ^C_{it+1} and τ^C_{it} the marginal tax rate calculated using year $t+1$ and year t tax schedules applied to pre-period (i.e., 2008) values of g_{i2008} and z_{i2008}.

The first stage regression, therefore, is given by:

$$\ln\left(\frac{1-\tau_{it+1}}{1-\tau_{it}}\right) = \alpha + \omega \ln\left(\frac{1-\tau^C_{it+1}}{1-\tau^C_{it}}\right) + \theta_1 \ln z_{it} + \theta_2 \ln g_{it} + \delta_t + \delta_i + \nu_{it} \tag{2.21}$$

We estimate the first- and second-stage equations using corporate income tax filings for tax years 2008-2010, such that the ETI estimates leverage reform-induced changes in marginal tax rates over the two key years of the rate reform: the move from a taxable income-based to a gross income-based schedule in 2009, and the additional marginal tax rate cuts in 2010.

[46] The pre-reform marginal tax rates come directly from the schedule. As shown in Figure 9, the 2009 reform introduced a non-linear schedule to determine the total taxes due T_{it} of taxpayers with gross income between IDR 4.8 billion and IDR 50 billion, whereby a taxpayer with g IDR billion in gross income paid $\frac{r^*}{2}$ over a $(4.8/g)$ share of its taxable income, and r^* over the remaining amount:

$$T_{it} = \frac{r^*}{2}\left(\frac{4.8 \text{ billion}}{g_{it}}\right) z_{it} + r^*\left[1 - \left(\frac{4.8 \text{ billion}}{g_{it}}\right)\right] z_{it}$$

The marginal tax rate τ^{Post}_{it} for these taxpayers is therefore obtained by differentiating T_{it} with respect to z_{it}. We calculate the MTR for an additional dollar of taxable income z_{it} holding gross income g_{it} constant.

Following the standard practice in the literature, in our main specifications, we exclude taxpayers reporting zero taxable income in years 2008-2010 (and therefore undefined log taxable income).[47] We separately examine extensive margin effects (i.e., moving from 0 taxable income to positive taxable income).

Appendix Figure B.14 presents this reform-induced variation visually with a heat map of the change in predicted marginal tax rates (specifically, $\tau_{it+1}^{C} - \tau_{it}^{C}$) as a function of taxpayers' 2008 gross and taxable income, and indicates with a scatter plot where taxpayers fall along this variation. Panel A shows that the 2008-2009 schedule change induced a rich pattern of differential tax rate cuts (light green to blue areas) and differential tax rate increases (yellow to red areas), while the 2009-2010 schedule change induced differential but more tenuous tax rate cuts. Table B.14 presents alternative estimates of the ETI when only the 2008-2009 schedule change is used in estimation, and when we use lagged instruments as suggested by Weber (2014), among other specification robustness.

As the ETI estimates will be used to benchmark the tax administration effects, we use the same sample and balancing weights as in Section 2.4. In addition to the overall impacts, we also estimate ETIs separately for MTO and PTO taxpayers in order to assess the extent of differential responsiveness to tax rate changes under the different administration regimes. The fact that we are using the entropy-balancing weights implies that the difference in ETIs between MTO and non-MTO firms can be interpreted as the effect of being in the MTO on the firm's ETI, holding characteristics of the firm constant.

2.5.2 Results

First-stage

Table 2.6 presents the results. Panel A shows the first stage from estimating equation (2.21). Column (1) shows the results for all taxpayers. The first stage is quite strong – the coefficient of the actual marginal tax change on the predicted marginal tax change is 0.980, and the first-stage F-statistic is over 3,000. Columns (2) and (3) show that the first-stage is virtually identical for both MTO and non-MTO firms.

[47] Another reason that the literature typically excludes taxpayers with zero taxable income is that marginal tax rates (and therefore any variation in these rates) are based on positive taxable income thresholds (as was the case in Indonesia's pre-2009 ETI schedule). These papers typically also exclude taxpayers with small levels of taxable income altogether (e.g., Auten and Carroll, 1999; Gruber and Saez, 2002; Weber, 2014).

The elasticity of taxable income

The second-stage ETI estimates, from estimating equation (2.20), are shown in Panel B. Overall, for all firms, we estimate an elasticity of taxable income with respect to the net-of-tax rate of 0.59. This estimate is substantially larger than the estimate from Gruber and Rauh (2007) using Compustat data in the United States (0.2), but very close to the net-of-tax rate estimate from Dwenger and Steiner (2012) using a pseudo-panel of German corporate taxpayers' average tax rates (0.6). It is considerably smaller, however, than Bachas and Soto (2018)'s estimate from Costa Rica, which focuses on much smaller firms.[48][49]

Applying standard formulas, we can calculate the marginal excess burden of raising the top corporate income tax rate using this elasticity. We slightly modify the notation in Section 2.3 to account for the fact that we have a progressive tax schedule, and so we consider changes to the top marginal rate; derivations largely following Saez et al. (see 2012) and Keen and Slemrod (2017) are provided in Appendix B.4. The marginal excess burden of taxation is:

$$-\frac{dB}{dR} = \frac{\varepsilon\tau\rho}{1-\tau-\varepsilon\tau\rho} \tag{2.22}$$

where $\rho = \left(\frac{z^m}{z^m-\bar{z}}\right)$ is the Pareto parameter (which we calculate as 1.33 in our data).[50] This captures the additional loss to the taxpayer above and beyond the taxes paid, for each additional dollar of revenue raised. Our estimates imply that the marginal excess burden per dollar raised is 0.51; that is, each dollar of taxes raised causes an additional burden of 0.51 cents on taxpayers.

We can also return to the welfare framework above to use the estimated ETI to compute optimal marginal tax rates as a function of v, the marginal cost of public funds. Modifying

[48]The tax rate, and variation used, is somewhat different in these studies. Both Gruber and Rauh (2007) and Dwenger and Steiner (2012) estimate the elasticity with respect to the average effective tax rate, generating variation by changes in depreciation schedules and other treatments of capital expenditure, holding the statutory rate fixed. By contrast, our setting is unusual in that we have direct policy variation in statutory marginal rates that differs across firms. We, therefore, estimate the elasticity directly with respect to the statutory marginal rate.

[49]Appendix Table B.13 displays effects of the MTR reform on additional outcomes, showing that the reform had an effect on both the intensive (ETI) and extensive (reports any positive taxable income) margins. Our estimate for the extensive margin elasticity is 0.425 (0.069). The reform also had no effect on VAT payments, employment, or gross income.

[50]In Indonesia's pre-2009 system, with a progressive marginal tax system, this formula applies exactly, and one can calculate $a = \frac{z^m}{z^m-\bar{z}}$, where $\bar{z}$ is the taxable income threshold over which the top rate applies, and z^m is the average taxable income conditional on it being above $\bar{z}$. Our estimates here thus apply to the pre-2009 system. In the 2009-and-after system, this estimate is only approximate since a change in the marginal tax rate applies to everyone, but with discounts depending on gross income.

equation (2.9) to take into account the fact that we are considering a top marginal tax rate change, the top optimal tax rate is given by $\tau^* = \frac{1}{1+\rho\varepsilon\frac{v}{v-1}}$. When $v \to \infty$, this formula yields the revenue maximizing Laffer rate, $\tau^* = \frac{1}{1+\rho\varepsilon}$. Our estimates imply that the revenue-maximizing top rate is 56 percent in this context, substantially higher than the top 30 percent marginal tax rate observed throughout the period we study. We can reject that Indonesia is above the revenue-maximizing rate (p-value < 0.01). More generally, the 30 percent top rate observed in this period would be optimal if the marginal value of public funds $v = 1.5$, so any higher valuations would suggest that increasing the corporate tax rate is optimal. For example, a value of $v = 2$, so the social value of public funds is twice that of private funds (which could happen if public goods are underprovided in many developing countries), would yield an optimal top tax rate of 39 percent.

Robustness

Appendix Table B.14 shows that estimated ETI is robust to specification choices. In particular, we explore: unweighted estimates (column 2); estimates where balancing weights are re-estimated conditional on the sample of taxpayers with non-zero taxable income throughout 2007-2010 (column 3); estimates restricting the estimation to the sample of taxpayers that have positive taxable income for all years 2007-2010 (column 4); estimates using lagged data for instrument and controls and the same set of firms as in column 4 (column 5);[51] estimates without taxpayer fixed effects but including baseline controls (column 6); estimates with no baseline controls but with taxpayer fixed effects (column 7); and estimates using only the 2008-2009 change in reported income and tax rates (column 8). In the specifications in columns (6) and (8), where we exclude taxpayer fixed effects, we include sector fixed effects instead, since the tax change may differ systematically by sector. We also include a dummy for the firm's MTO status. Finally, in columns (9) and (10), we split the sample by those

[51] That is, applying the 2009 and 2010 schedules to 2007 — instead of 2008 — gross and taxable income data when constructing the marginal tax rate change instrument; and controlling for 2007 (instead of 2008) log taxable and log gross income for the 2008–2009 change, and for 2008 (instead of 2009) log taxable and log gross income) for the 2009-2010 change. As argued in Weber (2014), constructing the reform-induced marginal tax rate changes using lagged (rather than base-year) data addresses the possibility that ETI estimates might be inconsistently estimated (in particular, too small) due to mean-reversion in taxpayers' taxable income. As shown in column (4) of Appendix Table B.14, however, if anything this alternative specification produces a slightly *smaller*, although much less precise, ETI point estimate than our main specification, which is the opposite of the finding in Weber (2014). This suggests that either taxable income mean reversion is limited among the Indonesian firms in the analyzed period, or that the variation induced by Indonesia's marginal tax rate schedule reform is so heterogeneous across taxpayers (as seen in Appendix Figure B.14) that it is on average uncorrelated with transitory income shocks that induce mean-reversion in ETI estimates, providing more exogeneity in tax rate changes than typically observed in the literature.

taxpayers predicted to have a tax cut in 2008-2009 and those taxpayers predicted to have a tax increase in 2008-2009.

Most of these estimates are very similar. Note that the estimates without taxpayer fixed effects (columns 6 and 8) are somewhat larger – the ETI rises to 1.036 and 0.977, respectively. While these are higher, they still indicate that Indonesia is below the Laffer rate on taxes – even using the highest estimate across all our specifications (1.036), the revenue-maximizing tax rate is 42 percent.

Finally, we explore whether tax cuts or tax increases drive our findings. Columns (10) and (11) suggest that our results are largely driven by comparing taxpayers receiving a large tax cut in 2008-2009 with those receiving a smaller tax cut in the same years. For this sample, the estimated ETI is 0.625, almost identical to the full sample effect. Appendix Figure B.16 shows these results in event-study form graphically year-by-year, plotting the change in marginal tax rate (Panel A) and the impacts on taxable income (Panel B) for those predicted to have large vs. small tax cuts. The plotted regression coefficients are conditional on controls that mimic the specification in equation (2.20) (taxpayer fixed effects, and year dummies interacted with 2008 log gross income and 2008 log taxable income), and similarly weighted by MTO balancing weights. For those predicted to have tax increases, column (11) of Table B.14 shows that results are statistically imprecise, although the point estimated for the ETI is positive. The reason is that there is much less variation in the tax increase for this sample (over 90 percent of taxpayers experiencing an increase face an increase smaller than 4 percentage points), and the sample size is 60 percent smaller.

Complements or substitutes: Does improved tax administration affect the elasticity of taxable income?

We next investigate whether improved tax administration changes the sensitivity of taxable income to the tax rate. As discussed by Slemrod and Kopczuk (2002) and Keen and Slemrod (2017), the sign of the effect is ex-ante ambiguous. For example, improved tax administration may reduce the elasticity of taxable income by making concealment activities more costly. On the other hand, greater tax administration may also make firms more responsive to changes in the tax rate. For example, if firms pay only a share λ of their taxes owed (i.e., pay a tax rate $\lambda\tau$), then the elasticity with respect to the statutory tax rate τ would be higher as λ increases.

We can combine the two sources of variation to estimate this cross-elasticity. Specifically, we weight taxpayers by the weights developed in Section 2.4.1, so that we are analyzing firms moved to the MTO in 2007 with comparable control firms who were still serviced by regular tax offices. We then estimate equation (2.20) to calculate the elasticity of taxable income separately for the weighted sample of MTO and non-MTO firms, in order to estimate how improved tax administration affects this elasticity.

Columns (2) and (3) of Table 2.6 present the results. We find no statistically significant difference in the elasticity of taxable income for firms that have been moved to the MTO, compared to similar firms who remain in primary tax offices, though the point estimates suggest that the elasticity is smaller in firms moved to the MTO.[52]

2.5.3 Comparing changes in tax administration and tax rates

Suppose the government wants to raise additional tax revenue. Should it do so by raising tax rates, or improving tax administration? To investigate this, we focus on corporate income taxation in particular, and use our estimates of improved tax administration from Section 2.4 and our estimates of the elasticity of taxable income to shed some light on this question. First, we can calculate revenue neutral alternatives – that is, we can estimate how much the government would have had to increase the top corporate income marginal tax rate of 30 percent in 2007 in order to achieve the same additional revenue as the MTO tax office reorganization. Second, we can use these estimates, combined with the theory discussed above, to give conditions under which doing so by improving tax administration is likely to be welfare-improving relative to doing so by raising tax rates.

How much would tax rates have to rise to generate the MTO impact?

Recall that in Section 2.3.1, we derived equation (2.12), which gives the relationship between marginal tax rate changes and changes in administration holding revenue constant. The key parameters in equation (2.12) are $\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}$, the empirically estimated change in tax revenue (net of administration costs) from the introduction of the MTO estimated in Section 2.4, $\varepsilon_{1-\tau}$, the estimated elasticity of taxable income with respect to the net of tax rate estimated in Section 2.5.2, and τ, the marginal tax rate from which we are starting. To take this to the data, we modify this equation slightly to account for the fact that we have a progressive tax schedule, and therefore are considering changes to the top rates (see Appendix B.4). We can therefore calculate $\frac{d(1-\tau)}{d\alpha}|_R$ as a function of our estimates of $\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}$ from Section 2.4.2, $\varepsilon_{1-\tau}$ from Section 2.5.2, the Pareto parameter (ρ), and the

[52] Appendix Table B.15 shows that these results are robust to using actual MTO treatment status, and to whether these elasticities are estimated without MTO balancing weights.

marginal tax rate (τ). The results using this calculation are shown in Table 2.7. We provide the MTO estimate used in column (1), and provide estimates of the tax rate changes needed if applied to MTO firms only in column (2) and all taxpayers in column (3).[53] As shown in column (1), the tax changes needed to match the MTO effect are large. In particular, matching the tax administration effect on corporate income tax revenues could not have been accomplished by raising the marginal corporate income tax rate of MTO taxpayers in our analysis sample only while keeping that rate below the revenue-maximizing rate of 56% (column (2)). Alternatively, if the government were to tax all firms in the analysis sample (including those in PTOs), then matching the MTO effect of corporate income tax revenues would require raising the top marginal corporate income tax rate by 8 percentage points.[54]

It is worth emphasizing that these counterfactual tax increases would only replace the additional corporate income tax generated by the MTO. As shown in Table 2.1, corporate income taxes represent only about 15 percent of the additional tax revenue generated by the MTO. To generate the same amount of total income tax generated by the MTOs (i.e., including individual withholding and other taxes) would have required raising the corporate income tax on *all* taxpayers by 17 percentage points.[55]

Conditions for improving tax administration to be welfare-improving, relative to raising tax rates

The theoretical framework also suggests a related calculation to assess whether raising revenue through improved tax administration is welfare-improving on the margin relative to raising revenue through higher tax rates. Recall that equation (2.14) gives the welfare change on the margin from shifting to increased tax administration and reducing marginal tax rates, holding government revenue constant. Modifying this equation to account for the

[53]The MTO estimates in Table 1 were in real terms (2007 rupiah). However, since the tax changes are in nominal terms, we provide the MTO effect in nominal terms.

[54]Appendix Table B.14 presents robustness checks on this calculation corresponding to the various robustness checks on the ETI estimate described above. In addition, as an additional bounding exercise, we consider what would happen if in fact the ETI was zero; that is, if there was no behavioral response whatsoever. In this case, we calculate that to match the MTO CIT revenue effect, the 2006 top marginal income tax rate would have had to be raised by 6 percentage points on all taxpayers (as opposed to by 8 percentage points at our estimated ETI of 0.590).

[55]Appendix Table B.16 presents alternative counterfactual tax rate increases based on extrapolating the MTO effect and the tax base to all taxpayers in the 19 regions. Since the extrapolated MTO effect is likely a lower bound (that is, it scales linearly with the number of MTO taxpayers rather than proportionally with their size, whereas the income subject to the marginal tax rate — $N\left(z^m - \bar{z}\right)$ in equation B.9 — increases proportionally with taxpayer size), the extrapolated counterfactual tax rate changes are also lower bounds.

fact that the tax increase applies only to the top bracket yields:

$$dW = \left(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}\right)\frac{1}{1-\frac{\tau}{1-\tau}\rho\varepsilon_{1-\tau}} - \frac{d\gamma}{d\alpha} \tag{2.23}$$

The first term, $\left(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}\right)\frac{1}{1-\frac{\tau}{1-\tau}\rho\varepsilon_{1-\tau}}$, is essentially the change in the tax rate given in equation (2.12) multiplied by average taxable incomes z, which we can estimate (see previous section). We do not, however, observe $\frac{d\gamma}{d\alpha}$, the change in a firm's private compliance costs associated with the MTO.

Nevertheless, there are several reasons to think that, in our context, equation (2.23) is positive, which implies that the welfare implications from using improved tax administration to raise more revenue on the margin, rather than higher tax rates, would be positive. First, applying our estimates from Section (2.5.2), the $\frac{1}{1-\frac{\tau}{1-\tau}\rho\varepsilon_{1-\tau}}$ term is 1.51 in our context. This term is the marginal efficiency cost of funds, equal to 1 + the excess burden calculated in equation (2.22). This term captures how much more efficient it is to raise funds via tax administration rather than via tax rates, in terms of lost deadweight-costs of taxation (other than the private costs of compliance $\frac{d\gamma}{d\alpha}$). The fact that $\frac{1}{1-\frac{\tau}{1-\tau}\rho\varepsilon_{1-\tau}}$ is 1.51 implies that equation (2.23) would be positive even if revenue gains from improved administration were only 63 percent of additional compliance costs. Second, the fact that the net revenue effect of the MTO, $\left(\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}\right)$, is so large – two orders of magnitude larger than what it costs the government to administer it $\left(\frac{da}{d\alpha}\right)$ – suggests that it may also be large relative to the *change* in compliance costs associated with the intervention.

Third, the intervention we study was actually an attempt to *reduce* compliance costs, not increase them, by improving customer service for taxpayers (e.g., answering questions, etc). As described in Section 2.2.1, anecdotal evidence from an ACNielsen survey of firms finds higher "satisfaction" with tax office interactions at MTOs compared to PTOs. One might imagine, then, that the MTO intervention raised the marginal costs of evasion while at the same time lowering the *level* of compliance costs. In such a case, the net change in firm compliance costs could be negative even if the marginal cost of evasion increased.

2.6 Conclusion

There is often a debate on whether to invest limited funds in improving tax administration, and how the returns from doing so differ from other policy levers such as changes to the tax rate. To study this, we estimate the impacts of two nationwide reforms in Indonesia—a cheap but expansive administration reform that differentially affected medium sized firms, and a change in corporate tax rates. We find that increasing the intensity of tax administration by moving the top firms in each region into special "Medium-Sized Taxpayer Offices," with similar structures and procedures, but much higher staff-to-taxpayer ratios, more than *doubled* tax revenue from affected firms. While there are concerns that new reforms may initially have impacts, but then fade over time as firms learn to evade, we actually find the opposite: impacts increase over the subsequent six years.

We find that one reason why these MTOs may have been so successful is that it flattened the relationship between enforcement and firm size, suggesting that governments that are designing tax administration reforms should be concerned not only with the level of enforcement, but also how the enforcement level changes as firms evolve. This finding suggests that differential tax enforcement on larger firms, which could be optimal for a tax authority facing limited resources and trying to maximize its tax intake in a static sense, may also contribute to the large number of very small firms in developing countries (Hsieh and Olken, 2014).

While this was a large-scale reform, its costs as a fraction of increased revenue were minuscule — about 1.5 percent — implying that this investment had a considerable overall return. In fact, the increase in tax rates needed to achieve a similarly sized effect would be quite large. Using non-linear changes to the corporate income tax schedule, we estimate an elasticity of taxable income of 0.59. Using this ETI to compare the two approaches, we calculate that the increased revenue from MTO taxpayers due to improvements in tax administration is equivalent to raising the marginal corporate tax rate on *all* firms by about 8 percentage points. Given these estimates, improved tax administration is likely to be the preferred approach unless the compliance costs imposed on taxpayers are extremely high. These results may also help explain why so many developing countries have been moving the largest taxpayers into separate offices with more intensive tax administration, such as the ones we study here, and more generally, why many developing country governments are increasingly investing in improved tax administration.

Figure 9: Change in corporate income tax schedule

Panel A: Tax schedule prior to 2009
(Marginal tax rates vs. taxable income)

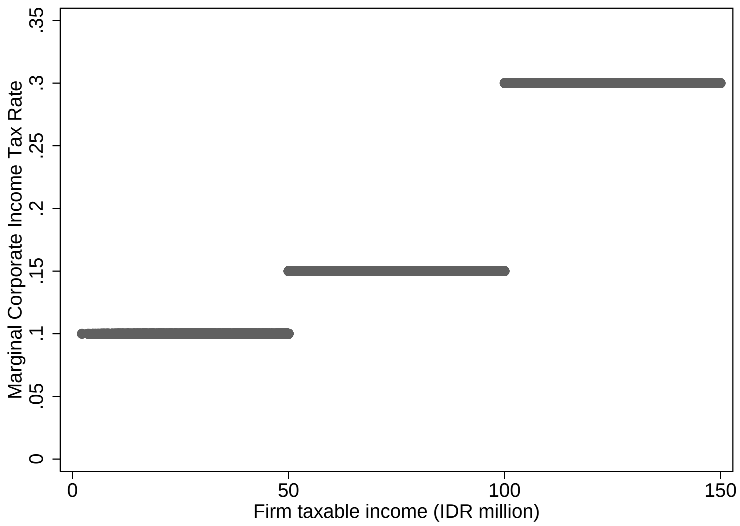

Panel B: Tax schedule 2009 and later
(Marginal tax rates vs. gross income)

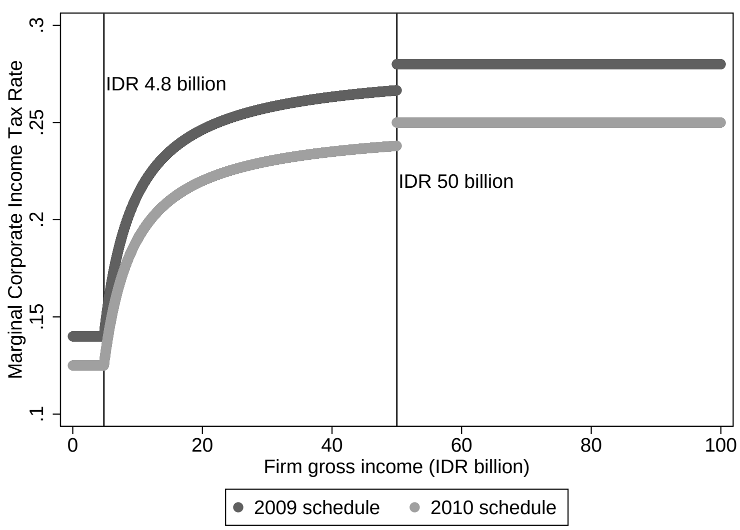

Notes: This figure shows corporate income tax rate schedules from before (Panel A) and after (Panel B) Indonesia's 2009 corporate income tax rate reform. Pre-reform rates were based on taxable income cutoffs. Post-reform rates were based on gross income cutoffs. In both periods, corporate income tax rates were applied to taxable income.

Figure 10: Example of an increase in enforcement α on reported and real output

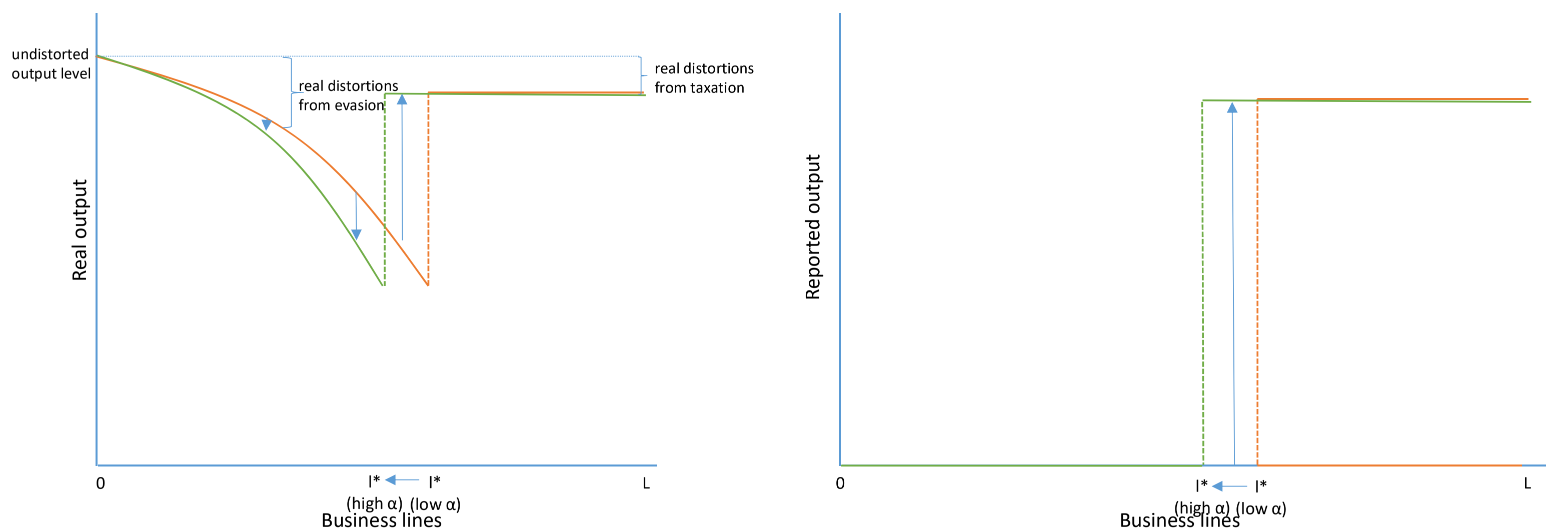

Notes: See Section 2.3.2.

Figure 11: MTO effect on total taxes paid

Panel A: MTO vs. non-MTO weighted annual averages

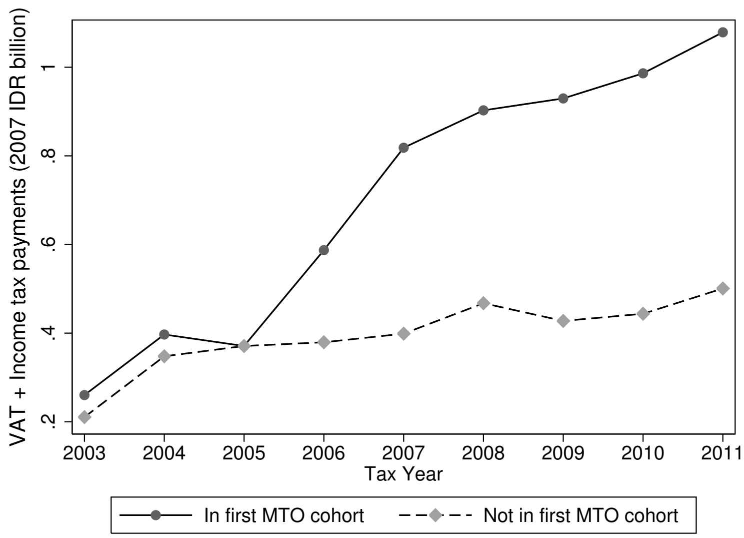

Panel B: Year-by-year estimates

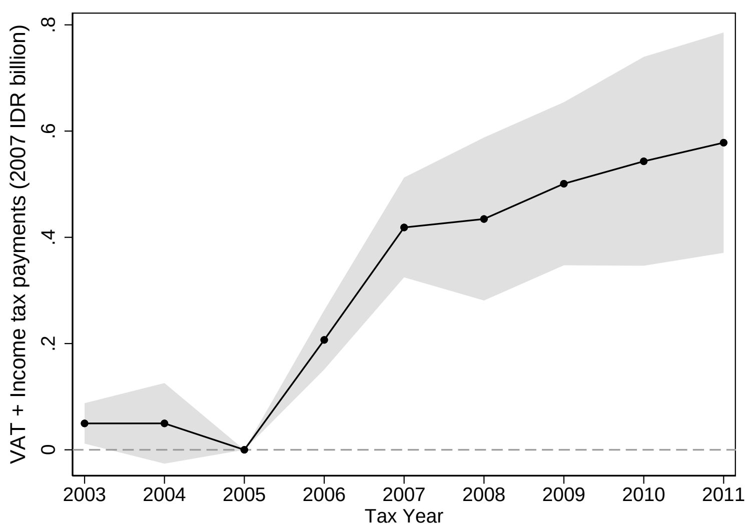

Notes: This figure shows annual weighted averages by MTO 2007 assignment group (Panel A) and year-by-year weighted regression estimates of the effect of MTO 2007 assignment on total taxes paid (Panel B). Regression coefficients are year-by-year reduced form effects of MTO treatment, and are estimated by interacting the MTO assignment dummy variable M_{iFC} in equation (2.16) with year dummies, while omitting the interaction and main effect dummies for base year 2005. The weights used in both panels are taxpayer-specific, fixed across all analyses, and constructed by applying Hainmueller (2012)'s entropy-balancing methodology to the MTO assignment formula inputs (gross income and total taxes paid) for tax year 2005. Taxpayer-level total taxes paid data are from the Treasury, and include payments from all branches of the same corporate entities. IDR values are deflated to 2007 IDR using Indonesia's GDP deflator. Solid lines are point estimates; shaded areas are 95% confidence intervals based on standard errors clustered at the taxpayer level.

Figure 12: MTO effect on reported income

Panel A: Gross income

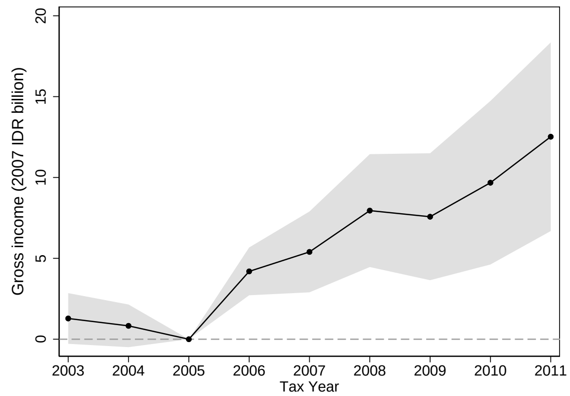

Panel B: Taxable income

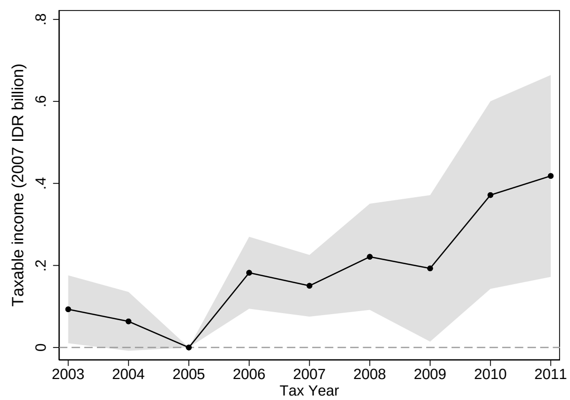

Panel C: Corporate income taxes due

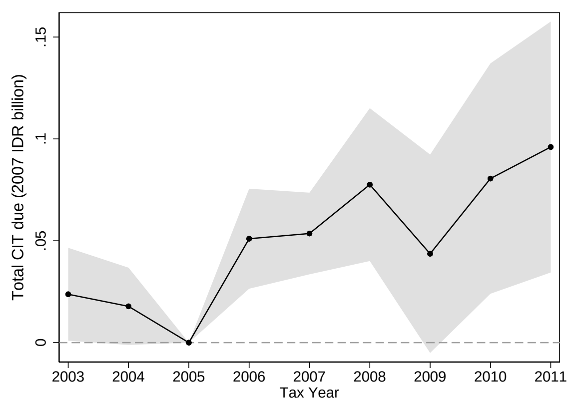

Notes: See notes to Figure 11. Reported income data are from tax filing form SPT 1771 (annual corporate income tax return), and are reported by the taxpayer headquarters on behalf of all branches of the same corporate entity.

Figure 13: Audit and assessment as a function of total taxes paid and permanent workers

Panel A: As a function of total taxes paid

Panel B: As a function of permanent workers

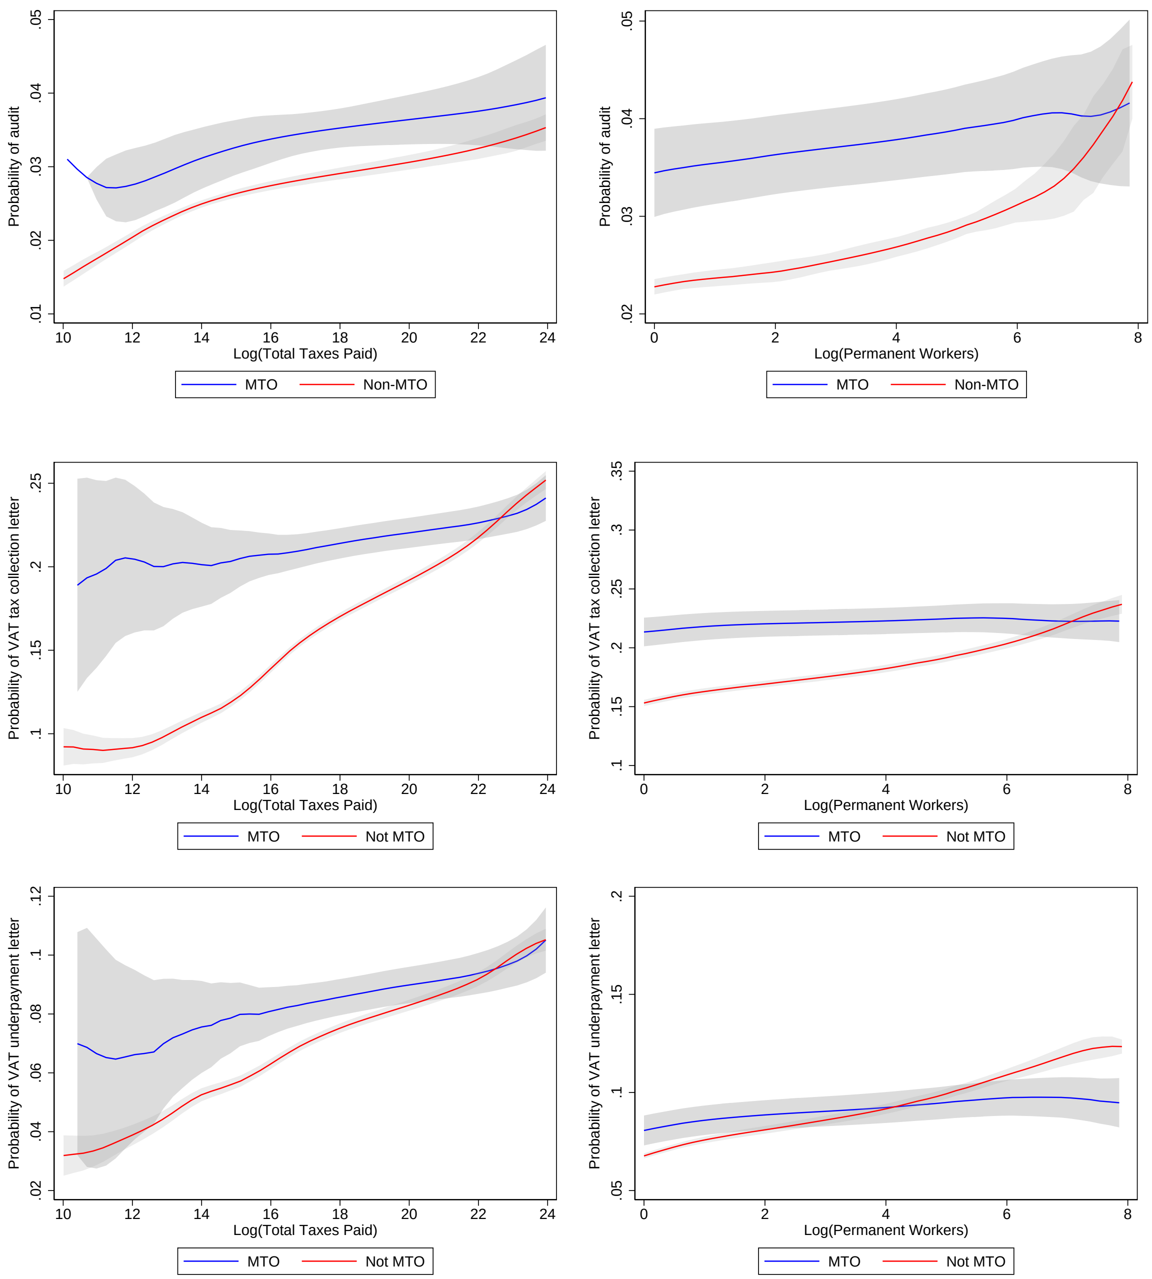

Notes: This figure shows estimates of the probability of audit and VAT tax assessment (receipt of tax collection letter or underpayment letter) as a function of taxpayer log total taxes paid (left, Panel A) and log permanent workers (right, Panel B). Shaded areas indicate 95% confidence intervals. Panels A and B show local linear regression estimates using an Epanechnikov Kernel of bandwidths 4 and 2, respectively. All plots are based on weighted post-MTO assignment data. Probability of audit is based on 2009-2011 audit data. Probability of VAT collection letter and of VAT underpayment letter are based on 2006-2011 tax assessment letters data. Firm employment data are from corporate employment tax withholding form SPT 1721.

Table 2.1: MTO treatment effect on tax payments, reported income, and tax collection rate

	Weighted means				MTO treatment effect		
	Pre-treatment			Treated post-treatment counterfactual	Reduced Form	IV	IV as % of post-treatment counterfactual
	Untreated	Treated	N				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Tax Payments (2007 IDR billion)							
VAT	0.26	0.26	163,572	0.27	0.240	0.371	137%
					(0.050)	(0.078)	
Corporate Income Tax	0.05	0.06	163,572	0.07	0.048	0.074	111%
					(0.009)	(0.014)	
Other income taxes	0.06	0.06	163,572	0.07	0.052	0.080	113%
					(0.011)	(0.017)	
Total	0.37	0.37	163,572	0.41	0.340	0.525	128%
					(0.062)	(0.096)	
Panel B: Reported Income (2007 IDR billion)							
Gross income	13.03	13.03	136,601	12.04	5.754	9.131	76%
					(1.375)	(2.181)	
Taxable income	0.39	0.46	137,585	0.50	0.150	0.238	47%
					(0.045)	(0.072)	
Corporate Income Tax due	0.09	0.12	137,586	0.13	0.041	0.065	51%
					(0.012)	(0.020)	
Profit margin (net income/ gross income)	0.06	0.07	110,492	0.07	0.001	0.001	--
					(0.002)	(0.003)	
Panel C: Tax Collection Rate							
CIT paid/ CIT due	0.92	0.67	113,480	0.83	0.054	0.088	--
					(0.131)	(0.214)	

Notes: This table presents estimates of the MTO treatment effect on tax payments, reported income, and Corporate Income Tax (CIT) collection rate. Columns (1)-(2) show pre-treatment (specifically, tax year 2005) weighted means for untreated and treated taxpayers, respectively. Column (3) shows number of observations in each regression. Column (4) shows post-treatment weighted means for the treated group absent treatment (that is, counterfactual means), and is computed by subtracting the MTO IV treatment effect in Column (6) from the treated group's realized post-treatment weighted mean. Column (5) presents estimates of the effect of being assigned to MTO in 2007 (that is, the reduced form) according to equation (2.16), while column (6) presents the IV estimates of MTO treatment as specified in equation (2.17). Column (7) benchmarks the IV effects in column (6) as a percentage of counterfactual means in column (4). Means in columns (1), (2), and (4) and estimates in columns (5)-(6) are all weighted by the same taxpayer-specific balancing weights. Weights are constructed by applying Hainmueller (2012)'s entropy-balancing methodology to the MTO assignment formula inputs (gross income and total taxes paid) for tax year 2005. Tax payments data are from the Treasury and include payments from all branches of the same corporate entity. Reported income data are from tax filing form SPT 1771 and are reported by the taxpayer headquarters on behalf of all branches of the same corporate entity. IDR values are deflated to 2007 IDR using Indonesia's GDP deflator. Standard errors are clustered at the taxpayer level.

Table 2.2: MTO treatment effect on reported employment

	Weighted means				MTO treatment effect		
	Pre-treatment						
	Untreated	Treated	N	Treated post-treatment counterfactual	Reduced Form	IV	IV as % of post-treatment counterfactual
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Total workers	93.31	167.37	117,049	162.53	6.960	12.646	--
					(12.032)	(21.865)	
Permanent workers	36.36	43.80	117,049	49.19	5.705	10.365	21%
					(3.309)	(6.009)	
Temporary workers	56.95	123.57	117,049	113.34	1.256	2.281	--
					(11.650)	(21.168)	
Total wage bill (2007 IDR billion)	1.11	1.33	117,049	1.37	0.182	0.330	24%
					(0.077)	(0.139)	
Permanent workers	0.69	0.81	117,049	0.92	0.106	0.193	21%
					(0.055)	(0.100)	
Temporary workers	0.41	0.52	117,049	0.44	0.075	0.136	--
					(0.053)	(0.097)	
Average yearly wage (2007 IDR million)	16.26	16.20	117,049	14.99	1.286	2.337	16%
					(0.553)	(1.002)	

Notes: See notes to Table 2.1. Firm employment and wage data are from corporate employment tax withholding form SPT 1721, and exclude tax year 2008, for which data are not available. Average yearly wage is computed as total wage bill divided by total workers, and is not reported separately for permanent vs. temporary workers as many firms have zero temporary workers. See Data Appendix for details.

Table 2.3: Impacts of MTO on corporate income tax corrections and VAT underpayment letters

	Weighted means				MTO treatment effect		
	Pre-treatment						
	Untreated	Treated	N	Treated post-treatment counterfactual	Reduced Form	IV	IV as % of post-treatment counterfactual
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Corporate Income Tax Corrections							
Filed any corrections	0.13	0.06	163,572	0.07	0.076	0.118	177%
					(0.008)	(0.012)	
Corrected this tax year's figures	0.21	0.36	163,572	0.24	-0.052	-0.080	-33%
					(0.012)	(0.019)	
Panel B: VAT tax assessment letters							
Tax collection letter	0.21	0.25	163,572	0.22	-0.004	-0.007	--
					(0.012)	(0.018)	
Underpayment letter	0.12	0.12	163,572	0.08	0.001	0.001	--
					(0.009)	(0.014)	

Notes: See notes to Table 2.1. This table presents estimates of the MTO treatment effect on tax filing corrections and VAT tax assessments.

Table 2.4: Enforcement, firm size, and the MTO: cross-sectional evidence

	Outcome		
	Audited	Received VAT Collection Letter	Received VAT Underpayment Letter
	(1)	(2)	(3)
Panel A: Measuring firm size as total taxes paid			
Assigned to MTO in 2007	-0.002	0.001	0.000
	(0.009)	(0.009)	(0.007)
Ln(Total Taxes Paid)	0.012	0.027	0.011
	(0.002)	(0.002)	(0.001)
Ln(Total Taxes Paid) x Assigned to MTO in 2007	-0.008	-0.016	-0.003
	(0.003)	(0.003)	(0.002)
N	52,763	111,940	111,940
Panel B: Measuring firm size as permanent workers			
Assigned to MTO in 2007	0.054	0.106	0.042
	(0.016)	(0.016)	(0.011)
Ln(Permanent Workers)	0.014	0.028	0.023
	(0.005)	(0.004)	(0.003)
Ln(Permanent Workers) x Assigned to MTO in 2007	-0.014	-0.022	-0.013
	(0.006)	(0.006)	(0.004)
N	42,792	73,043	73,043
Panel C: Measuring firm size as total workers			
Assigned to MTO in 2007	0.037	0.115	0.034
	(0.016)	(0.016)	(0.011)
Ln(Total Workers)	0.013	0.024	0.021
	(0.004)	(0.003)	(0.003)
Ln(Total Workers) x Assigned to MTO in 2007	-0.008	-0.021	-0.009
	(0.005)	(0.005)	(0.004)
N	43,202	74,125	74,125
Years	2009-2011	2006-2011	2006-2011
Year FE	Yes	Yes	Yes
Firm FE	No	No	No

Notes: This table presents cross-sectional regression estimates of the effect of MTO 2007 assignment on the slope of several measures of enforcement as a function of several measures of taxpayer size. Regression coefficients for alternative measures of enforcement are presented in columns (1)-(3). Regressions are separately estimated in Panels A through C given alternative measures of taxpayer size, and including the regressors listed on the left-most column of each panel. All regressions are weighted by the same taxpayer-specific balancing weights as in the MTO treatment effect and ETI estimation analyses. Standard errors are heteroskedasticity-robust.

Table 2.5: Enforcement, firm size, and the MTO: difference-in-differences estimates

	Outcome	
	Received VAT Collection Letter	Received VAT Underpayment Letter
	(1)	(2)
Panel A: Measuring firm size as total taxes paid		
Assigned to MTO in 2007 x (Year>2005)	-0.041	-0.022
	(0.016)	(0.012)
Ln(Total Taxes Paid)	0.016	0.003
	(0.003)	(0.002)
Ln(Total Taxes Paid) x Assigned to MTO in 2007	0.008	0.010
	(0.005)	(0.003)
Ln(Total Taxes Paid) x Assigned to MTO in 2007 x (Year>2005)	-0.018	-0.011
	(0.005)	(0.004)
N	168,541	168,541
Panel B: Measuring firm size as permanent workers		
Assigned to MTO in 2007 x (Year>2005)	0.069	0.067
	(0.024)	(0.016)
Ln(Permanent Workers)	0.042	0.020
	(0.013)	(0.014)
Ln(Permanent Workers) x Assigned to MTO in 2007	-0.004	0.006
	(0.016)	(0.015)
Ln(Permanent Workers) x Assigned to MTO in 2007 x (Year>2005)	-0.026	-0.022
	(0.007)	(0.005)
N	126,417	126,417
Panel C: Measuring firm size as total workers		
Assigned to MTO in 2007 x (Year>2005)	0.068	0.056
	(0.026)	(0.018)
Ln(Total Workers)	0.019	0.008
	(0.006)	(0.005)
Ln(Total Workers) x Assigned to MTO in 2007	0.001	0.003
	(0.009)	(0.007)
Ln(Total Workers) x Assigned to MTO in 2007 x (Year>2005)	-0.020	-0.014
	(0.006)	(0.005)
N	128,553	128,553
Years	2003-2011	2003-2011
Firm FE	Yes	Yes
Year FE	Yes	Yes

Notes: This table presents taxpayer-level difference-in-differences regression estimates of the effect of 2007 MTO assignment on the slope of several measures of enforcement as a function of several measures of taxpayer size. Regression coefficients for alternative measures of enforcement are presented in columns (1)-(2). Regressions are separately estimated in Panels A through C given alternative measures of taxpayer size, and including the regressors listed on the left-most column of each panel. All regressions are weighted by the same taxpayer-specific balancing weights as in the MTO treatment effect and ETI estimation analyses. Standard errors are clustered at the taxpayer level.

Table 2.6: Estimated elasticity of taxable income w.r.t. the net of tax rate

	Instrument: Reform-induced change in marginal tax rate		
		Separate by MTO status	
	All taxpayers	MTO	Not MTO
	(1)	(2)	(3)
Panel A: First Stage			
Endogenous:	0.980	0.981	0.982
Δ Ln(Net-of-tax rate)	(0.010)	(0.018)	(0.010)
F-statistic	3,629.32	1,112.23	3,250.73
N	12,790	964	11,826
Panel B: IV (ETI estimates)			
Outcome:	0.590	0.348	0.779
Δ Ln(Taxable Income)	(0.198)	(0.379)	(0.216)
P-value of difference		0.322	
Taxpayer FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes

Notes: This table presents instrumental variable (IV) estimates of the corporate Elasticity of Taxable Income (ETI) based on Indonesia's 2009 corporate income tax schedule reform and 2010 marginal tax rate cut. Panel A presents first stage effects of the reform-induced predicted change in marginal tax rates on realized marginal tax rates according to equation (2.21). Realized marginal tax rates are computed according to the schedule rules described in Section 2.2.1. Panel B presents IV estimates of the effect of log marginal net of tax rates on log taxable income (that is, ETI estimates) according to equation (2.20). The estimation sample is composed of the same taxpayers as in the MTO treatment effect analyses, and consists of data for the years immediately surrounding the reform (2008-2010). All regressions are weighted by the same taxpayer-specific weights as in the MTO treatment effect analyses. In addition to taxpayer and year fixed effects, all regressions control for base year log taxable income and base year log gross income. The p-value of the test for difference between the MTO and Non-MTO ETIs is shown between columns (2) and (3). Standard errors are clustered at the taxpayer level.

Table 2.7: Counterfactual CIT income tax increases to match MTO effects

		MTR raise needed to generate MTO effect on total revenue	
	MTO IV treatment effect (IDR billion)	Taxing MTO taxpayers	Taxing all taxpayers
	(1)	(2)	(3)
Corporate Income Tax	0.091	xx	8 pp
Total Income Taxes	0.180	xx	17 pp

Notes: This table presents estimates of by how much Indonesia would have had to raise its 2006 top marginal corporate income tax (CIT) rate of 30 percent in order to generate the same total revenue gains as the MTO effect, following calculations described in the text. Counterfactuals in columns (2)-(3) are computed by plugging the MTO treatment effect in column (1) and the ETI estimate of 0.590 from Table 2.6 into equation (2.7). Counterfactuals are displayed as "xx" whenever it is not possible to raise the respective amount of tax revenues without exceeding the revenue-maximizing tax rate of 56 percent. Because taxpayers' behavioral response to marginal tax rate increases (and therefore the ETI) are with respect to nominal (not real or IDR-deflated) values, the MTO treatment effect used for the counterfactual and displayed in column (1) is MTO effect on nominal (not real or IDR-deflated) corporate income taxes and total income taxes. The remaining non-schedule inputs of equation (2.7) are computed from the taxpayer-level data depending on which sample of taxpayers is assumed to received the counterfactual tax rate increase. Column (2) assumes only MTO taxpayers would be taxed; while column (3) assumes all taxpayers in the analysis sample would be taxed. All estimates are conditional on taxpayers who filed SPT 1771 (Corporate Income Tax return form), from where the non-schedule inputs of equation (2.7) are computed.

Chapter 3

Charter Schools and Suspensions: Evidence from Massachusetts Chapter 222

Abstract

I evaluate the impact of Massachusetts Chapter 222—a policy that limited charter schools' ability to suspend students—on student suspensions and test scores. Comparing charter attendance effects before vs. after Chapter 222, I find that Chapter 222 reduced charter suspensions by roughly 10 percentage points, but had no impact on charter learning. I then use variation in lottery offers and applicants' pre-lottery suspensions to separate the effect of suspensions from that of charter attendance on test scores. Suspensions appear to be unrelated to achievement in charters, while the causal effect of charter attendance on test scores is large and positive.[1]

[1]I am indebted to Josh Angrist and Parag Pathak for invaluable guidance, and to Alberto Abadie, Chris Ackerman, Isaiah Andrews, Ben Olken, Sarah Cohodes, Esther Duflo, Helen Ho, Peter Hull, David Martin, Elizabeth Setren, Camille Terrier, and participants of the MIT Economics Labor Lunch and Development Tea for valuable feedback. Anran Li provided expert research assistance. School Effectiveness and Inequality Initiative (SEII) Assistant Director Eryn Heying provided invaluable administrative support. I am especially grateful to Carrie Conaway, Matt Deninger, Alison Bagg, Pierre Lucien, and staff from the Massachusetts Department of Elementary and Secondary Education for providing access to administrative student data.

3.1 Introduction

Student suspensions in charter schools are common yet controversial. In Massachusetts, the debate over school suspensions has centered around urban charter schools, which increase students' test scores but suspend more often than traditional public schools ("TPS") (Angrist et al., 2013). Suspensions are most prevalent in grades 5–8, most suspended students are Black or Hispanic, and suspended students are typically removed from school for a day.[2] Yet there is no evidence on whether charter suspensions harm, improve, or have no effect on student learning.

This paper leverages Massachusetts Chapter 222, a policy that limited charters' use of suspensions, to estimate the effect of suspensions on charter students' test scores. Chapter 222 was signed in August 2012 and took effect in school year 2015.[3] Under the policy, principals are required to take several steps before suspending or expelling a student, such as sending written notifications to parents and meeting with parents to discuss the circumstances that led to the suspension. Schools must also ensure that students who were excluded from school for disciplinary reasons can make academic progress during the classroom removal period, a requirement that previously applied only to students with special needs.

I analyze the effect of Chapter 222 using a Difference-in-Differences Instrumental Variables empirical strategy. Specifically, I compare the outcomes for charter vs. TPS students before and after Chapter 222 in a sample of Boston charter middle school applicants, where charter attendance is randomly assigned via lottery. I find that by the end of school year 2017 Chapter 222 reduced the causal effect of charter attendance on suspensions by 10 percentage points, nearly halving the pre-Chapter 222 gap in suspensions between charters and TPS. In contrast, the policy had no sizable or statistically significant effect on charter math test scores.

To understand how Chapter 222 reduced suspensions without affecting test scores, I use variation in lottery offers from charters of varying disciplinary environments, and heterogeneity in applicants' pre-lottery suspensions, to separately identify the effects on test scores of charter suspensions vs. those of charter attendance. Consistent with the observed impact of Chapter 222, the causal effect of suspensions on charter students' test scores is zero. Conversely, the causal effect of charter attendance on suspended students' test scores is positive, large, and similar to the effect on non-suspended students.

[2]See Appendix Figures C.1 and C.2 for a breakdown of suspension rates.

[3]Throughout, I refer to school years after their spring semester year (e.g., school year 2015 refers to Fall 2014 and Spring 2015).

My findings contribute to a large literature on the effects of charter attendance. Many lottery-based studies have documented large positive effects of charter attendance on test scores—see, for example, Hoxby and Murarka (2009); Dobbie and Fryer Jr (2011); Angrist et al. (2010, 2012, 2016); Abdulkadiroğlu et al. (2011); Setren (2017)—with the largest test score gains in this literature come from *No Excuses* charters (Chabrier et al., 2016). *No Excuses* charters are characterized by an "emphasis on discipline, school uniforms, cold-calling, strict adherence to school-wide standards, and the use of Teach For America alumni" (Angrist et al., 2013).[4] And while a large body of evidence shows that Boston's *No excuses* charters significantly improve learning (Angrist et al., 2016; Abdulkadiroğlu et al., 2016, 2011), they also suspend more. As OLS effects of suspensions on test scores are negative, this raises the question of whether *No excuses* Boston charters' success is because of, in spite of, or unrelated to high suspension rates. This paper addresses this knowledge gap.

3.2 Empirical strategy

I estimate the effect of Chapter 222 on student outcomes using a Difference-in-Differences Instrumental Variables ("DD-IV") approach. This approach compares the outcomes of charter to TPS students, before and after Chapter 222 took effect, in a sample of charter school applicants—where charter attendance is randomly assigned via lottery. The effect of Chapter 222 on student outcome Y_{it} is coefficient γ in the following second stage regression:

$$Y_{it} = \quad \alpha + \beta D_{it} + \gamma \left[D_{it} \times 1_{\{t>t^*\}}\right] + \zeta' X_i + \delta_t + \delta_{g(i,t)} + \epsilon_{it}, \tag{3.1}$$

where D_{it} is a dummy indicating whether charter applicant i was enrolled in a charter in school year t;[5] $1_{\{t>t^*\}}$ is a dummy indicating whether year t is after Chapter 222's effective year; δ_t and $\delta_{g(i,t)}$ are year and grade fixed effects; and X_i is a vector of applicant-level demographics and baseline grade covariates, including a fixed effect for application year and a fixed effect for the set of charter schools to which i applied (the applicant's "risk set"). Conditioning on risk sets is necessary because the probability of winning any charter lottery depends on the set of charters to which the applicant applies; controlling for student baseline covariates X_i reduces the variance of point estimates.

Since charter attendance is itself a treatment, γ can also be interpreted as Chapter 222's

[4]More specifically, these five variables are most predictive of a school self-identifying as *No Excuses*.

[5]As in Angrist et al. (2013), I define a student to be enrolled in a charter for the whole school year even if the student only attended the charter for a single day in that year.

impact on the charter attendance treatment effect. The first stage regressions for charter attendance before and after Chapter 222 are

$$D_{it} = \theta + \iota Z_i + \kappa \left[Z_{it} \times 1_{\{t>t^*\}} \right] + \nu' X_i + \lambda_t + \lambda_{g(i,t)} + \mu_{it}, \tag{3.2}$$

$$D_{it} \times 1_{\{t>t^*\}} = \xi + \pi Z_i + \rho \left[Z_{it} \times 1_{\{t>t^*\}} \right] + \varphi' X_i + o_t + o_{g(i,t)} + \upsilon_{it}, \tag{3.3}$$

where Z_i is a dummy for whether applicant i received a lottery offer from any charter; and λ and o coefficients are the same set of fixed effects as in Equation 3.1. The key assumption for a causal interpretation of γ is that potential outcomes of charter vs. TPS students would have followed parallel trends but for Chapter 222, or equivalently, that the charter attendance *effect* would have remained constant after Chapter 222 had the policy not taken effect.

3.3 Data and results

I implement the DD-IV empirical strategy from Section 3.2 by linking administrative data on student enrollment, demographics, test scores, and disciplinary records to the list of Boston charter middle school lottery applicants for cohort years 2005–2014. These students were in middle school grades between 2006 and 2017. Appendix Table C.11 lists the charter schools and cohort years in the sample. The Data Appendix describes each of the data sources and linking procedures, which follow Setren (2017) and Angrist et al. (2016, 2013).

In order for the exercise's results to be interpreted as the casual impact of Chapter 222, the potential outcomes of treated vs. untreated students must follow parallel trends. Figure 14 presents a visual check that parallel trends in student outcomes does indeed hold prior to Chapter 222. It plots year-by-year estimates of charter attendance effects—relative to school year 2012—on a dummy for whether a student is ever suspended (in-school or out-of school). Despite year-to-year variation, the charter attendance effect prior to Chapter 222's was not statistically different from the baseline year 2012, displaying no pre-trends. However, the charter attendance effect on suspensions starts to decline in 2013, the first school year following Chapter 222's signing. On levels, Appendix Table C.4 shows that up to 2012 applicants who attended charters by virtue of winning the lottery were on average 22 percentage points more likely to be suspended out-of-school than lottery-losing counterparts attending TPS. By 2017 I cannot reject that the charter attendance effect on the probability of a suspension was zero (or, alternatively, 22 percentage points lower than the 2012 estimate), suggesting that Chapter 222 closed the charter vs. TPS gap in out-of-school suspension probability within five years of its signing.

Figure 15 replicates the exercise in Figure 14 for math test scores, also showing no pre-trends. However, in contrast to the effect of Chapter 222 on charter suspensions, Figure 15 shows that that Chapter 222 had no statistically significant impact on math test scores at charters. Instead, as shown in Appendix Table C.4, the charter attendance effect on math remained steadily large and positive at 0.566 standard deviations throughout the period.[6]

Table 3.1 presents a formal quantification of Chapter 222's effect. Columns (1) and (2) report estimates of β and γ, respectively, from Equation 3.1 using Chapter 222's signature year as the key policy year ($t^* = 2012$). Column (3) replicates this exercise using the year Chapter 222 took effect as the key policy year instead ($t^* = 2014$).[7] These charter attendance effects are consistent with lottery-based charter attendance effects reported elsewhere in the literature. Attending a charter school by virtue of winning the lottery caused a 20.5% increase in a student's probability of suspension (out-of- or in-school) at any point in their middle school years relative to lottery applicants who attended TPS. Lottery-induced charter attendees also experienced substantial average gains in test scores: 0.591 standard deviations for math and 0.348 standard deviations for English.[8]

Overall, Chapter 222 reduced the probability of being suspended at a charter school by 9.5 percentage points, with no statistically significant effect on math test scores. Table 3.1 also shows that Chapter 222's impact on charter discipline was primarily driven a reduction in out-of-school suspensions, and that the policy's effect on the charter probability of suspension was slightly smaller (6.3 percent) when measured relative to its signing. Consistent with Figure 14,this difference shows that the effect of Chapter 222 grew over time as schools adopted the policy. Finally, note that Table 3.1 omits the estimated effect of Chapter 222 on English test scores because—as indicated by Appendix Figure C.3—the charter attendance effect on English scores were on a steady and positive before the introduction of Chapter

[6]Appendix Figure C.3 plots the equivalent estimates for English, showing that the charter attendance effect on test scores followed a steady positive trend prior to the introduction of Chapter 222. Due to pre-trends, the Difference-in-Differences estimates for the effect of Chapter 222 on English test scores are be reliable. Therefore, I focus the analysis of Chapter 222's effect on n math test scores and suspensions outcomes only, but report point estimates for all year-by-year outcomes in Appendix Table C.4, along with information on first stages and sample sizes.

[7]Appendix Tables C.8 and C.9 show pooled and year-by-year covariate balance regression results, respectively, documenting that charter lottery offers were as good as randomly assigned. Appendix Table C.10 shows no differential attrition by charter lottery offer status in the sample.

[8]As shown in Appendix Table C.2, these large average gains during middle school reflect the fact charter attendance effects grow with years of charter attendance, with the effects being smallest—though already substantial—in the first year after the lottery (0.400 standard deviations for math, 0.227 standard deviations for English), and largest in the fourth year after the lottery (0.814 math, 0.716 English).

222, such that the difference-in-differences estimate for English is not reliable.[9]

3.4 Mechanisms

Table 3.1 suggests that Boston charter schools' suspensions practices are orthogonal to their ability to deliver large test score gains, since they delivered these gains even after reducing suspensions. In this section, I investigate two mechanisms through which Chapter 222 could have reduced charter suspensions without reducing the charter attendance effect on test scores.

The first possibility is that the reduction in charter suspensions induced by Chapter 222 benefited students who would have otherwise been suspended (for example, by keeping them in class), and harmed non-suspended students (for example, by not removing distracting behavior from classrooms), such that on average the charter attendance effect on test scores remained constant. The second possibility is that suspensions had no effect on learning, meaning that they were both inconsequential to suspended students' test scores and unnecessary for non-suspended students' learning gains. While the first mechanism requires that the effect of a charter suspension on test scores be negative, the second requires that it be zero.

Differentiating between these two mechanisms therefore requires identifying the causal effect of suspensions on charter students. The key identification challenge in this analysis concerns student selection into suspension. In particular, unlike charter attendance, which is as good as randomly assigned via lottery, suspensions are not randomly assigned, and the behaviors that lead to suspensions are often a consequence of complex unobserved factors that also negatively affect students' learning, such as problems at home (Steinberg and Lacoe, 2017). [10] I next describe how I address this identification challenge by combining lottery

[9]While the total number of suspended students is too small to allow for a breakdown of Chapter 222's impact by suspension offense type, comparing the rates of suspension in charter relative to TPS before and after Chapter 222 suggests that Chapter 222's primary incidence was on its intended offense type: non-drug, non-violent, and non-criminal offenses. Out-of-school suspensions for this offense type declined both in Boston TPS schools and in charters, though with a more pronounced decline for charters (5 percentage points) than for TPS (1 percentage point). The other offense types along with their average percent incidence are: criminal offenses, violent offenses, and bullying, harassment or property offenses. Less than one percent of students in either charters or TPS are suspended out-of-school under these offense types.

[10]As shown in Appendix Table C.5, this pattern holds true among charter applicants. Suspended students who attend Boston TPS score 0.163 and 0.150 standard deviations lower in math and English, respectively, than their non-suspended peers. In charter schools, the suspended vs. non-suspended test score gap is 0.110 standard deviations for math and 0.096 for English. The key identification question is whether these gaps are in fact caused by suspensions, or by omitted factors outside of the school's control.

offers with students' pre-lottery suspension records.

3.4.1 The effect of suspension on charter students' test scores

I start my investigation of the causal effect of suspension on charter students' test scores by pooling outcomes of applicants from all grades into a simple additive effects regression. In this regression, the effect of attending a charter and the effect of being suspended have a linear and additive effect on test scores:

$$Y_{ik} = \quad \phi + \chi_1 D_{ik} + \chi_2 S_{ik} + \omega' X_i + \varepsilon_{ik}, \tag{3.4}$$

where D_{ik} is a dummy indicating whether applicant i attended any charter, S_{ik} is a dummy for whether applicant i was suspended (out-of or in-school) in the kth year after the lottery, and X_i is defined as in Equation 3.1.

To overcome the challenge that suspensions are not randomly assigned, I leverage cross-charter variation in suspension rates to separately identify the effect of suspension from the effect of any-charter attendance on test scores. The key idea is that if a student wins a charter lottery to a school with a stricter disciplinary environment, then she is more likely to be suspended in that charter, but continues to experience the same charter attendance effect as students who attended other charters. Note that this approach implicitly assumes that — suspension decisions aside— the effect of charter attendance on test scores is homogenous across charters.

To implement this strategy, I instrument both suspensions and charter attendance using each applicant's vector of charter lottery offers (rather than with a single dummy indicating an offer from any charters). Since charters' disciplinary codes and educational philosophies are determined at the charter network level, I instrument D_{ik} and S_{ik} with a full set of charter school network dummies, each indicating whether the applicant received an offer from one of the charter schools in the respective network.[11]

Columns (1) and (2) of Table 3.2 present estimates of χ_1 and χ_2 from Equation 3.4. Column (1) shows that being suspended does not significantly impact a student's test scores, while Column (2) shows that attending a charter increases math test scores by 0.314 standard deviations (standard error 0.047), and English test scores by 0.107 (0.047) standard deviations. These findings suggest that being suspended does not meaningfully affect a student's

[11] A total of nine networks span the fifteen charter schools for which lottery records are available, listed in Appendix Table C.11. Some networks have only one charter school operating in Boston.

test scores, but attending a charter does.

3.4.2 Robustness

A potential concern with the estimates in Columns (1) and (2) Table 3.2 is that the individual charter lottery offers may not provide enough variation in suspension treatment assignment, as suggested by the 9.674 first-stage F-statistic on the suspension treatment, slightly below the rule-of-thumb F-statistic of 10 for rejection of weak identification. To overcome a weak first stage, I create additional instruments for D_{ik} and S_{ik} using applicants' pre-lottery suspensions data. Specifically, I first estimate suspension propensity scores σ_i for each applicant i, and create new instruments for D_{ik} and S_{ik} by interacting the individual network lottery offers with σ_i.[12] Each applicant's σ_i is also included as a control in Equation 3.4 and its corresponding first stage equations,[13] such that χ_1 and χ_2 are estimated among students of similar suspension propensity.

I present results for this alternative exercise in Appendix Table 3.2. Columns (1)–(3) present estimates of χ_1 and χ_2 where interactions of the charter dummies with σ_i are added to the set of instruments to Equation 3.4. Since suspensions are rare in the pre-lottery years, most students have a very low σ_i. As a result, the average realized probability of suspension post-lottery is higher than the predicted suspension probability at low σ_i values. To address the concern that rare suspensions in the pre-lottery data might underestimate the lower values of σ_i, in Columns (4)–(6) I dichotomize the σ_i distribution with a dummy variable indicating high suspension propensity, and use that dummy variable instead of σ_i itself to generate the interactions with charter network offer dummies. The high suspension propensity dummy includes all applicants with σ_i greater than 0.2 on a scale of 0 to 1, which is the cutoff above which the post-lottery suspension probability averages fifty percent.[14]

The results in Columns (3) and (4) of Table 3.2 are qualitatively similar to those in

[12]Setren (2017) uses a similar methodology to analyze the impact of special education re-classification in charters. I estimate the student-specific suspension propensity in two steps. First, I use the sample of non-charter applicant 5th–8th graders attending TPS schools to estimate a logistic regression of an out-of-school suspension dummy on various student demographics and disciplinary records. I then predict σ_i for each charter applicant using the applicant's demographics and baseline grade disciplinary data and the logit coefficients of each predictor from the regression estimated in the TPS sample. Finally, I include the student-specific propensity score as a control in the second stage and first stage regressions.

[13]I present the distribution of σ_i among charter applicants in Appendix Figure C.4. Panel A plots this distribution separately by suspension treatment post-lottery, showing that σ_i is predictive of suspensions, while Panel B plots it separately for lottery winners and losers, showing that σ_i is balanced across charter lottery offer status.

[14]Since suspensions are rare in the pre-lottery data, $\sigma_i = 0.20$ is the 95th percentile of the σ_i distribution.

Columns (1) and (2), but the first stage in Column (3) suggests a much stronger identification of suspension effects (with an F-statistic of 15) and near zero suspension effect point estimates of -0.047 (0.149) for math and 0.065 (0.145) for English. Finally, I now augment Equation 3.4 with the interaction term $D_{ik} \times S_{ik}$ in order to test whether suspensions have different effects on test scores if they happen at charters versus at TPS:

$$Y_{ik} = \quad \phi + \chi_1 D_{ik} + \chi_2 S_{ik} + \chi_3 (D_{ik} \times S_{ik}) + \omega' X_i + \varepsilon_{ik}. \tag{3.5}$$

With Equation 3.5 I can separately estimate the effect of suspension on charter students $(\chi_1 + \chi_3)$ and the effect of charter attendance on suspended students $(\chi_2 + \chi_3)$. Table 3.3 presents estimates of χ_1, χ_2, and χ_3 using the dummied instruments as in Columns (3)–(4) of Appendix Table 3.2.[15] While the first-stage F-statistics suggest weak identification of coefficients χ_2 and χ_3, I find some evidence that the net-positive effect of charter suspensions shown in Column (6) of Appendix Table 3.3 is driven by a large positive effect of charter attendance on suspended students (Column 5) and a zero causal effect of suspensions on charter students' test scores (Column 4). Finally, consistent with the hypothesis that the causal effect of suspensions depends on which school the suspended student attends, Column (1) suggests that the causal effect of being suspended at a TPS school on test scores is large and negative.

Taking all robustness exercises into account, the results from Table 3.2 stand: suspensions have no meaningful or statistically significant effect on charter students' math test scores.

3.5 Conclusion

Previous charter lottery studies have documented large test score gains from charter attendance, with the largest gains driven by charters who suspend more students Angrist et al. (2013); Chabrier et al. (2016). This paper leverages Boston charter middle school lotteries and Chapter 222, a Massachusetts policy aimed at reducing charter suspensions, to understand whether suspensions are a key component of charters' ability to deliver large average learning gains.

I find that Chapter 222 successfully reduced the charter attendance effect on suspensions by 10 percentage points three years after the policy took effect, nearly closing the charter to Traditional Public Schools suspensions gap. However, I find no evidence that charters'

[15] Appendix Table C.7 presents robustness checks using the same instruments as in Columns (1)–(2) of Table 3.2.

reduction in suspensions reduced the charter attendance effect on student test scores. I then investigate the mechanisms behind Chapter 222's effects, finding suggestive evidence that the causal effect of suspensions on Boston charter students' test scores is zero, whereas the causal effect of charter attendance on suspended students' test scores is large and positive. Overall, these findings indicate that non-suspended students' large test score gains at charter schools are *not* obtained at the expense of suspended students' learning.

Figure 14: 2SLS DD effects on probability of suspension

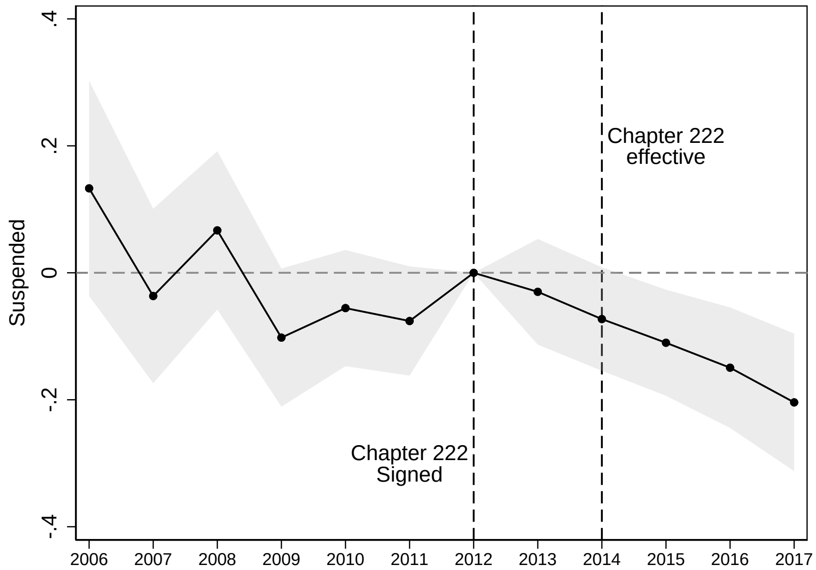

Note: This figure shows 2SLS difference-in-differences estimates of charter attendance on whether the student was ever suspended in the school year, relative to base year 2012. The treatment is a charter attendance dummy, interacted with year dummies, and omitting school year 2012. The instrument is a charter offer dummy interacted with the same dummies. The specification controls for applicant risk sets, grade and outcome year dummies, and baseline covariates.

Figure 15: 2SLS DD effects on MCAS math test scores

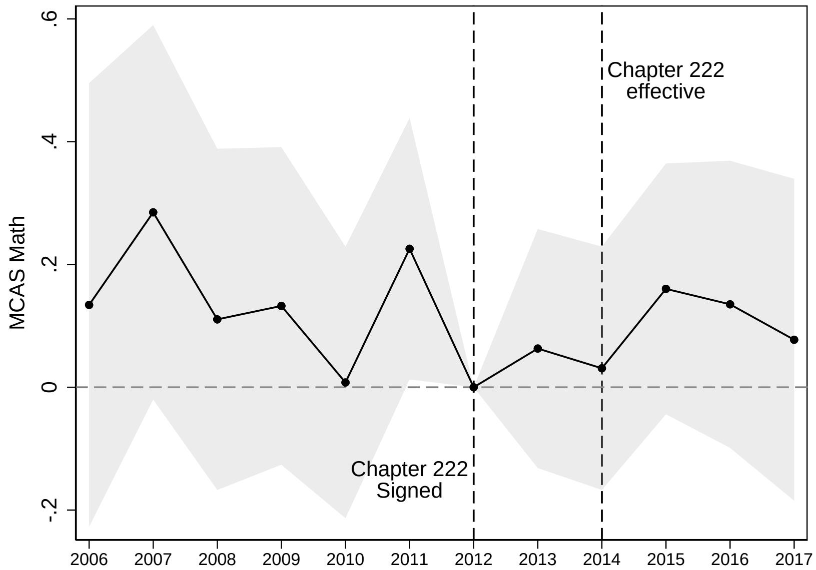

Note: See notes to Figure 15. See the Data Appendix for details on Massachusetts' MCAS tests.

Table 3.1: Charter/Suspension reform interactions

	Chapter 222 signed		Chapter 222 effective	
	Ever attended charter	Ever attended charter x (After 2012)	Ever attended charter	Ever attended charter x (After 2014)
	(1)	(2)	(3)	(4)
Panel A: Second stage				
Disciplinary outcomes				
Suspended	0.216***	-0.063**	0.205***	-0.095***
	(0.017)	(0.026)	(0.014)	(0.027)
Out-of-school	0.185***	-0.046*	0.182***	-0.093***
	(0.017)	(0.025)	(0.014)	(0.026)
In-school	0.072***	-0.016	0.065***	-0.007
	(0.009)	(0.014)	(0.008)	(0.015)
MCAS test scores				
MCAS Math	0.613***	-0.013	0.591***	0.072
	(0.041)	(0.062)	(0.034)	(0.065)
Panel B: First stage				
Charter offer	0.443***	-0.006***	0.399***	-0.014***
	(0.009)	(0.002)	(0.007)	(0.002)
Charter offer x (After 222)	-0.120***	0.330***	-0.091***	0.342***
	(0.012)	(0.008)	(0.012)	(0.010)
F-statistic	3,511	3,623	3,511	3,623
N	29,096		29,096	

Notes: This table displays coefficients from Two Stage Least Squares Difference-in-Differences (2SLS DD) regressions whose goal is to quantify the effect of Chapter 222 on charter attendance 2SLS treatment effects. The 2SLS DD procedure is implemented as a two-endogenous variable, two-instrument 2SLS regression, where the treatment variables are a dummy for whether the charter applicant ever attends charter, and an interaction of this dummy with whether the outcome variable year is in or after Chapter 222. Columns (1)-(2) displays results using Chapter 222's signature year (2012) to construct the interaction dummy, whereas Columns (3)-(4) use Chapter 222's effective year (2014). All regressions control for fully-saturated charter application risk sets and baseline grade covariates. Since charter applicants enter the sample in different years and at different grades, all regressions include outcome year, grade, and years-since-charter-lottery fixed effects. Sanderson-Windmeijer (2015) F-stats based on Angrist and Pischke (2009) are displayed for each the excluded instruments in the first stage regression. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table 3.2: Charter and suspension effects identified using cross-school lottery variation and predicted suspensions

			Interacting individual charter offer instruments with			
			Applicant suspension propensity score		Dummy for suspension propensity score > 0.20	
	Suspended	Attended charter	Suspended	Attended charter	Suspended	Attended charter
	(1)	(2)	(3)	(4)	(5)	(6)
MCAS test scores						
Math	0.076	0.314***	-0.120	0.355***	-0.047	0.336***
	(0.177)	(0.047)	(0.140)	(0.042)	(0.149)	(0.043)
English	0.230	0.107**	0.061	0.151***	0.065	0.139***
	(0.174)	(0.047)	(0.144)	(0.044)	(0.145)	(0.044)
First stage F-statistics						
F-statistic	9.674	25.613	6.264	20.885	15.065	22.404
Degrees of freedom	9	9	19	19	19	19
N			8,149		8,149	

Notes: The instruments in Columns (1)-(2) are the charter network-specific offer dummies, while those in Columns (3)-(4) are the charter network-specific offer dummies plus interactions with an applicant-specific suspension propensity score. The instruments in Columns (4)-(6) are the charter network-specific offer dummies plus interactions with a dummy indicating whether the applicant-specific suspension propensity is above 0.20 (in a scale of 0 to 1). See Table C.6 and Figure C.4 for details on the estimation and distribution of the student suspension propensity. All regressions control for applicant risk set dummies and all baseline grade covariates listed in Table C.6. Sanderson-Windmeijer (2015) F-stats based on Angrist and Pischke (2009) are displayed for each the excluded instruments in the first stage regression. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table 3.3: Suspension effect in charters vs. Charter effect on suspended

	Instruments: individual charter offers plus interactions with applicant suspension propensity score > 0.20				
	Treatments				
	Suspended	Attended charter	Attended charter x Suspended	Suspension effect in charters (1)+(3)	Charter effect on suspended (2)+(3)
	(1)	(2)	(3)	(4)	(5)
MCAS test scores					
Math	-0.629*	0.291***	0.589*	-0.040	0.879***
	(0.375)	(0.050)	(0.332)	(0.154)	(0.312)
English	-0.861**	0.057	0.967***	0.105	1.024***
	(0.345)	(0.052)	(0.318)	(0.152)	(0.295)
First stage F-statistics					
F-statistic	3.881	28.263	4.846	--	--
Degrees of freedom	18	18	18	--	--
N			8,149		

Notes: This table displays coefficients from an over-identified three-endogenous variables 2SLS regressions of charter attendance and suspensions on applicant test scores. The treatment variables are a dummy for whether the charter applicant ever attends charter, a dummy for whether the applicant is ever suspended, and the interaction of these two dummies. The instruments are charter network-specific offer dummies, and interactions of these dummies with a dummy indicating whether the applicant-specific suspension propensity score is above 0.20 (in a scale of 0 to 1). See Table C.6 and Figure C.4 for details on the estimation of the student suspension propensity. Columns (1)-(3) display the second state coefficients of the effect of each treatment on the respective outcome variables, while Columns (4)-(5) report net average treatment effects of interest, obtained as linear combinations of the estimated coefficients. All regressions control for applicant risk set dummies and baseline grade covariates. Sanderson-Windmeijer (2015) F-stats based on Angrist and Pischke (2009) are displayed for each the excluded instruments in the first stage regression. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Appendix A

Appendices for Chapter 1

A.1 Chapter 1 Appendix Figures and Tables

Figure A.1: 1991 share of Brazilian exports by sector

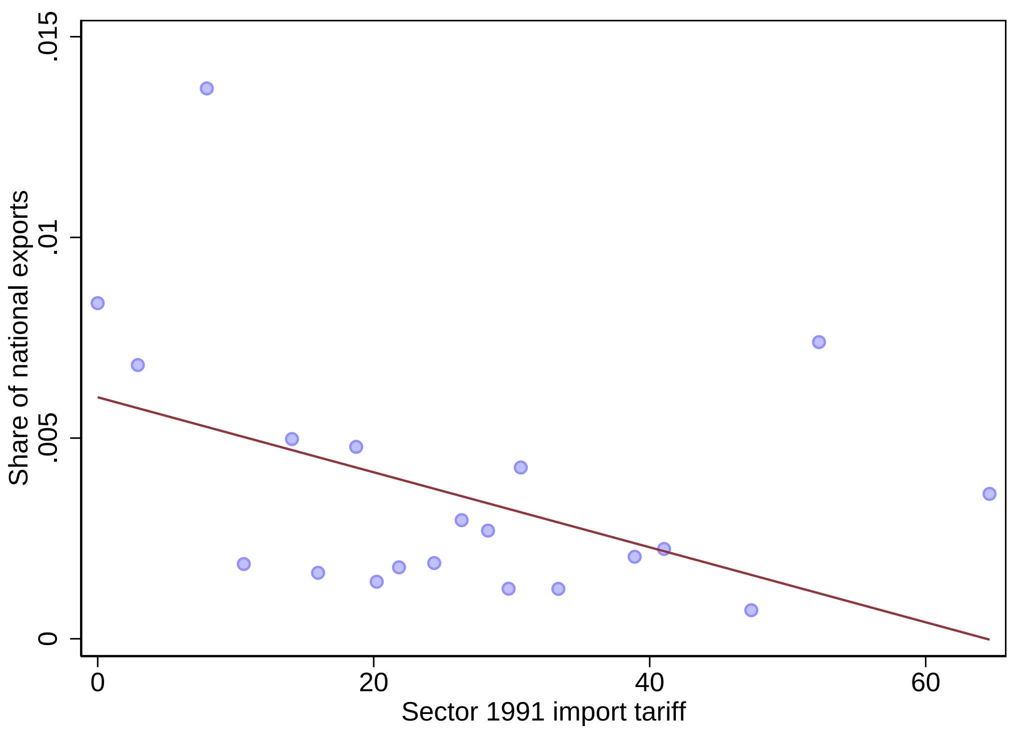

Note: This figure plots a binned scatter of the relationship between a sector's 1991 share of the value of national exports (measured in USD) with the sector's 1991 import tariff. Exports and tariffs are measured at the CNAE95 level (for a total of 285 tradable sectors).

Figure A.2: Statistical significance across placebo regressions

Panel A: 1,000 placebo shock regressions on 1997-1991 change in wage premia

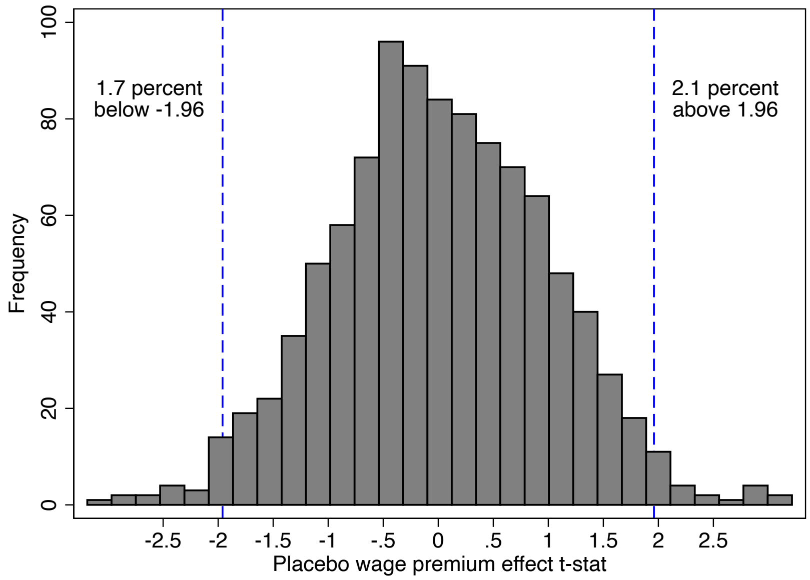

Panel B: 1,000 placebo shock regressions on 1997-1991 change in log employment

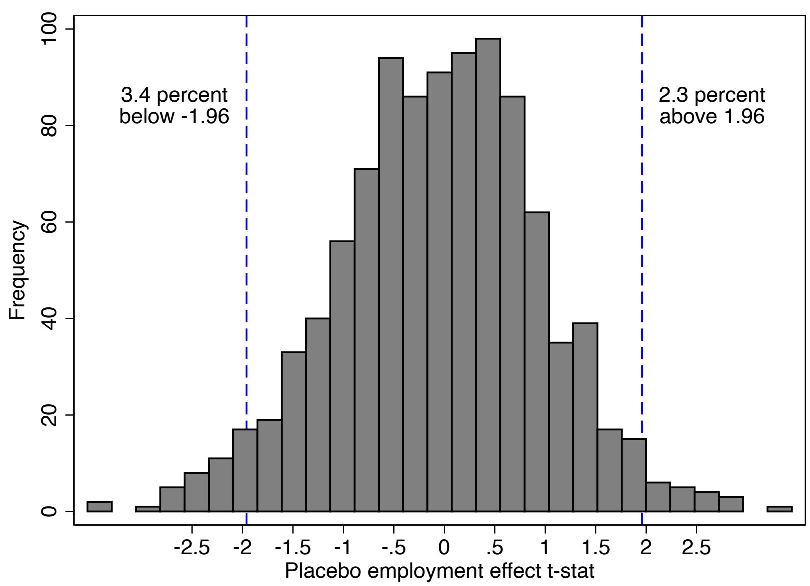

Note: This figure plots the distribution of t-statistics for 1,000 placebo regressions where the regressor is ΔICE_m defined as in equation 1.1 but using randomly drawn shocks, drawn from a normal distribution centered at zero.

Figure A.3: Worker labor supply decision

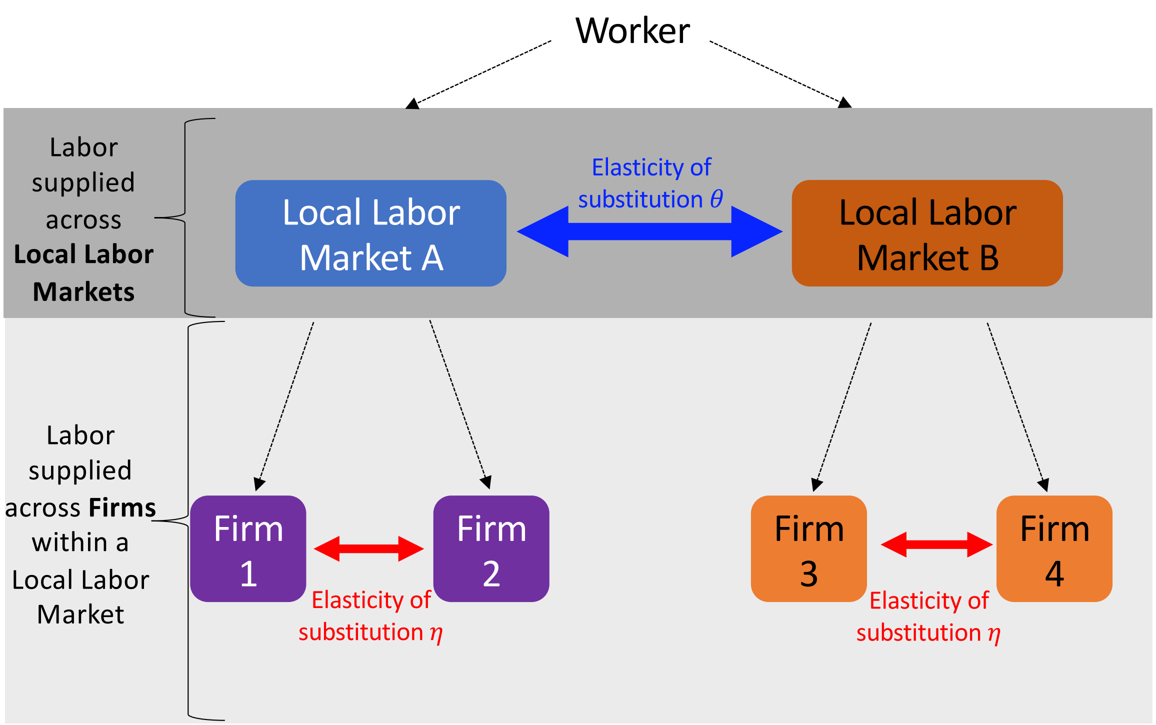

Note: This figure displays a diagram of worker's labor supply decision according to the discrete choice labor supply framework presented in Section 1.3.

Figure A.4: 1990-1991 cross-microregion transitions conditional on switching firms (Top 50)

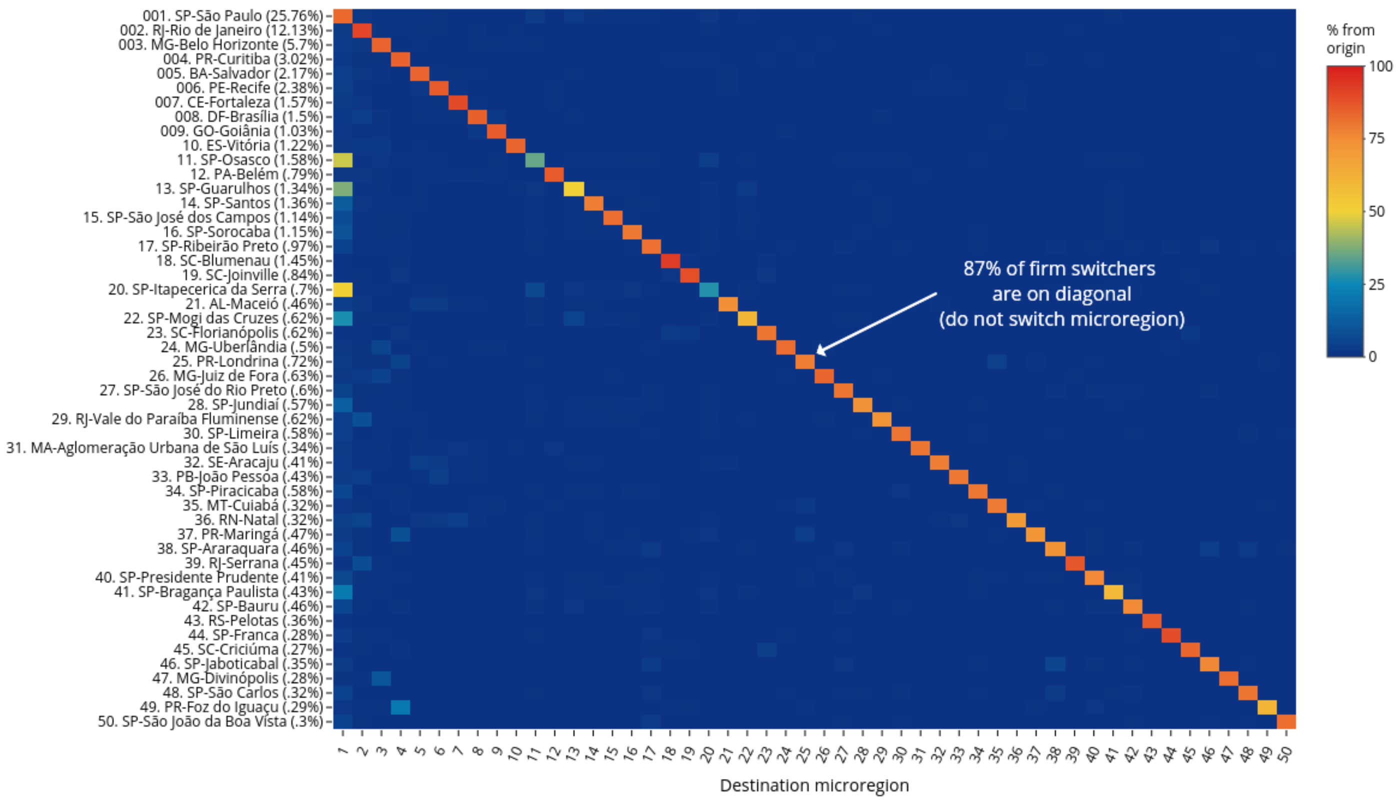

Note: This figure plots worker microregion to microregion transitions, among workers who switched employers between 1990 and 1991, for the top 50 microregions by number of workers at origin. Each row lists the origin microregion (with percent of total workers indicated in parentheses), while each column lists the destination microregion.

Figure A.5: 1990-1991 cross-occupation transitions conditional on switching firms (Top 50)

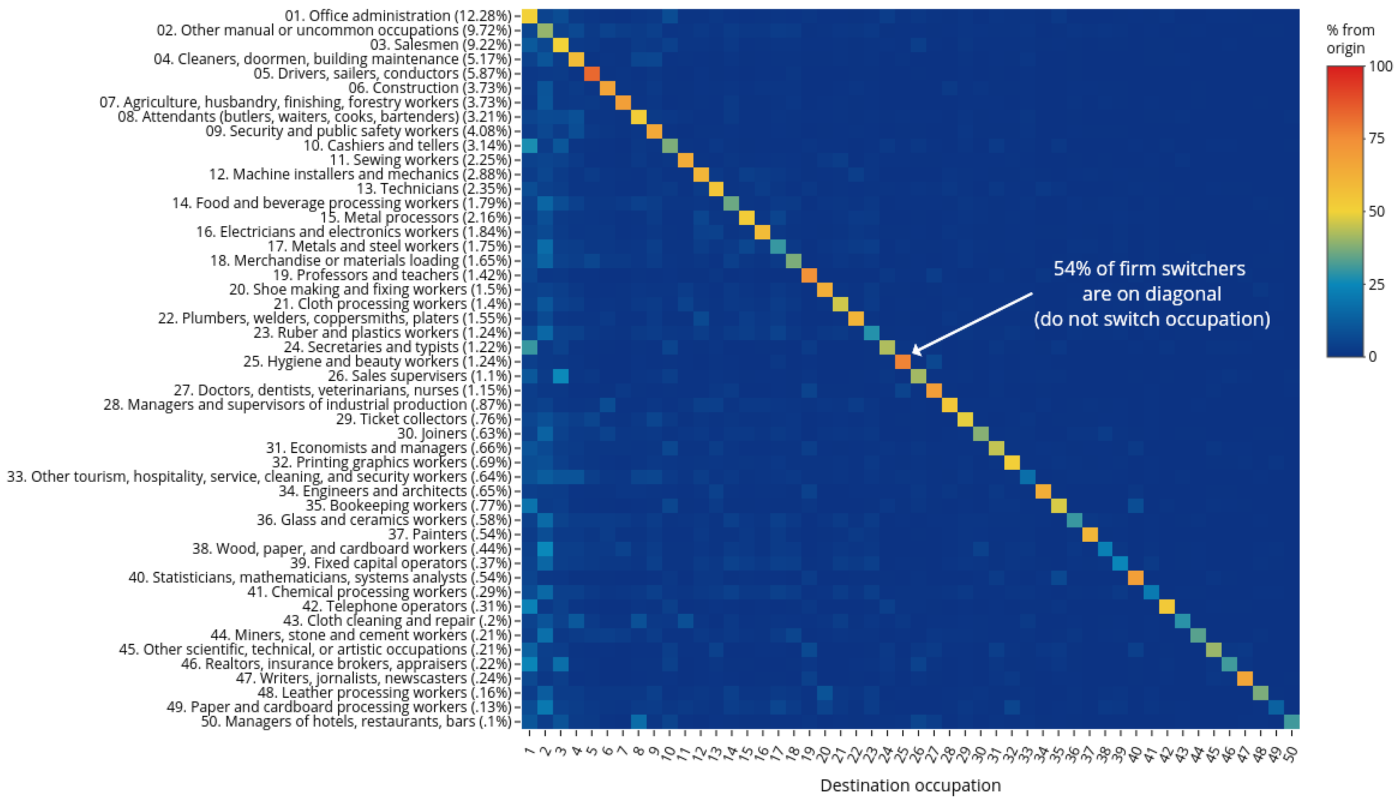

Note: This figure plots worker occupation group to occupation group transitions, among workers who switched employers between 1990 and 1991, for the top 50 occupation groups (2-digit CBO94) by number of workers at origin. Each row lists the origin occupation group (with percent of total workers indicated in parentheses), while each column lists the destination occupation group.

Figure A.6: 1990-1991 cross-sector transitions conditional on switching firms (Top 50)

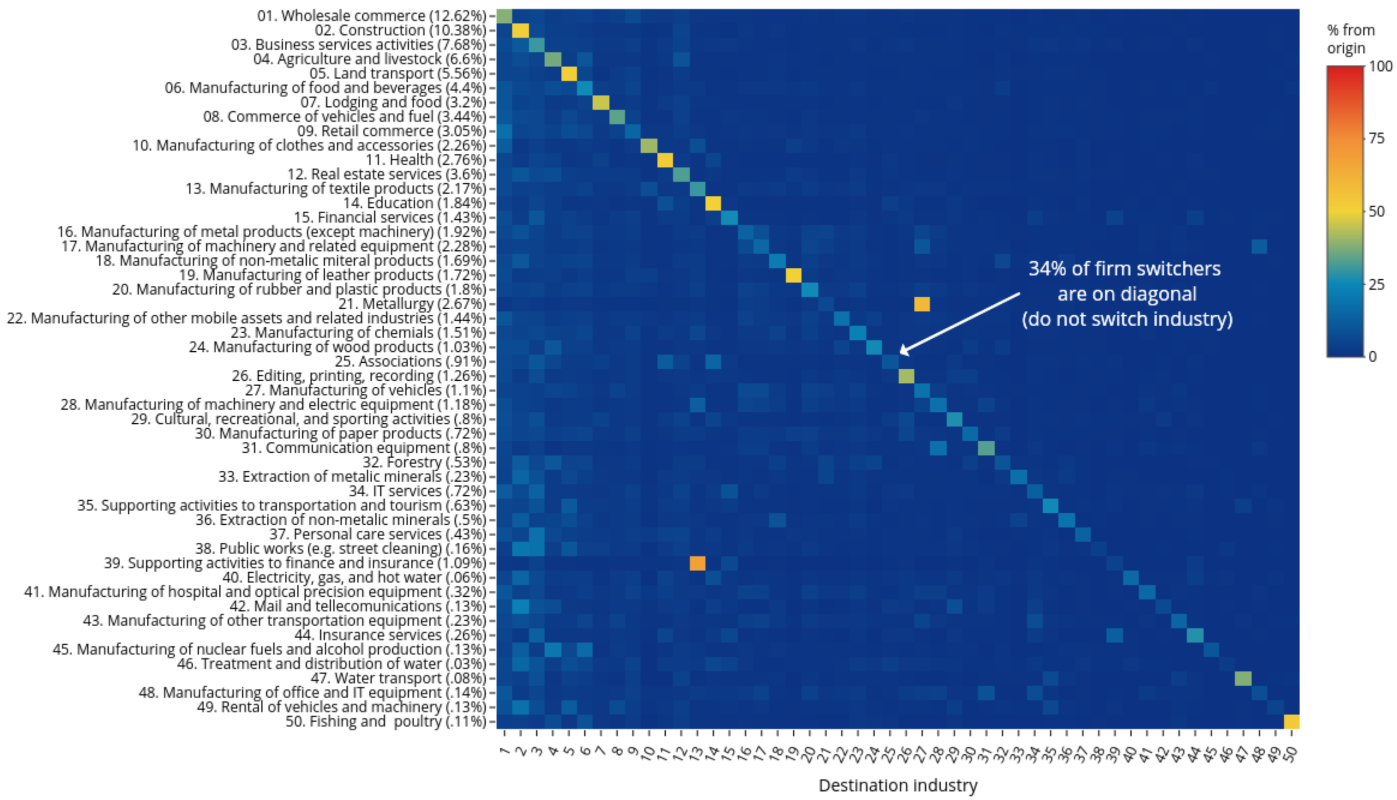

Note: This figure plots worker sector group to sector group transitions, among workers who switched employers between 1990 and 1991, for the top 50 sector (2-digit CNAE95) groups by number of workers at origin. Each row lists the origin sector group (with percent of total workers indicated in parentheses), while each column lists the destination sector group.

Figure A.7: Local labor market concentration

Panel A: Number of firms

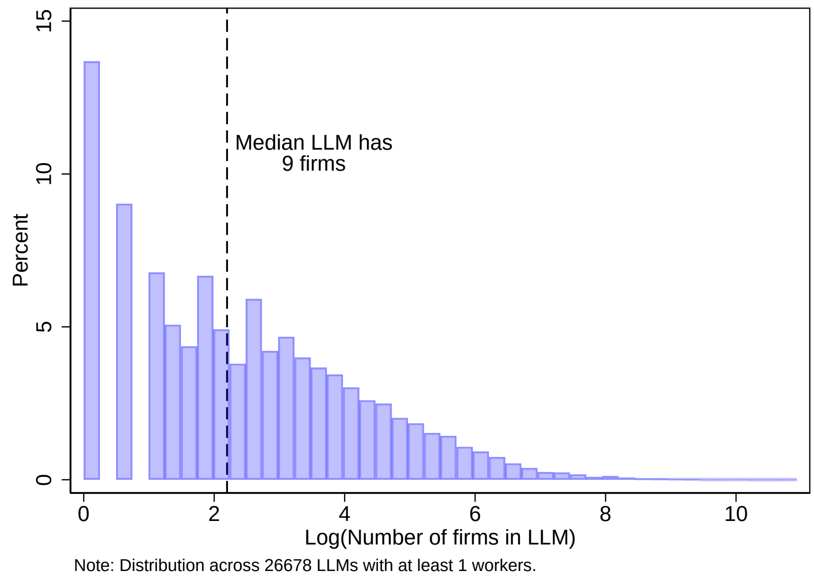

Panel B: Employment share of largest 2 firms

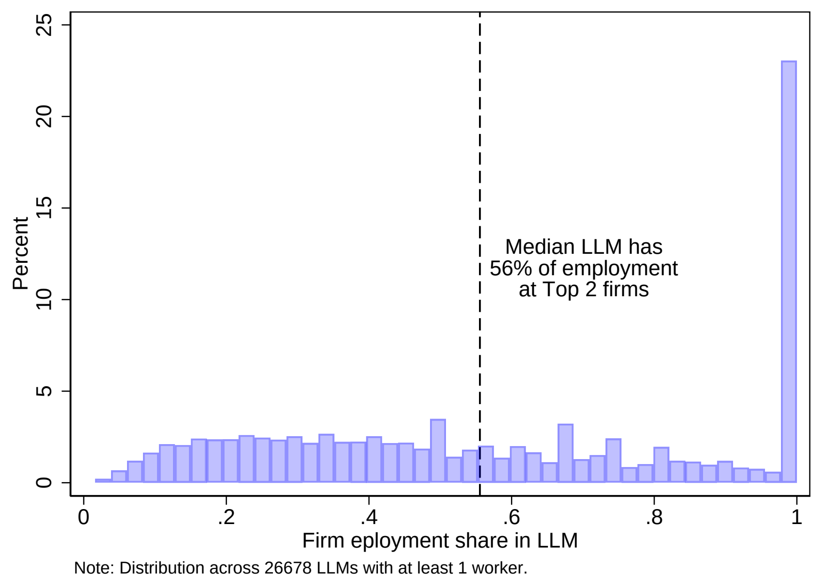

Note: This figure plots the 1991 distributions of number of firms (Panel A), and employment share of the largest 2 firms (Panel B) across local labor markets. Local labor markets are defined as a microregion $\times$ occupation group cell. See Appendix A.2 for details on the definitions of microregion and occupation group.

Figure A.8: Effect of import competition on exporter vs. non-exporter size

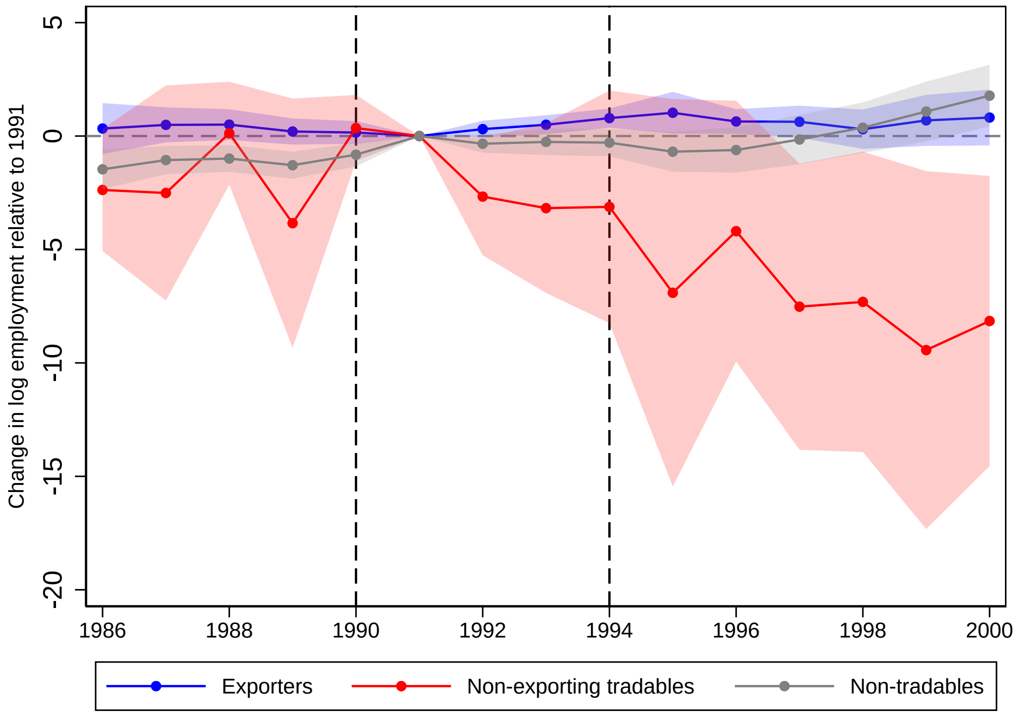

Note: This figure plots coefficients of three regressions about the cumulative effect of the change in import competition exposure: on changes in log employment of exporters; on changes in log employment of non-exporting tradables; and on changes in log employment of non-tradables. Each point is a β_k coefficient from equation 1.2. Dotted lines indicate the beginning and end of the tariff reductions reform. So that all differences reflect a change from a future year to a past year, for the pre-reform years the outcome is the 1991 log employment minus each respective year's log employment, whereas for the post-reform years the outcome is each respective year's log employment minus the 1991 log employment. All regressions are weighted by 1991 employment. Standard errors are two-way clustered by microregion and occupation group.

Figure A.9: Effect of import competition on local labor market wages

Panel A: Relative to trend

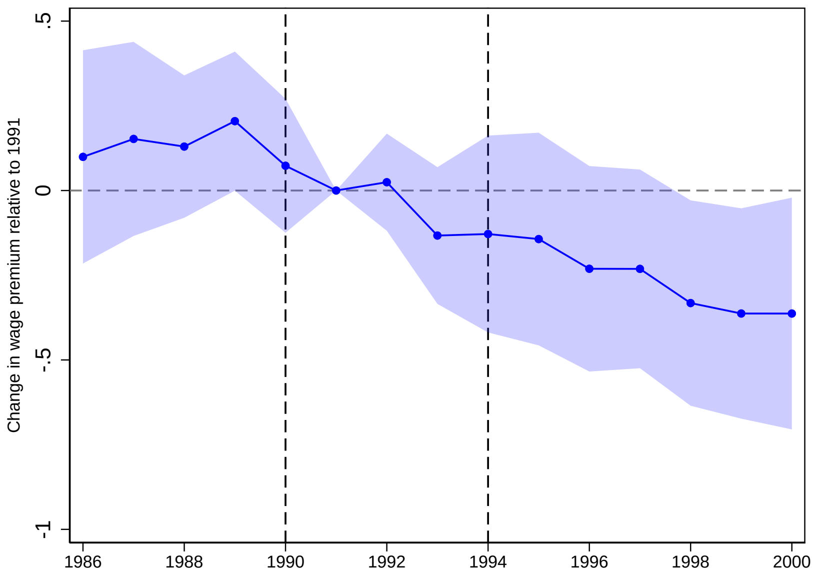

Panel B: Relative to base year

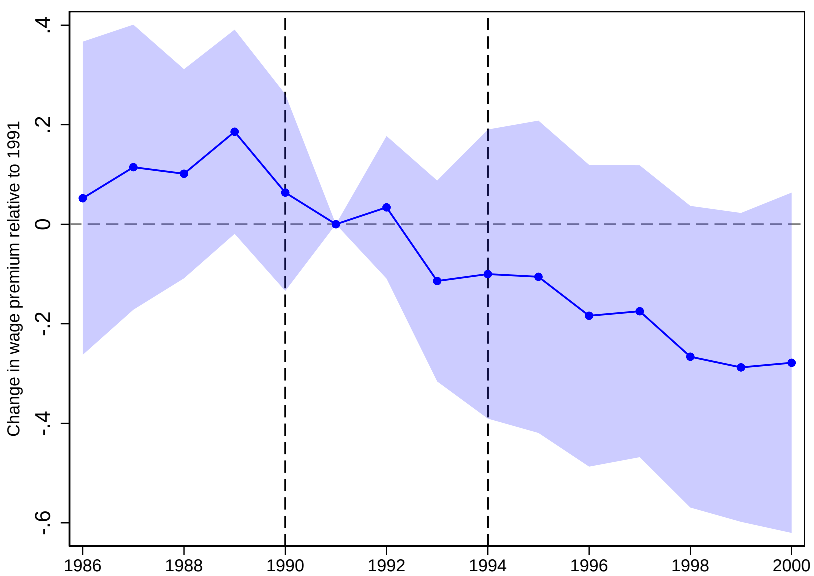

Note: See notes to Figure 4.

Figure A.10: Effect of import competition on employment

Panel A: Relative to trend

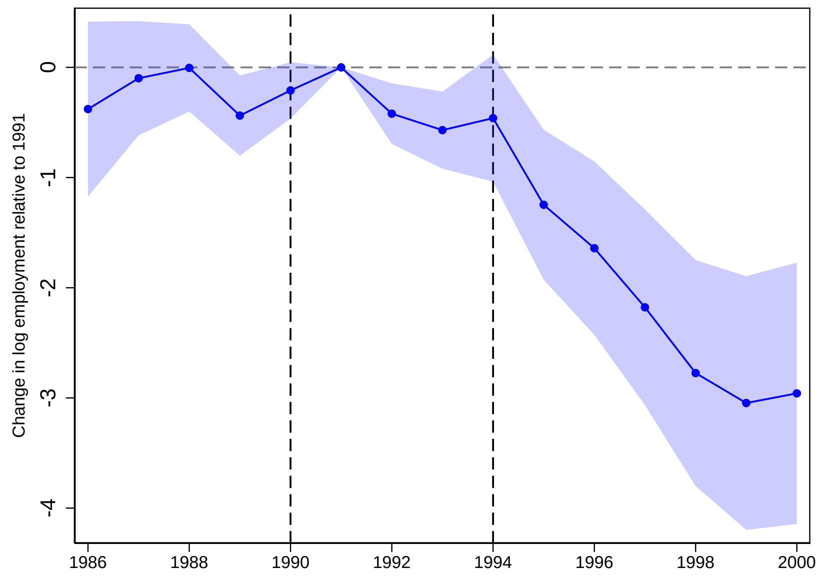

Panel B: Relative to base year

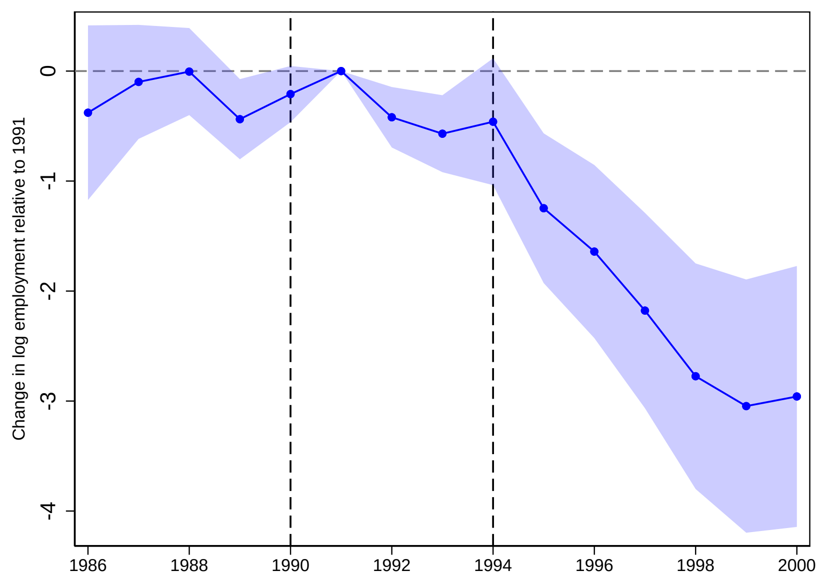

Note: See notes to Figure 4.

Figure A.11: Brazil's 1990-1994 tariff reduction reform: variation across 285 sectors

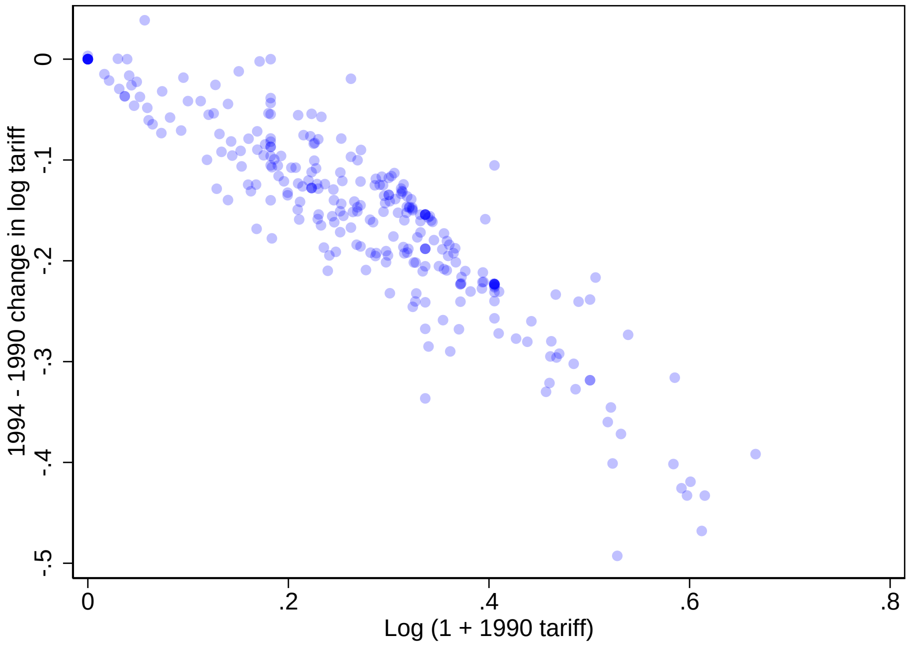

Note: This figure plots import tariff reductions from Brazil's 1990-1994 import tariff reform. See Section 1.4 for details.

Figure A.12: Correlation between firm labor market power measures, employment, and wage premia

Panel A: Wage markdowns

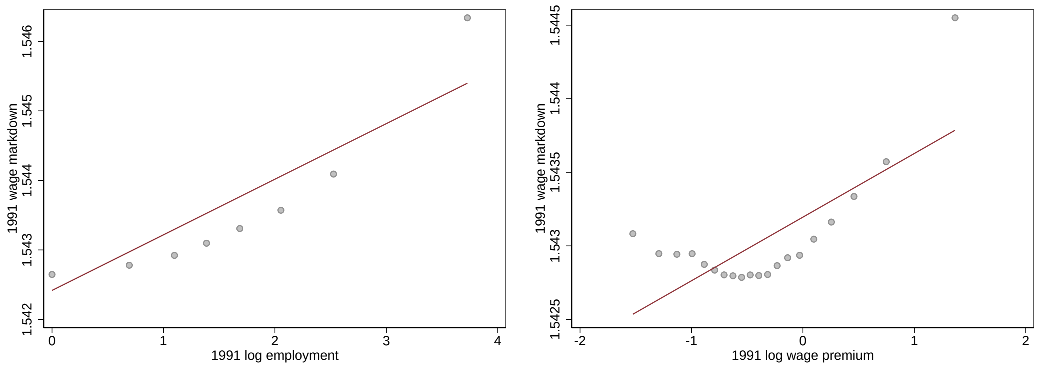

Panel B: Wage take-home share

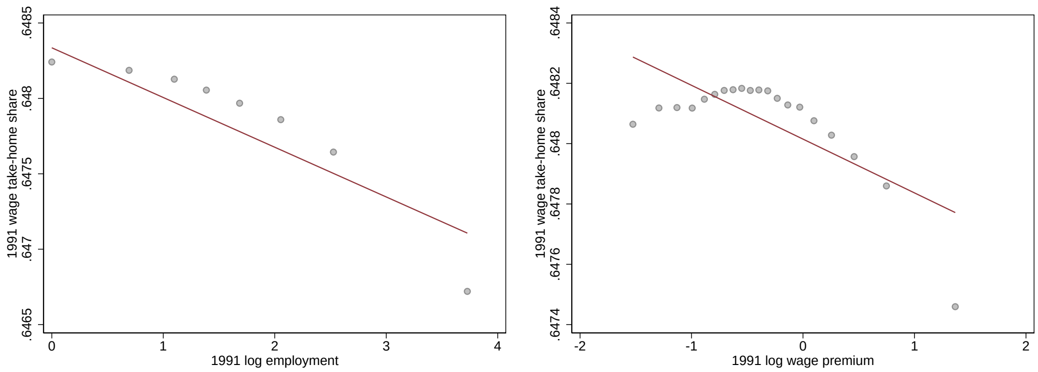

Note: This figure plots correlations between estimated firm wage take-home shares and base year employment and firm wage premia. See equation 1.13 for the expression of wage take-home shares (i.e., the inverse of wage markdowns), and Section A.2 for details on wage premia estimation.

Figure A.13: Effect of tariff reductions on firm-market-level employment and wage premia

Panel A: Wage premia

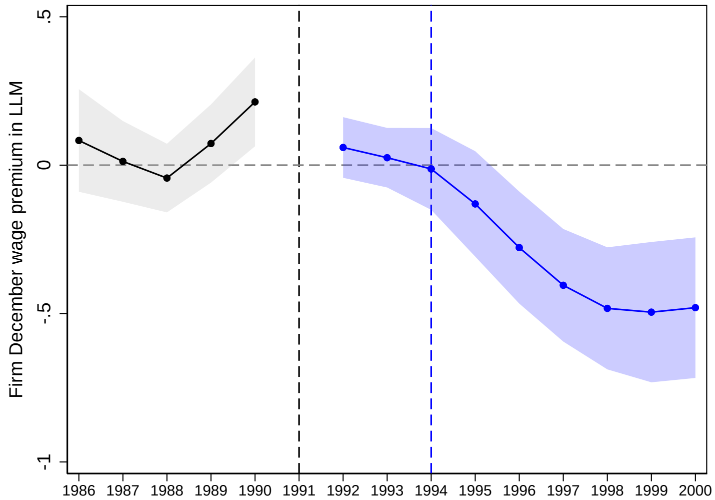

Panel B: Employment

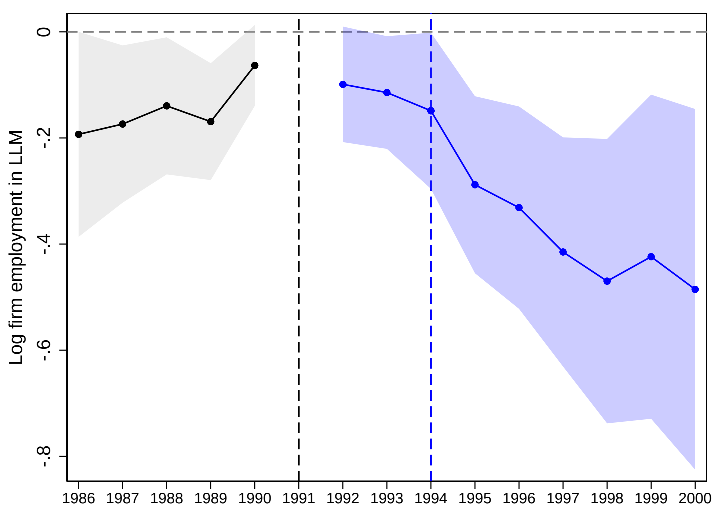

Note: This figure plots coefficients of regressions of firm-level changes in log employment (from each year to the base year of 1991) on minus $\ln\left(\frac{1+\tau_{1994}}{1+\tau_{1990}}\right)$, which is the firm-level change in import competition exposure, separately estimated for each year. Dotted lines indicate the beginning and end of the tariff reductions reform. So that all differences reflect a change from a future year to a past year, for the pre-reform years the outcome is the 1991 log employment minus each respective year's log employment, whereas for the post-reform years the outcome is each respective year's log employment minus the 1991 log employment. All regressions are weighted by 1991 firm employment. Standard errors are clustered at the sector level.

Figure A.14: Wage take-home shares: 1991 vs. 1997

Note: This figure plots the cross-firm national distribution of wage take-home shares (i.e., the inverse of wage markdowns) in the base year before the reform (1991) and the mid-year of the post-reform period (1997).

Figure A.15: Pre-reform distribution of firm size and wages

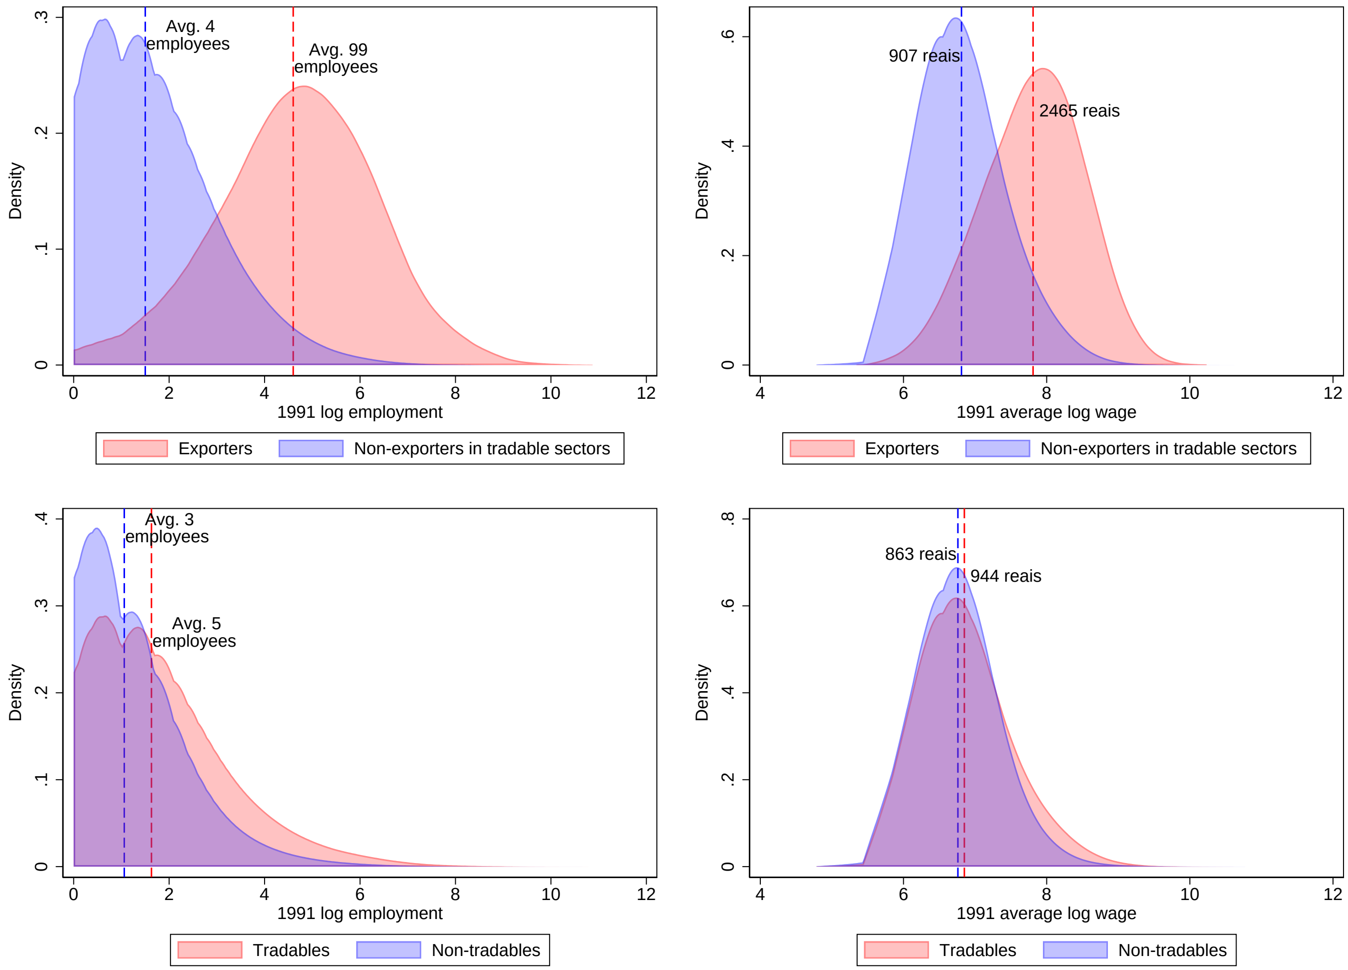

Note: This figure plots pre-reform distributors of firm log employment and log December monthly wages for exporters, non-exporters, and non-tradables. Wages are reported in 2017 reais.

Figure A.16: Regional concentration vs. informality

Panel A: Wage bill concentration

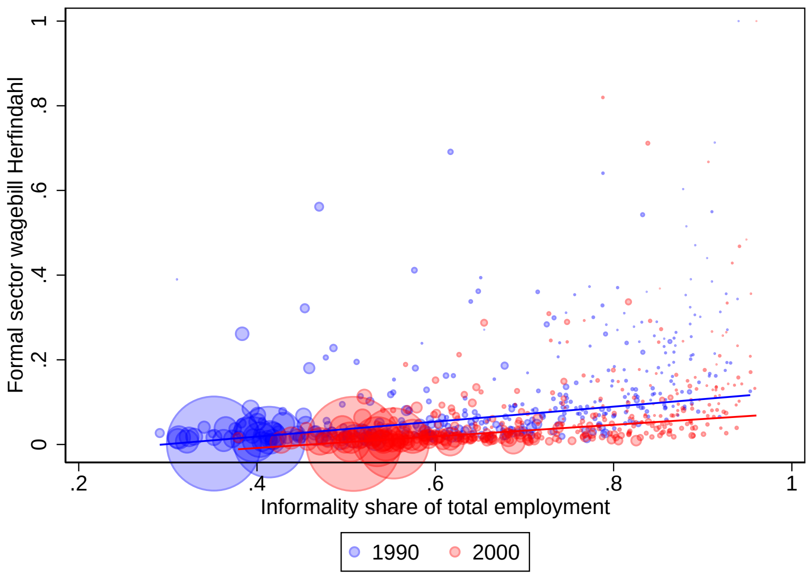

Panel B: Employment concentration

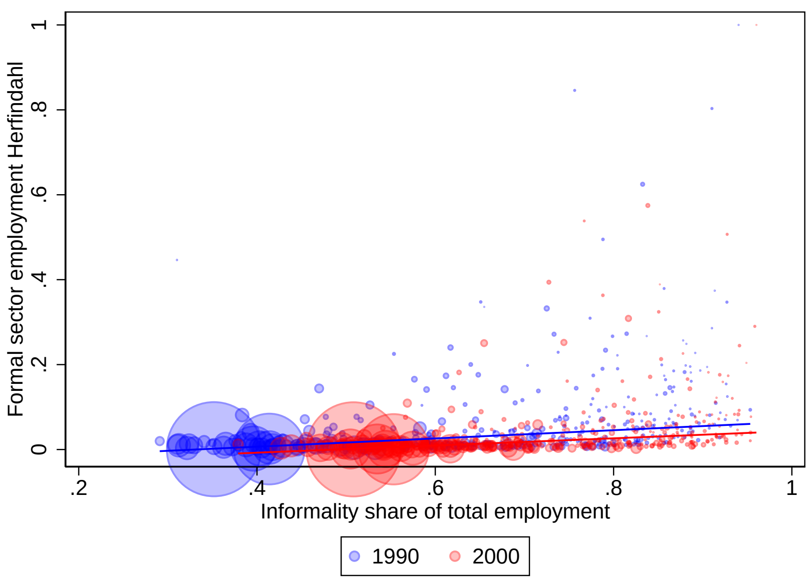

Note: This figure plots microregion-level concentration measures computed from RAIS against microregion-level measures of informality share from the 1990 and 2000 census. Census data was obtained from the supplemental materials to Dix-Carneiro and Kovak (2017).

Table A.1: Local labor market descriptive statistics

	Mean	Market percentile				
		10th	25th	50th	75th	90th
	(1)	(2)	(3)	(4)	(5)	(6)
Total market employment	659	6	16	59	252	921
Tradable sector	268	0	4	20	98	387
Exporters	235	0	1	9	64	294
Δ Import Competition Exposure	5%	0%	1%	3%	8%	14%
Average December earnings (2017 reais)	R$ 2,275	R$ 664	R$ 920	R$ 1,498	R$ 2,683	R$ 4,802
Numer of firms	115	3	6	16	55	183
Number of exporters	16	0	1	2	8	24
Employment Herfindahl	0.23	0.03	0.06	0.15	0.31	0.54
Wage Bill Herfindahl	0.28	0.04	0.09	0.21	0.39	0.63
Average firm size	6	1	2	3	5	10
Top 25 percentile firms	541	4	11	43	192	734
25th - 50th percentile firms	85	3	5	13	40	128
Bottom 25 percentile firms	58	2	4	9	29	93

Note: This table presents descriptive statistics across 21,242 Brazilian local labor markets defined as microregion × occupation group pairs. See Appendix A.2 for details on microregion and occupation group definitions.

Table A.2: Workers' labor market transition probabilities conditional on switching firms

Total workers transitioning to different firm in 1990-1991	1,108,130
Percent staying in…	
Microregion (557 groups of municipalities)	79%
Occupational group (CBO94 / 2-digit / 65 groups)	53%
Microregion x Occupational group Pair	43%
Highly specialized occupation (CBO94 / 5-digit / 2,357 occupations)	32%
Industry Group (CNAE95 / 2-digit / 59 groups)	33%
Sub-industry (CNAE95 / 5-digit / 614 sub-industries)	18%

Note: This table presents statistics on the probability that a worker remains in the same (microregion, occupation group, etc.) conditional on the worker having switched firms. All probabilities are conditional on workers remaining in the formal sector.

Table A.3: Estimates of the inverse elasticity cross-firm within-market substitution $\frac{1}{\eta}$: heterogeneity by sample

	Main specification (1)	All firms (2)
Panel A: First stage		
Δ Firm log employment in LLM	-0.611	-0.436
	(0.146)	(0.097)
First stage F	17.568	20.255
Panel B: Reduced form		
Δ Firm's wage premium in LLM	-0.331	-0.517
	(0.091)	(0.094)
Panel C: 2SLS		
Labor supply within-market cross-firm inverse elasticity of substitution	0.542	1.185
	(0.207)	(0.222)
Implied upper bound on wage take-home share	65%	46%
Local labor market (LLM) FE	Yes	Yes
Observations	759,612	955,471
Firms	328,462	382,274
Local labor markets	14,918	18,293

Note: This table shows robustness to estimation sample of estimates of the cross-firm within-market inverse elasticity of substitution, estimated using equation 1.15. The main specification includes firms that are unique producers in their local labor markets, plus non-tradable sector firms.

Table A.4: Estimates of the inverse elasticity cross-firm within-market substitution $\frac{1}{\eta}$: robustness to clustering

	Main specification (1)	Clustered at local labor market (2)	Clustered at microregion (3)
Panel A: First stage			
Δ Firm log employment in LLM	-0.611	-0.611	-0.611
	(0.146)	(0.085)	(0.095)
First stage F	17.568	52.285	41.360
Panel B: Reduced form			
Δ Firm wage premium in LLM	-0.331	-0.331	-0.331
	(0.091)	(0.050)	(0.052)
Panel C: 2SLS			
Labor supply within-market cross-firm inverse elasticity of substitution	0.542	0.542	0.542
	(0.207)	(0.116)	(0.110)
Local labor market (LLM) FE	Yes	Yes	Yes
Observations	759,612	759,612	759,612
Firms	328,462	328,462	328,462

Note: This table shows robustness to clustering of estimates of the cross-firm within-market inverse elasticity of substitution, estimated using equation 1.15. Standard errors in the main specification are clustered at the CNAE95 economic sector level.

Table A.5: Estimates of the inverse elasticity cross-firm within-market substitution $\frac{1}{\eta}$: robustness to wage and shock

	Main specification	Using December average wage (instead of premium)	Using effective rate of protection (instead of tariff)
	(1)	(2)	(3)
Panel A: First stage			
Δ Firm log employment in LLM	-0.611	-0.611	-0.294
	(0.146)	(0.146)	(0.106)
First stage F	17.568	17.568	7.758
Panel B: Reduced form			
Δ Firm wage premium in LLM	-0.331	-0.377	-0.184
	(0.091)	(0.094)	(0.063)
Panel C: 2SLS			
Labor supply within-market cross-firm	0.542	0.616	0.626
inverse elasticity of substitution	(0.207)	(0.224)	(0.324)
Implied upper bound on wage take-home share	65%	62%	61%
Local labor market (LLM) FE	Yes	Yes	Yes
Observations	759,612	759,612	758,597
Firms	328,462	328,462	328,163
Local labor markets	14,918	14,918	14,870

Note: This table shows robustness to wage and tariff shock definition of estimates of the cross-firm within-market inverse elasticity of substitution, estimated using equation 1.15. The main specification uses firm wage premia as the wage variable and import tariffs as the shock.

Table A.6: Estimates of the inverse elasticity cross-market substitution $\frac{1}{\theta}$: robustness to sample

	Main specification (1)	All firms (2)
Panel A: First stage		
Δ LLM employment index	-0.582	-0.607
	(0.159)	(0.266)
First stage F	13.364	5.212
Panel B: Reduced form		
Δ LLM wage premium index	-0.049	-0.049
	(0.127)	(0.127)
Panel C: 2SLS		
$\frac{1}{\theta} - \frac{1}{\eta}$	0.084	0.081
	(0.232)	(0.233)
Panel D: Cross-market inverse elasticity of substitution		
$\frac{1}{\theta}$	0.627	0.623
	(0.105)	(0.107)
Implied lower bound on wage take-home share	61%	62%
Observations (Local labor markets)	13,965	13,965

Note: This table shows robustness to estimation sample of estimates of the cross-market inverse elasticity of substitution, estimated using equation 1.20. The main specification includes firms that are unique producers in their local labor markets, plus non-tradable sector firms.

Table A.7: Estimates of the inverse elasticity cross-market substitution $\frac{1}{\theta}$: robustness to clustering

	Main specifications		Additional trend controls	
	Level effect	De-trended effect	Level effect	De-trended effect
	(1)	(2)	(3)	(4)
Panel A: Labor market concentration				
Δ Log Wage Bill Herfindahl (based on wage premia)	0.432	0.630	0.507	0.706
	(0.294)	(0.294)	(0.348)	(0.348)
Δ Log Wage Bill Herfindahl	0.490	0.666	0.607	0.783
	(0.292)	(0.292)	(0.325)	(0.325)
Δ Log Employment Herfindahl	0.067	0.340	0.159	0.433
	(0.193)	(0.193)	(0.275)	(0.275)
Panel B: Log number of firms and log employment				
Δ Log number of firms	-0.216	-0.834	0.516	-0.102
	(0.418)	(0.418)	(0.362)	(0.362)
Δ Log total employment	-0.939	-2.177	0.088	-1.150
	(0.445)	(0.445)	(0.436)	(0.436)
Δ Exporter log employment	0.629	1.016	-0.075	0.312
	(0.356)	(0.356)	(0.524)	(0.524)
Δ Non-exporting tradables log employment	-7.522	-9.997	-4.037	-6.513
	(3.159)	(3.159)	(3.997)	(3.997)
Δ Non-tradables log employment	-0.152	-0.793	1.551	0.911
	(0.545)	(0.545)	(0.738)	(0.738)
Panel C: Log wage premium				
Δ Log wage premium	-0.660	-0.889	-0.119	-0.348
	(0.175)	(0.175)	(0.175)	(0.175)
Local Labor Market FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	No	No
Microregion x Year FE	No	No	Yes	Yes
Occupation x Year FE	No	No	Yes	Yes
Outcome is residual from regression on 1986-1990 trend	No	Yes	No	Yes
Observations	337,650	337,650	337,650	337,650

Note: This table shows robustness to clustering of estimates of the cross-market inverse elasticity of substitution, estimated using equation 1.20. Standard errors in the main specification are clustered at the microregion level.

A.2 Chapter 1 Data and Methods Appendix

Data on workers and firms: RAIS

Overview. I use Brazil's Relação Anual de Informações Sociais (RAIS) for years 1986 to 2000 as my source of information on workers and firms. RAIS is an administrative employer-employee linked dataset collected by the federal government for the purposes of administering workers' social security. Thus, RAIS covers all workers with signed worker cards (Carteira do Trabalho), namely the entirety of formal sector employment. Firms report RAIS once a year, reporting all workers who ever worked for the firm in the prior calendar year. Firms are required to report a rich set of information about each employment contract (e.g. occupation, admission date, separation date, etc.), as well as worker demographics (i.e. education, date of birth, and gender), separate by each establishment. The municipality of each establishment as well as the economic sector of the firm are also reported.

Wages. RAIS includes two wage variables for years 1986-2000: average monthly earnings and December monthly earnings. Both variables are reported as multiples of the national minimum wage. I match this data to IPEA (Instituto de Pesquisa Econômica Aplicada)'s real minimum wage series using dataset "Supplemental_materials_minwage_ipeadata[25-09-2017-04-00].csv", downloaded from IPEA's website on September 25, 2017.

Occupation codes. RAIS' occupation codes are 5-digit variables "CBO" (prior to 1994) and "CBO94" (1994 onwards). I focus on the first 2 digits to group workers into occupation groups. Both variables share the same data dictionary, with the only difference between them being phased-out and phased-in occupation codes. Supplementary material "Supplemental_materials_CBO94.xlsx" provides a complete list of all raw occupation codes, along with the total number of workers in each of them, labels, and flags for which codes were either "phased-out" or phased-in, which I identified based on whether the number of workers changing by more than 100 times between any two years. I then re-classify the first two digits of all phased-out and phased-in codes as "99 - Other occupations", a reclassification that affects roughly 10% of all workers.

Sector codes. RAIS' finest sector codes for 1986-2000 are 4-digit "IBGESUBATIVIDADE" (prior to 1995) and 5-digit "CNAE95" (1995 onwards). I focus on the 5-digit CNAE95 codes to map tariff shocks to firms in RAIS, For firms that exit the data prior to reporting any CNAE95 codes, I assign a CNAE95 code using the correspondence table "rais_ibgesubactivitytocnae10.dta," which I constructed using the pre-1995 and post-1995 codes of firms in business in both periods. To each IBGESUBATIVIDADE code I assign the most commonly reported CNAE95 code. Finally, throughout all years I use the first CNAE95 code ever reported by a firm as its official CNAE95 code.

Sample restrictions. I focus on workers employed as of December 31 of each year, and aged 18-65, and with positive December earnings. I exclude all public sector workers (i.e. "IBGESUBSETOR" 24). To make sure all government sector workers are excluded, I further a) exclude workers whose employer's economic activity was not marked as government, but which exert public sector occupations (i.e. 2-digit occupation codes 22 - Diplomats, 31 - Civil servants, and 37 - Post office); and b) focus only on workers whose labor contracts are not covered under the private sector's Consolidation of Labor Laws (i.e. contract types "1" prior to 1994 and "10" thereafter). Finally, following Dix-Carneiro and Kovak (2017) I exclude from all analyses the free trade zone of Manaus (i.e. microregion code 13007).

Data on tariff shocks: TRAINS

I use tariff data from UNCTAD's Trade Analysis Information System (TRAINS), which I download from the World Integrated Trade Solution (WITS)'s website: `https://wits.worldbank.org/`. I focus on the raw tariff data are available for Brazil at the 8-digit HS product level for years 1988 (the first year the data are available) through 2000. As outlined in Section 1.2, I compute a firm's tariff reduction shock is the change in log one plus a firms' CNAE95 sector code's nominal tariff between years 1990 and 1994. To map the product-level data to CNAE95, which is an economic activity code, I correspondence tables provided as supplemental materials to this paper. These tables are: a) correspondences between 8-digit product-level HS codes and 4-digit economic activity codes ISIC version 3.1 for each year, downloaded from WITS; b) correspondences between ISIC version 3.1 and CNAE95, downloaded from Brazil's Comissão Nacional de Classificação (CONCLA) website: `https://concla.ibge.gov.br/classificacoes/correspondencias/atividades-economicas`. The resulting dataset, also provided as supplemental materials, is a dataset of annual nominal tariffs. CNAE95 level-tariffs are then computed as simple averages of nominal tariffs across all product codes. For robustness exercises, I also compute each CNAE95's effective rate of protection (ERP), which net out the effect of tariffs on inputs. I calculate ERPs using Brazil's 1985 intersectoral tecnical coefficients matrix ("Tabela 20"), which is available at Nível 50 from Brazil's national accounts website: `https://www.ibge.gov.br/estatisticas-novoportal/economicas/contas-nacionais/9085-matriz-de-insumo-produto.html?&t=downloads`.

Other data

List of exporters. I classify firms as exporters during the reform period (1990-1994) by matching the list of exporters during that period to RAIS using firms' unique identifiers (CNPJ). The list of exporters was provided by the (extinct as of 2019) Ministry of Development, Industry, and Foreign Trade, currently a part of the Ministry of the Economy, in October 2018.

Census. I produced Appendix Figure A.16 by using data on informality at the microregion level from the 1990 and 2000 census, which I obtained from Dix-Carneiro and Kovak (2017)'s supplemental materials.

Methods: wage premia regressions

I estimate firm wage premia as firm fixed effects from a regression of log december earnings on firm fixed effects and the same controls as Dix-Carneiro and Kovak (2017), namely: a dummy for female; 4 age group dummies (25-29; 30-39; 40-49,50-64); 8 education group dummies (primary school, incomplete primary school, middle school, incomplete middle school, high school, incomplete high school, college, incomplete college). The omitted category is therefore males aged 18-24 with no formal education. I estimate local labor market wage premia as local labor market fixed effects from a regression of worker log december earnings on local labor market fixed effects and the same demographic controls as in the firm wage premia regressions.

A.3 Chapter 1 Model Appendix

Setup

I study the effect of import tariff reductions in a two-sector small "protected" (that is, open with tariffs) economy.

Household. A representative household takes prices, wages, and unearned income as given and solves:

$$\max_{\{c_s\},\{l_{sj}\}} \left(\sum_{s\in\{b,c\}} c_s^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}} - \left(\sum_{s\in\{b,c\}} \int_{j\in\Omega_s} l_{sj}^{\frac{\varepsilon+1}{\varepsilon}} dj\right)^{\frac{\varepsilon}{\varepsilon+1}}$$
$$\text{s.t.} \quad \sum_{s\in\{b,c\}} p_s c_s \leq \sum_{s\in\{b,c\}} \int_{j\in\Omega_s} w_{sj} l_{sj} dj + Z$$

where $\sigma > 1, \varepsilon > 0$ are elasticities of product demand and labor supply, respectively; c_b, c_c are consumption of goods b ("bananas") and c ("cell phones"), respectively; l_{sj} and w_{sj} are the labor allocated to and wage paid by firm j in sector s; p_s is the price of good $s \in \{b, c\}$ faced by the household; Ω_s is the set of domestic firms producing in sector s; and

$$Z = \sum_{s\in\{b,c\}} \int_{j\in\Omega_s} \pi_{sj} dj + \tau p_c I_c \tag{A.1}$$

is unearned income (total firm profits plus tariff revenues from protecting the cell phone sector). Since this is a small "protected" economy, prices faced by households and firms are given by world prices and import tariffs. Since the cell phone sector is protected, prices are:

$$p_s = \begin{cases} p_b & \text{if } s = b \\ p_c(1+\tau) & \text{if } s = c \end{cases}$$

where p_b,p_c are world prices for bananas and cell phones, respectively, and $\tau > 0$ is the import tariff. The FOCs give the expressions for good demand and labor supply:

$$c_s = \left(\frac{p_s}{P}\right)^{-\sigma} C, \quad s \in \{b, c\} \tag{A.2}$$
$$l_{sj} = \left(\frac{w_{sj}}{P}\right)^{\varepsilon} L, \quad \forall j, s \in \{b, c\} \tag{A.3}$$

where $P = \left(\sum_{s\in\{b,c\}} p_s^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ is the CES price index, $C = \left(\sum_{s\in\{b,c\}} c_s^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}$, and $L = \left(\sum_{s\in\{b,c\}} \int_{j\in\Omega^s} l_{sj}^{\frac{\varepsilon+1}{\varepsilon}} dj\right)^{\frac{\varepsilon}{\varepsilon+1}}$.

Firms. Firms produce with CRS technology using labor as the only input of production.

They take goods prices as given and compete monopsonistically for labor. Each firm solves:

$$\begin{aligned} \max_{q_{sj}, l_{sj}} \quad & p_s q_{sj} - w_{sj} l_{sj} \\ \text{s.t.} \quad & q_{sj} = \varphi_{sj} l_{sj} \end{aligned}$$

where $\varphi_{sj} = \gamma_s \xi_j$, $\xi_j \sim G(\cdot)$, with $\gamma_b > \gamma_c$ such that productivity in the banana sector is on average higher than productivity in the cell phone sector.[1] FOCs give the firms' wage setting equation:

$$w_{sj} = \frac{p_s \varphi_{sj}}{1 + \varepsilon^{-1}} \tag{A.4}$$

Market clearing. Finally, market clearing requires

$$l_{sj}^D = l_{sj}^S \quad \forall j, s \tag{A.5}$$

$$c_s = \int_{j \in \Omega_b} q_{bj} d_j + I_s - E_s, \quad s \in \{b, c\} \tag{A.6}$$

where l_{sj}^D indicates labor demand at firm j in sector s, and l_{sj}^S labor supply. I_s is imports of good s and E_s is exports of good s. Since by assumption the country has a comparative advantage in producing bananas, it exports bananas and imports cell phones, such that

$$I_b = 0 \tag{A.7}$$

$$E_c = 0 \tag{A.8}$$

Solution. We now count equations and unknowns. While we have eight independent equations (ignoring the s, j multipliers), we have nine unknowns: $c_s, q_{sj}, l_{sj}, w_{sj}, \forall j, s$, and C, L, E_b, I_c, Z. I therefore close the model by making one additional assumption, namely that aggregate labor supply is fixed at a constant $\mathcal{L}$:

$$\mathcal{L} = \sum_{s \in \{b,c\}} \int_{j \in \Omega_s} l_{sj} dj$$

We could now in principle derive expressions for each unknown as a function of model parameters, world prices, and the import tariff only. Instead, I focus attention on the

[1] In other words, the productivity distribution of the banana sector has the same shape as the distribution of the cell phone sector, but it's shifted by a constant.

solutions for wage and employment as a function of CES indices, model parameters, world prices, and import tariffs. These expressions are sufficient to derive predictions for the effect of a change in tariffs on labor market concentration.

Trade and labor market concentration

I study the comparative statics of a decrease in import tariffs on the wage bill Herfindahl of the small open economy described in Section A.3. In the context of this economy, the wage bill Herfindahl is defined as

$$HHI = \sum_{s\in\{b,c\}} \int_{j\in\Omega_s} \left(\frac{w_{sj}l_{sj}}{WL}\right)^2 dj$$

where $WL = \int_{j\in\Omega_s} w_{sj}l_{sj}dj$ is the market-level wage bill.

Special case. I start with the special case where $\xi_j = 1\forall j \implies \varphi_{sj} = \gamma_s\forall j \implies w_{sj} = w_s, l_{sj} = l_s\forall j$. In this case:

$$\begin{aligned} HHI &= \sum_{s\in\{b,c\}} \int_{j\in\Omega_s} \left(\frac{w_{sj}l_{sj}}{WL}\right)^2 dj \\ &= \int_{j\in\Omega_b} \left(\frac{w_bl_b}{WL}\right)^2 dj + \int_{j\in\Omega_c} \left(\frac{w_cl_c}{WL}\right)^2 d \\ &= \left(\frac{w_bl_b}{WL}\right)^2 \int_{j\in\Omega_b} dj + \left(\frac{w_cl_c}{WL}\right)^2 \int_{j\in\Omega_c} d \\ &= \left(\frac{w_bl_b}{WL}\right)^2 + \left(\frac{w_cl_c}{WL}\right)^2 \end{aligned} \tag{A.9}$$

where the last line follows by the assumption that each sector as a unit mass of firms. Note that HHI defines a circle of radius $\sqrt{HHI}$ centered at $(0,0)$ on the $\frac{w_bl_b}{WL} \times \frac{w_cl_c}{WL}$ plane, whose level sets can be plotted as shown in Figure 1. To see how a decrease in τ changes the HHI, I next lay over this graph the identity that the sum of sectoral employment shares must equal one:

$$\frac{w_bl_b}{WL} + \frac{w_cl_c}{WL} =1 \tag{A.10}$$

Finally, I use the model solution to identify the equilibria HHI before and after the tariff decrease using equations A.9 and A.10. Combining equations A.3 and A.4 gives

$$\frac{w_c l_c}{w_b l_b} = (1+\tau)\left(\frac{p_c \gamma_c}{p_b \gamma_b}\right)^{\varepsilon+1} \tag{A.11}$$

Thus, a tariff decline reduces $\frac{w_c l_c}{w_b l_b}$. Figure 1 shows that the implication of this reduction to HHI depends on in which HHI level set the economy is in the pre-reform equilibrium. To locate the pre-reform equilibrium HHI, I consider the simpler case of infinitely high tariffs (i.e. autarky). Solving the model in Section A.3 under autarky gives the following relationship between the wage bill, prices, and productivity:

$$\frac{w_c^a l_c^a}{w_b^a l_b^a} = \left(\frac{p_c^a \gamma_c}{p_b^a \gamma_b}\right)^{\varepsilon+1} \quad \text{and} \quad \frac{p_c^a}{p_b^a} = \left(\frac{1/\gamma_c}{1/\gamma_b}\right)^{\frac{\varepsilon+1}{\varepsilon+\sigma}}$$

Combining:

$$\frac{w_c^a l_c^a}{w_b^a l_b^a} = \left(\frac{\gamma_c}{\gamma_b}\right)^{\varepsilon+1-\frac{1}{\varepsilon+\sigma}} \tag{A.12}$$

Since $\gamma_c < \gamma_b$ and $\varepsilon > 0, \sigma > 1$, it follows that $\frac{w_c^a l_c^a}{w_b^a l_b^a} < 1$. This information thus places the line pinning down the autarky equilibrium HHI somewhere below the 45-degree line in Figure 1. This positioning is now sufficient to show that a tariff reduction increases HHI because the pre-reform c-to-b wage bill ratio is necessarily smaller than the autarky c-to-b wage bill ratio. So we now know that the pre-reform c-to-b wage bill ratio is below the 45-degree line. The last step is easily seen in Figure 1: a tariff reduction further reduces the wage bill ratio, moving the HHI to a higher level set. *Therefore, when a country has absolute and comparative disadvantage in the protected good, a unilateral tariff reduction increases labor market concentration.*

General case. More generally, from the model solutions we obtain

$$w_s l_s = \left[\frac{p_s \varphi_{sj}}{P\left(1+\varepsilon^{-1}\right)}\right]^{\varepsilon+1} L$$

From which:

$$
\begin{aligned}
HHI &= \sum_{s\in\{b,c\}} \int_{j\in\Omega_s} \left(\frac{w_{sj} l_{sj}}{WL}\right)^2 dj \\
&= \sum_{s\in\{b,c\}} \int_{j\in\Omega_s} \left(\frac{\left[\frac{p_s \varphi_{sj}}{P(1+\varepsilon^{-1})}\right]^{\varepsilon+1} L}{WL}\right)^2 dj \\
&= \frac{1}{(WL)^2}\left\{\underbrace{L^2\left[\frac{p_b}{P\left(1+\varepsilon^{-1}\right)}\right]^{2(\varepsilon+1)} \int \varphi_{bj}^{2(\varepsilon+1)} dj}_{\equiv\left(\tilde{w}_b\tilde{l}_b\right)^2} + \underbrace{L^2\left[\frac{p_c\left(1+\tau\right)}{P\left(1+\varepsilon^{-1}\right)}\right]^{2(\varepsilon+1)} \int \varphi_{cj}^{2(\varepsilon+1)} dj}_{\equiv\left(\tilde{w}_c\tilde{l}_c\right)^2}\right\} \\
&= \left(\frac{\tilde{w}_b\tilde{l}_b}{WL}\right)^2 + \left(\frac{\tilde{w}_c\tilde{l}_c}{WL}\right)^2
\end{aligned}
$$

which again describes a circle of radius $\sqrt{HHI}$ centered at $(0,0)$ on the $\frac{\tilde{w}_b\tilde{l}_b}{WL} \times \frac{\tilde{w}_c\tilde{l}_c}{WL}$ plane. In addition, the identity for the sum of sectoral employment shares now follows

$$
\frac{\tilde{w}_b\tilde{l}_b}{WL} + \frac{\tilde{w}_c\tilde{l}_c}{WL} = 1
$$

And the wage bill ratios are given by

$$
\begin{aligned}
\frac{\tilde{w}_c\tilde{l}_c}{\tilde{w}_b\tilde{l}_b} &= \left(\frac{L^2\left[\frac{p_c(1+\tau)}{P(1+\varepsilon^{-1})}\right]^{2(\varepsilon+1)} \int \varphi_{cj}^{2(\varepsilon+1)} dj}{L^2\left[\frac{p_b}{P(1+\varepsilon^{-1})}\right]^{2(\varepsilon+1)} \int \varphi_{bj}^{2(\varepsilon+1)} dj}\right)^{\frac{1}{2}} \\
&= \left(\frac{p_c\left(1+\tau\right)}{p_b}\right)^{\varepsilon+1} \left(\frac{\gamma_c^{2(\varepsilon+1)} \int \xi_j^{2(\varepsilon+1)} dj}{\gamma_b^{2(\varepsilon+1)} \int \xi_j^{2(\varepsilon+1)} dj}\right)^{\frac{1}{2}} \\
&= \left(\frac{p_c\gamma_c\left(1+\tau\right)}{p_b\gamma_b}\right)^{\varepsilon+1} \underbrace{\left(\frac{\int \xi_j^{2(\varepsilon+1)} dj}{\int \xi_j^{2(\varepsilon+1)} dj}\right)^{\frac{1}{2}}}_{=1}
\end{aligned}
$$

from which it follows that, as in the special case of identical idiosyncratic productivity, the c-to-b wage bill ratio is decreasing in tariffs. Completing the proof now requires placing the autarky equilibrium c-to-b wage bill ratio below the 45 degree line. Solving for the autarky price ratio gives:

$$
\begin{aligned}
\frac{p_c^a}{p_b^a} &= \left(\frac{\int \varphi_{bj}^{\varepsilon+1} dj}{\int \varphi_{cj}^{\varepsilon+1} dj} \right)^{\frac{1}{\varepsilon+\sigma}} \\
&= \left(\frac{1/\gamma_c}{1/\gamma_b} \right)^{\frac{\varepsilon+1}{\varepsilon+\sigma}} \left(\frac{\int \xi_j^{\varepsilon+1} dj}{\int \xi_j^{\varepsilon+1} dj} \right)^{\frac{1}{\varepsilon+\sigma}} \\
&= \left(\frac{1/\gamma_c}{1/\gamma_b} \right)^{\frac{\varepsilon+1}{\varepsilon+\sigma}}
\end{aligned}
$$

Solving for the autarky sectoral wage bill ratio equivalent gives:

$$
\frac{\tilde{w}_c^a \tilde{l}_c^a}{\tilde{w}_b^a \tilde{l}_b^a} = \left(\frac{p_c^a \gamma_c}{p_b^a \gamma_b} \right)^{\varepsilon+1} = \left(\frac{\gamma_c}{\gamma_b} \right)^{\varepsilon+1-\frac{1}{\varepsilon+\sigma}}
$$

which is the same wage bill ratio as in the simple case of identical idiosyncratic productivities.

Appendix B

Appendices for Chapter 2

B.1 Chapter 2 Appendix Figures and Tables

Figure B.1: Probability of MTO assignment

Panel A: As a function of match tax year (2005)
gross income

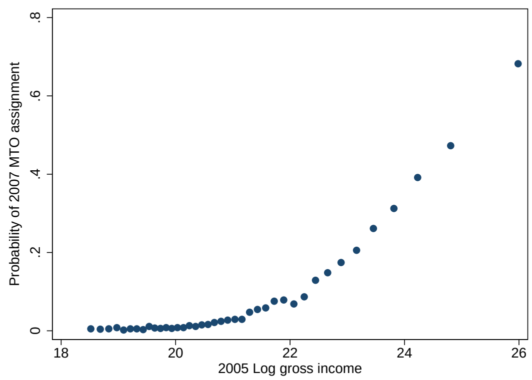

Panel B: As a function of match tax year (2005)
total taxes paid

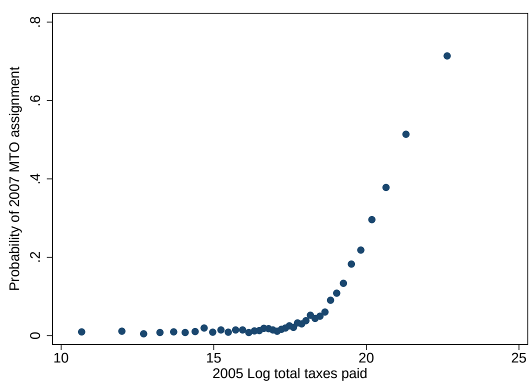

Notes: This figure displays the percent of taxpayers assigned to MTO in 2007 as a function of MTO assignment input variables (total taxes paid and gross income) for tax year 2005. Percentages are plotted against forty equal-sized bins of the 2005 log gross income and log total taxes paid distribution of taxpayers in eligible origin tax offices as of 2006.

Figure B.2: Joint distribution of taxpayer gross and taxable income

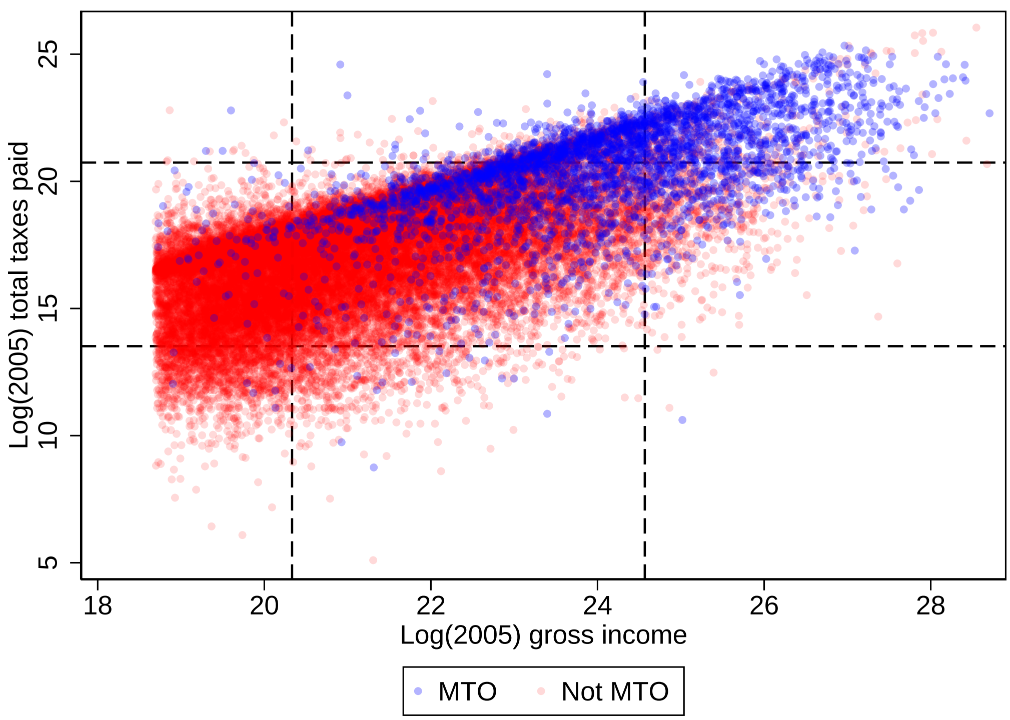

Notes: This figure shows the joint distribution of taxpayers' 2005 log gross income and 2005 log total taxes paid. Each blue dot is a taxpayer assigned to MTO's first cohort, while each red dot is a taxpayer not assigned to MTO's first cohort. Dotted black lines indicate the lower bound and upper bounds of the 2.5th-97.5th percentile common support.

Figure B.3: Common support restrictions for taxpayer size distributions

Panel A: Gross income

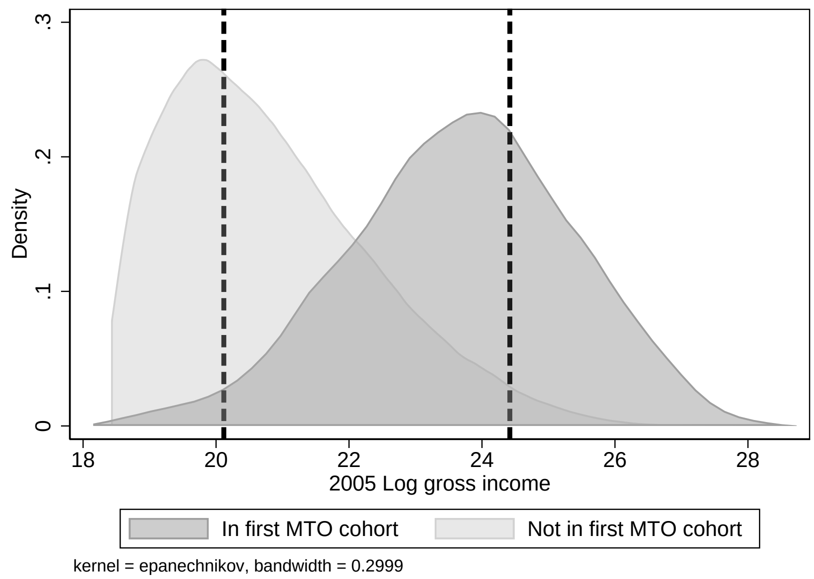

Panel B: Total taxes paid

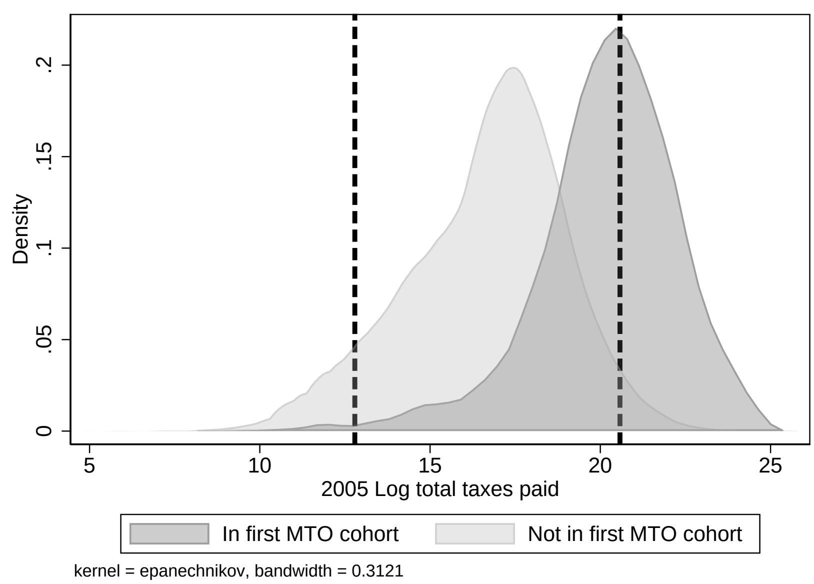

Notes: This figure shows the distributions of taxpayer 2005 log gross income and 2005 log total taxes paid by MTO 2007 assignment status. Dotted black lines indicate the lower bound and upper bound of the 2.5th-97.5th percentile common support.

Figure B.4: Taxpayer size distributions pre- and post- MTO creation

Panel A: Pre-MTO (tax year 2005)

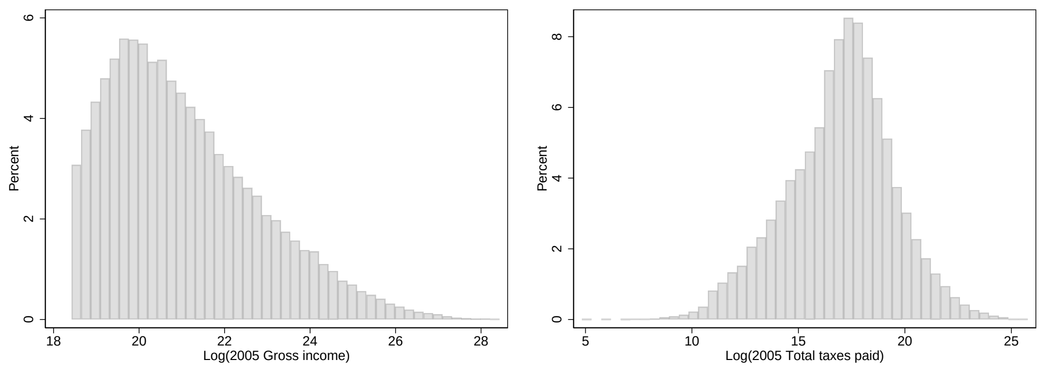

Panel B: Post-MTO (tax year 2011)

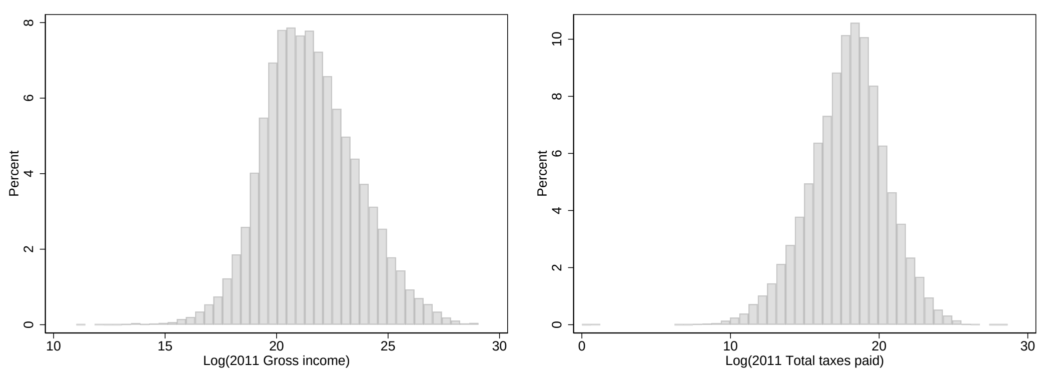

Notes: This figure shows the distributions of taxpayer log gross income and log total taxes paid before (tax year 2005) and after (tax year 2011) the creation of MTO. 2005 log gross income distribution is truncated at gross income sample restriction of IDR 100 million Rp (roughly USD 10,000).

Figure B.5: Effect of MTO first cohort assignment on year-by-year MTO treatment

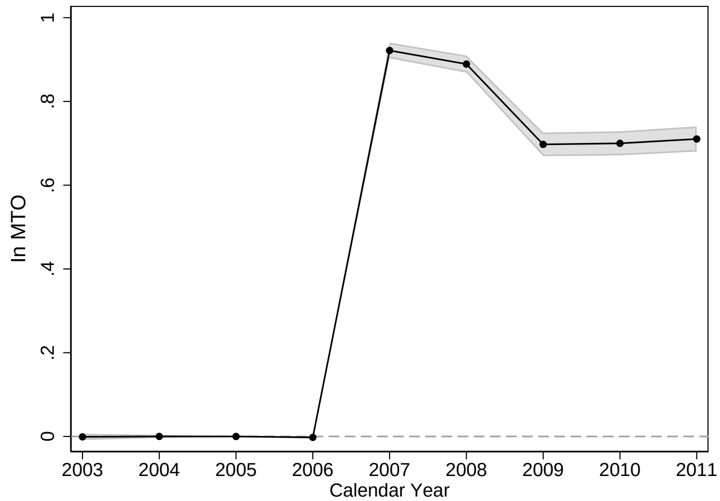

Notes: See notes to Figure 11.

Figure B.6: Effects on total taxes paid for MTO, larger non-MTO firms, and smaller non-MTO firms (weighted annual averages)

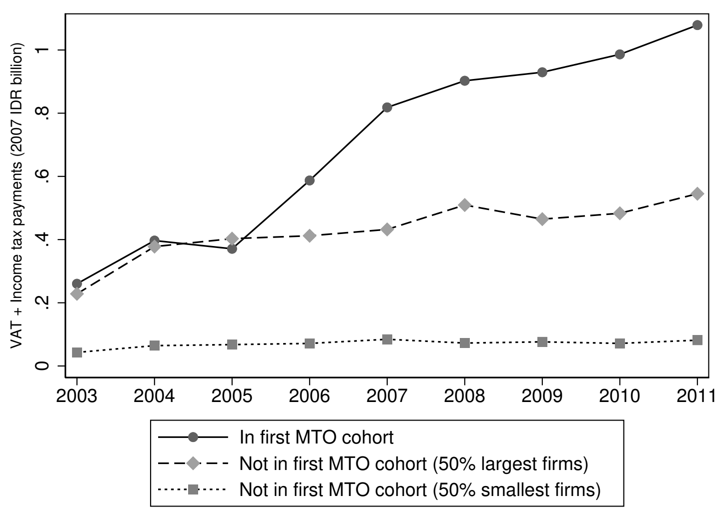

Notes: This figure shows weighted annual averages of total taxes paid, separately by taxpayer's assignment to the first MTO cohort. Data are weighted using entropy-balancing weights.

Figure B.7: MTO effect on detailed tax filing outcomes

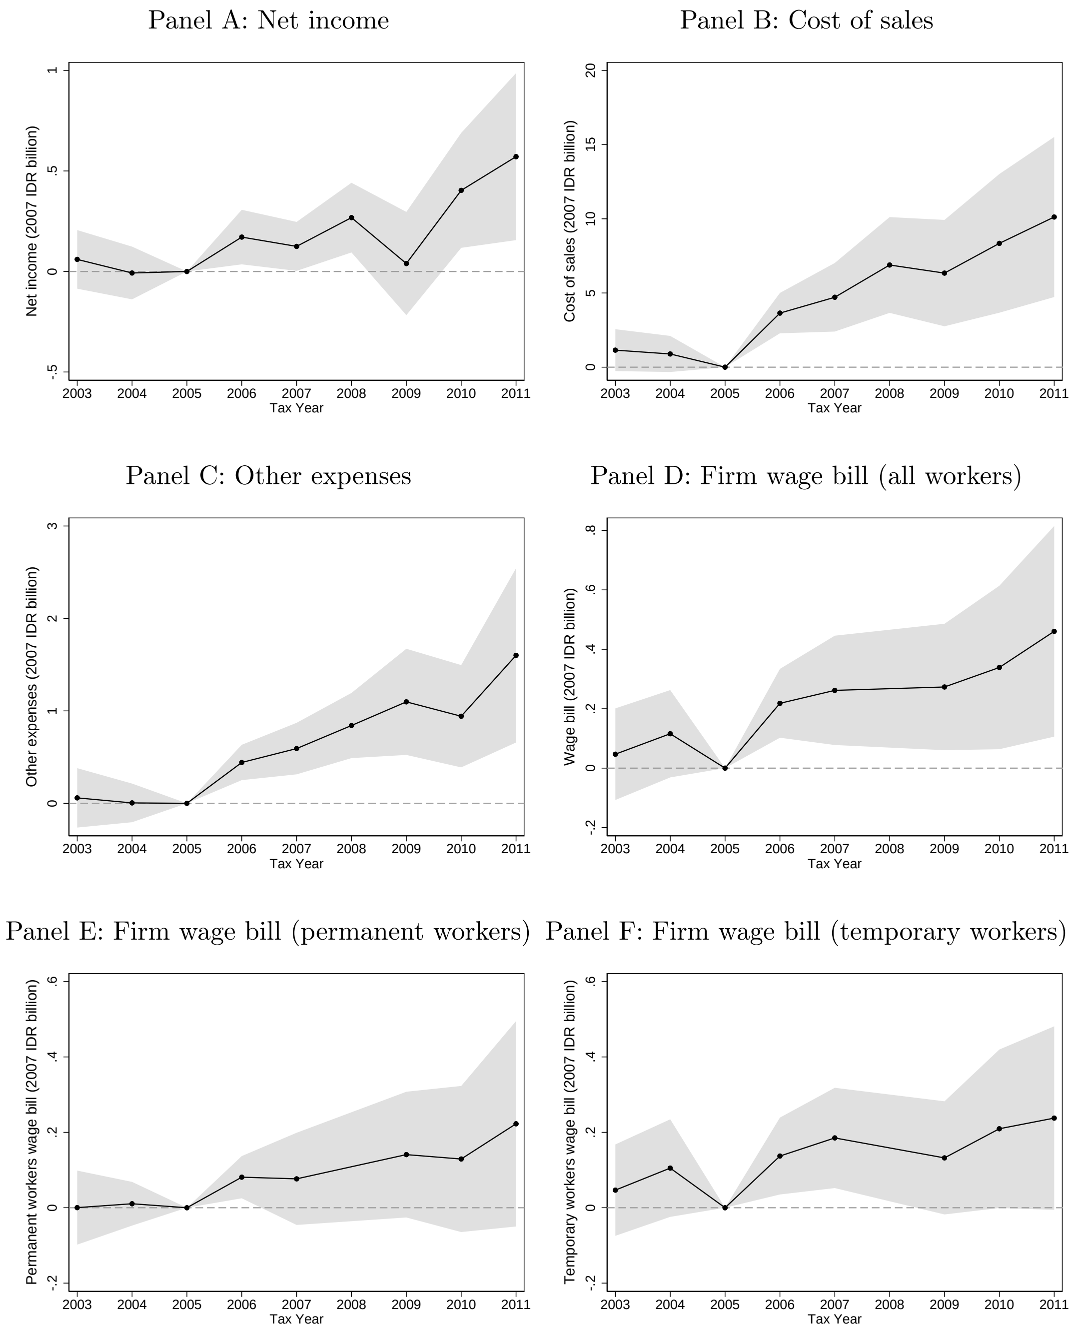

Notes: See notes to Figure 11, Table B.10, and Table 2.2.

Figure B.8: MTO effect on cost of sales, split by base year taxable income

Panel A: Taxpayers with zero taxable income in 2005

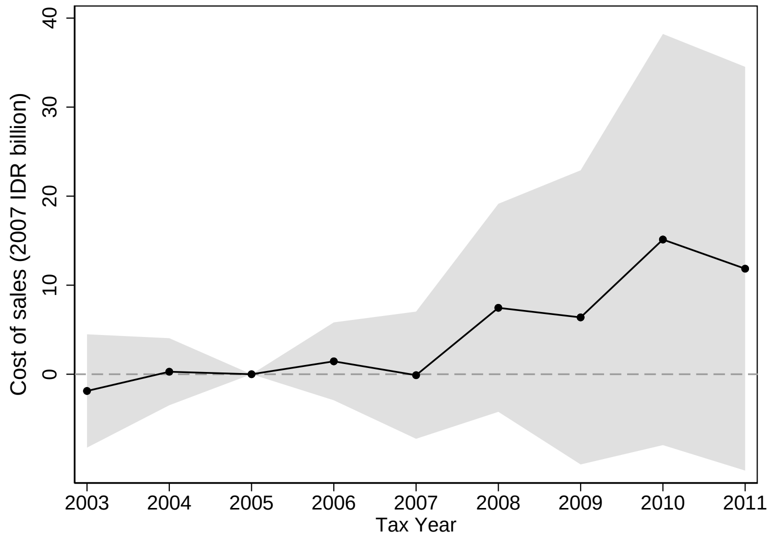

Panel B: Taxpayers with positive taxable income in 2005

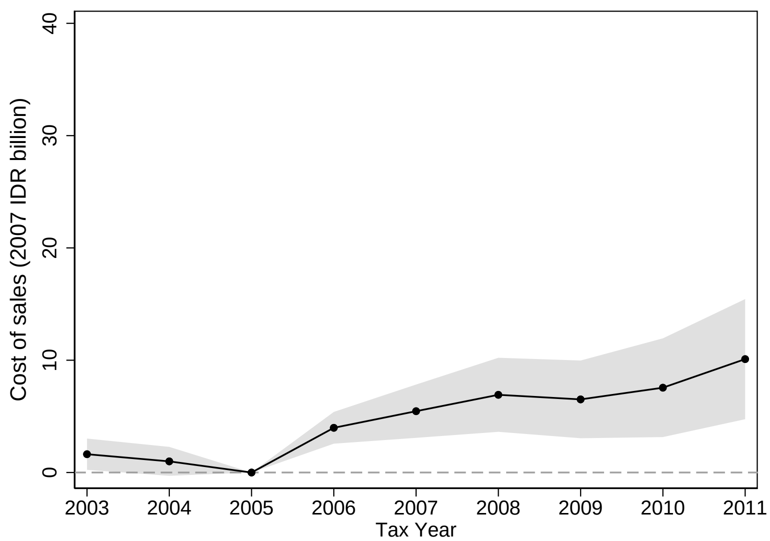

Notes: This figure shows year-by-year reduced form effects of MTO treatment on cost of sales by two groups: taxpayers with zero vs. positive 2005 taxable income.

Figure B.9: MTO effect on reported income: weighted average levels

Panel A: Gross income

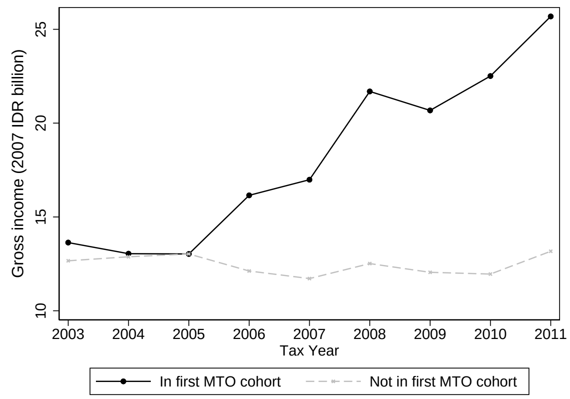

Panel B: Taxable income

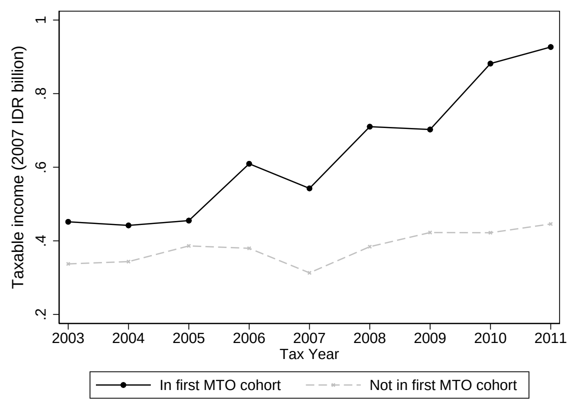

Panel C: Corporate income taxes due

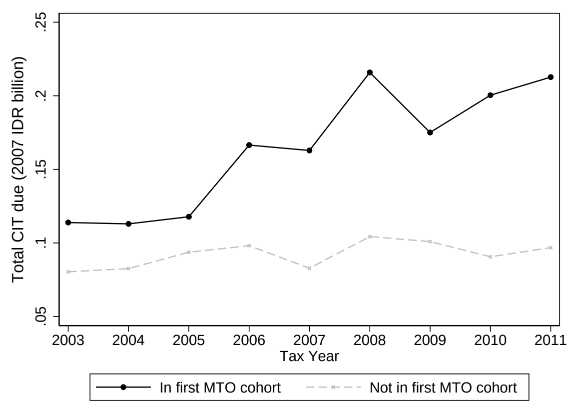

Notes: This figure shows weighted annual averages of reported gross income, taxable income, and total corporate income taxes due. Data are weighted using entropy-balancing weights.

Figure B.10: MTO effect on Employment

Panel A: All workers

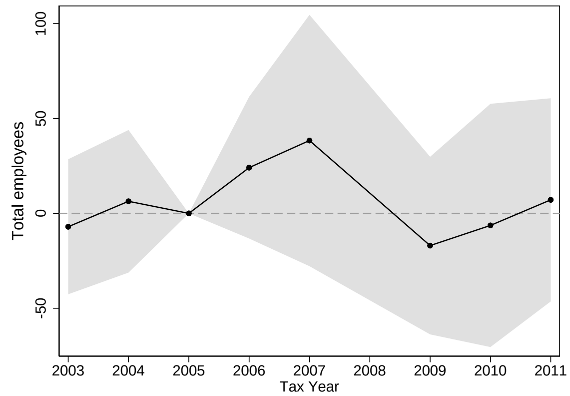

Panel B: Permanent

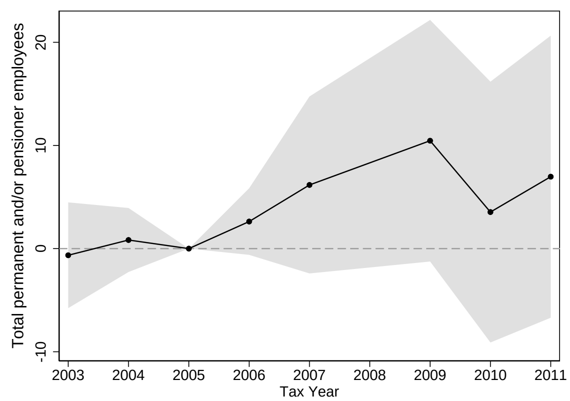

Panel C: Temporary

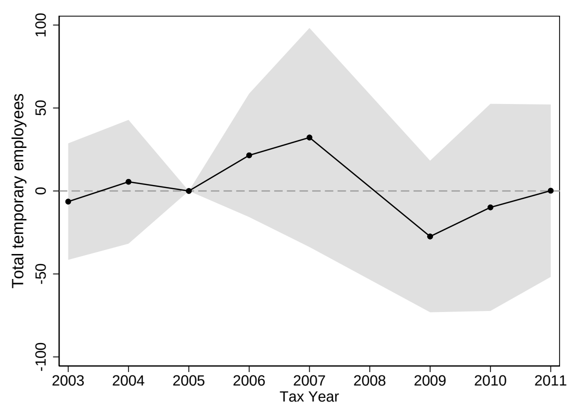

Notes: See notes to Figure 11. Firm employment data are from corporate employment tax withholding form SPT 1721. Employment data for tax year 2008 are not available. See Data Appendix for details.

Figure B.11: MTO effects, including MTOs started in 2005 and 2006

Panel A: Total Taxes Paid

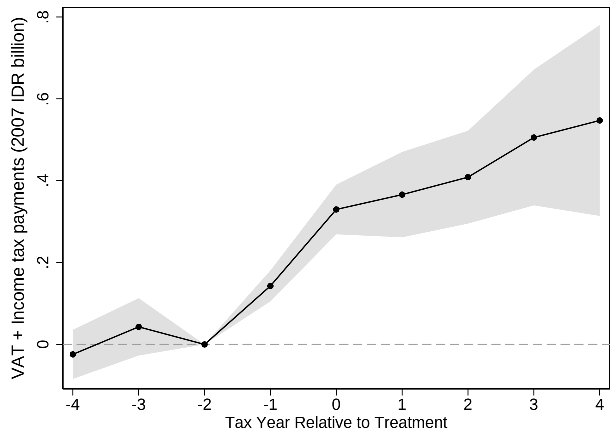

Panel B: Gross income

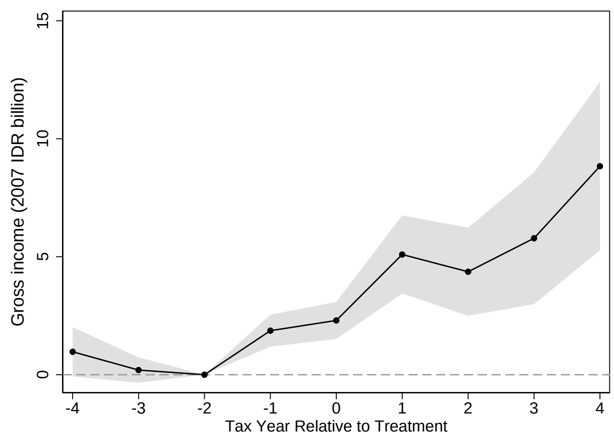

Notes: This figure shows year-by-year reduced form effects of MTO treatment on total taxes paid and gross income including the 5 MTOs created in 2005 and 2006 to the main sample of 13 MTOs created in 2007. Year-by-year effects are estimated relative to the year of MTO assignment by stacking the 2005, 2006, and 2007 MTO assignment cohorts, and slightly modifying equation (2.16) to be defined in relative years. In particular, year-by-year effects are coefficients on interactions of the MTO assignment dummy variable M_{i0} (equivalent to M_{iFC} in equation (2.16)) with year dummies, omitting the interaction and main effect dummies for base relative year -2 (the last tax year that filings would have been available to the tax office at the time of each MTO assignment in relative year zero). As MTO assignment occurred in different years, year fixed effects are also included. The stacked regression is weighted following the same balancing methodology as in Figure 11. Specifically, the weights used are taxpayer-specific and constructed by applying Hainmueller (2012)'s entropy-balancing methodology to the MTO assignment formula inputs (gross income and total taxes paid). The formula inputs are for tax year 2005 for the 2007 MTO cohort, tax year 2004 for the 2006 cohort, and tax year 2003 for the 2005 cohort. Taxpayer-level total taxes paid data are from the Treasury, and include payments from all branches of the same corporate entities. Reported income data are from tax filing form SPT 1771 (annual corporate income tax return), and are reported by the taxpayer headquarters on behalf of all branches of the same corporate entity. IDR values are deflated to 2007 IDR using Indonesia's GDP deflator. Solid lines are point estimates; dashed lines are 95% confidence intervals based on standard errors clustered at the taxpayer level.

Figure B.12: Probability of placebo treatment assignment among non-MTO taxpayers

Panel A: As a function of match tax year (2005) gross income

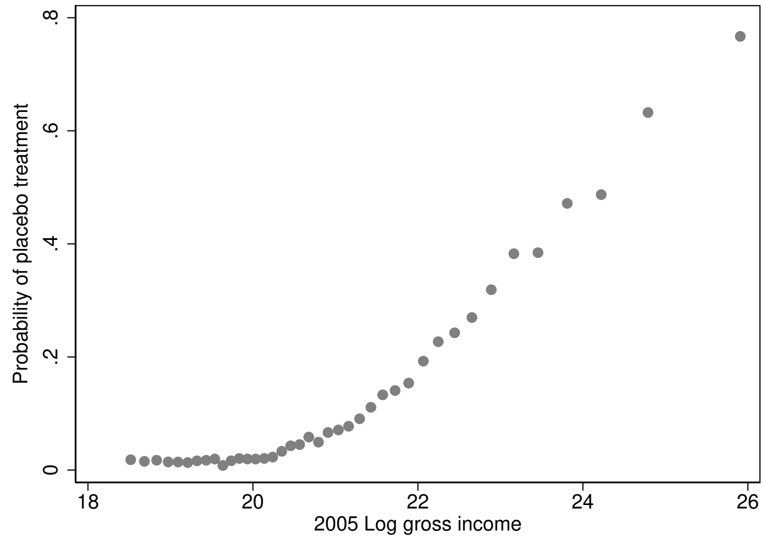

Panel B: As a function of match tax year (2005) total taxes paid

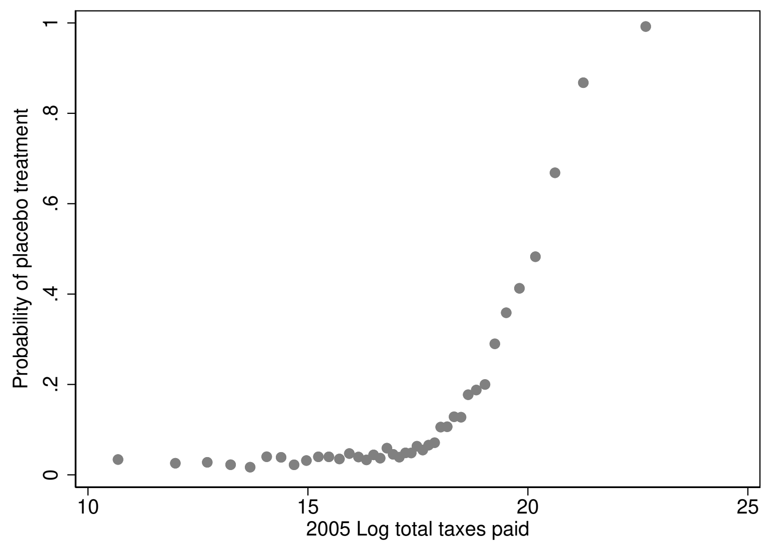

Notes: See Section 2.4.2. This figure displays the percent of non-MTO taxpayers assigned to a placebo treatment as a function of the placebo treatment input variables (total taxes paid and gross income) for tax year 2005.

Figure B.13: Placebo effect on Total taxes paid

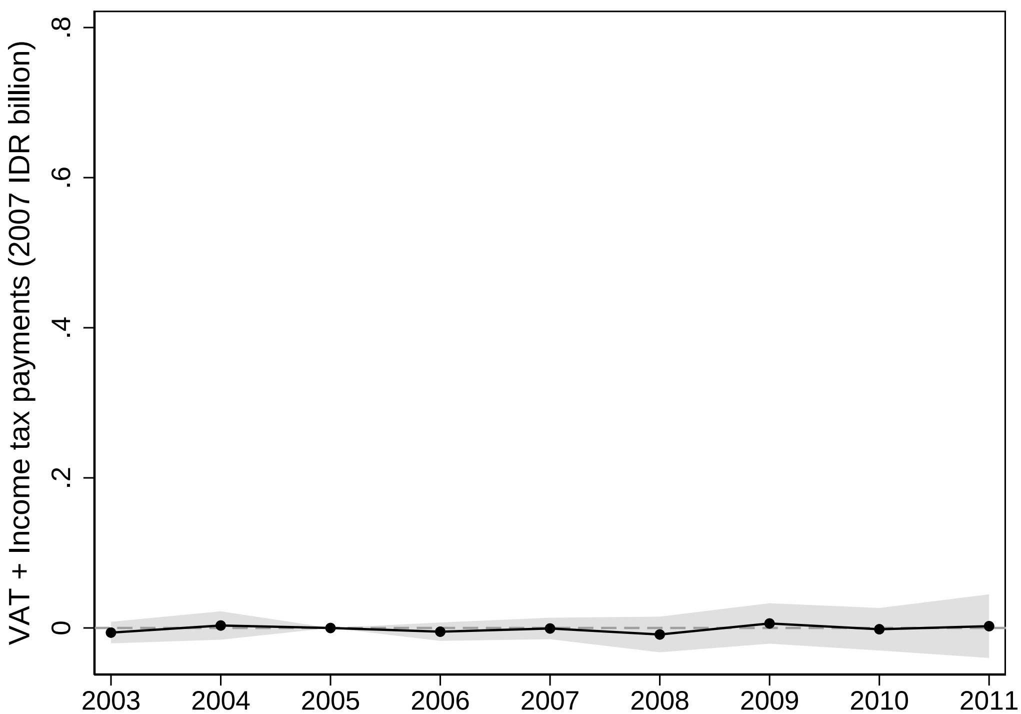

Notes: See notes to Figures 11 and B.12. Y-axis displays the same scale as the MTO effect on total taxes paid.

Figure B.14: Taxpayer density along MTR variation from 2009 corporate tax rate reform

Panel A: 2008-2009 schedule change

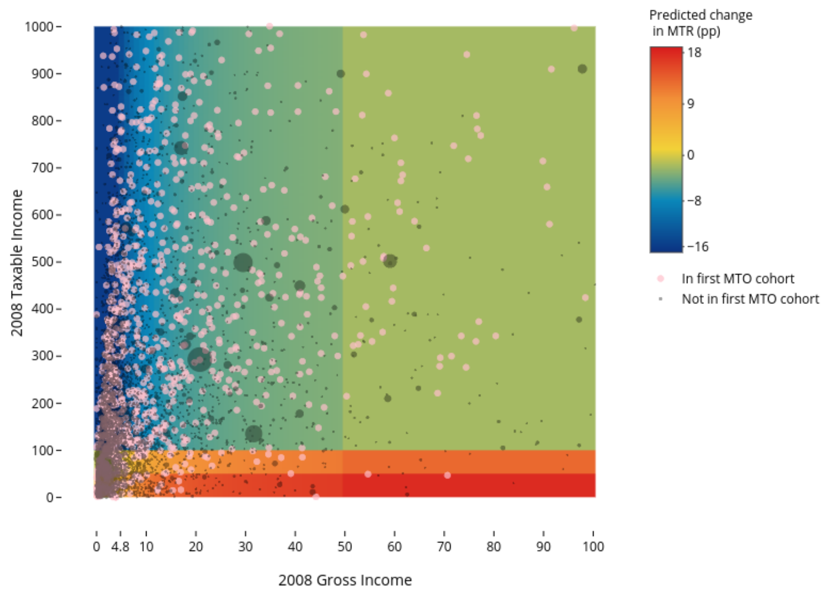

Panel B: 2009-2010 tax rate cut

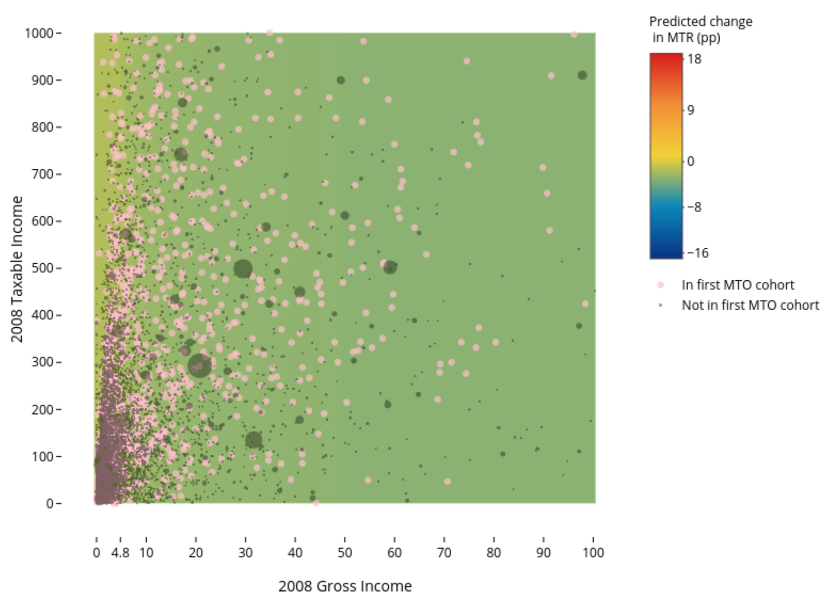

Notes: This figure displays the marginal tax rate change variation induced by Indonesia's 2009 tax rate reform within the ETI estimation analysis sample. Scatter plot marker sizes are proportional to taxpayer-specific entropy-balancing weights. See Section 2.5 for a detailed description of how predicted marginal tax rates are computed.

Figure B.15: MTO effect on corporate income tax marginal tax rate

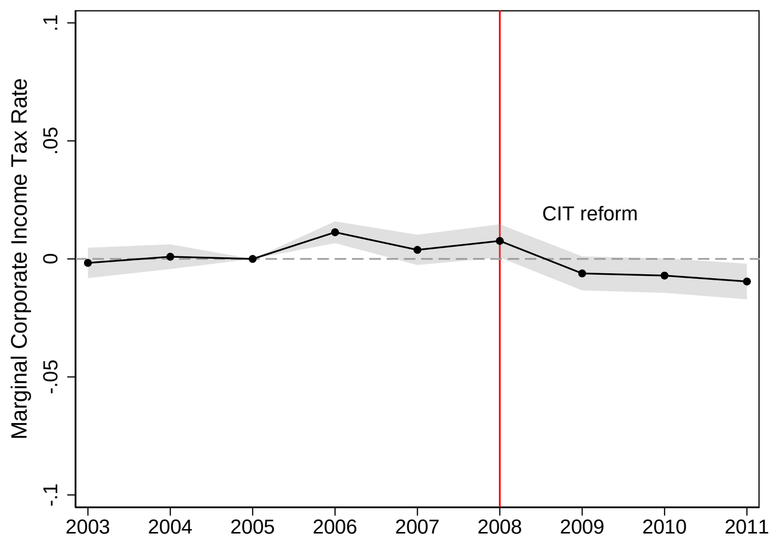

Notes: See notes to Figure 11. This figure plots year-by-year coefficients of the effect of MTO assignment on the marginal corporate income tax rate faced by taxpayers. Taxpayers' MTR are measured according to the MTR schedules presented in Section 2.2.1. The red line indicates the year of the MTR schedule reform.

Figure B.16: Effect of a large predicted tax cut in 2009 on marginal tax rates and log taxable income

Panel A: Marginal Tax Rates

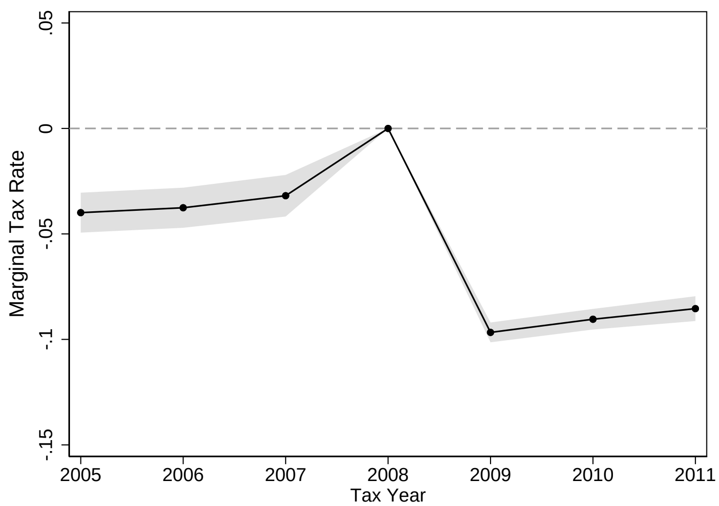

Panel B: Log taxable income

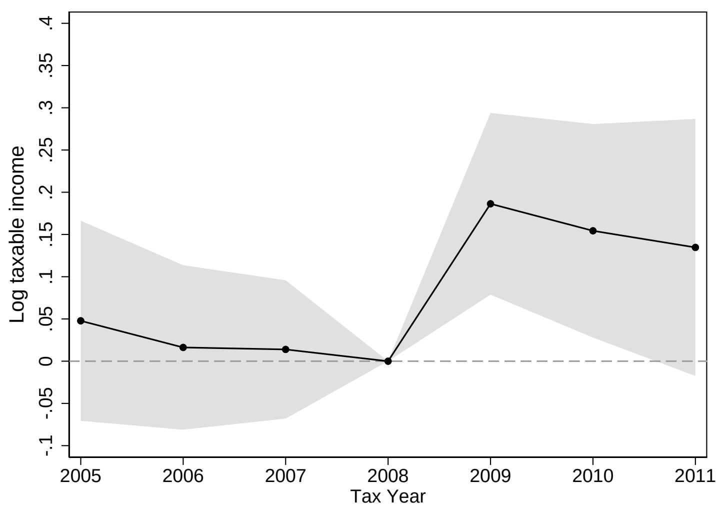

Notes: This figure presents regression coefficients from a regression of each outcome (i.e., marginal tax rate or log taxable income) on year dummies (omitting 2008) interacted with a dummy indicating whether the taxpayer was predicted to experience a large tax rate cut in 2009 (top 50% of tax rate cuts, with an average of 15 percentage points cut). The regression also includes taxpayer fixed effects and controls for the 2008 log gross income and 2008 log taxable income, each interacted with year dummies (and again, omitting 2008). The control group was predicted to experience a small tax cut (bottom 50% of tax cuts, with an average of 3 percentage points cut). The regression is thus conditional on taxpayers who were predicted to experience a tax cut in 2009, and is weighted by MTO balancing weights. Standard errors are clustered at the taxpayer level.

Figure B.17: Bunching at notch before and after MTR schedule reform

Panel A: After 2009 MTR schedule reform

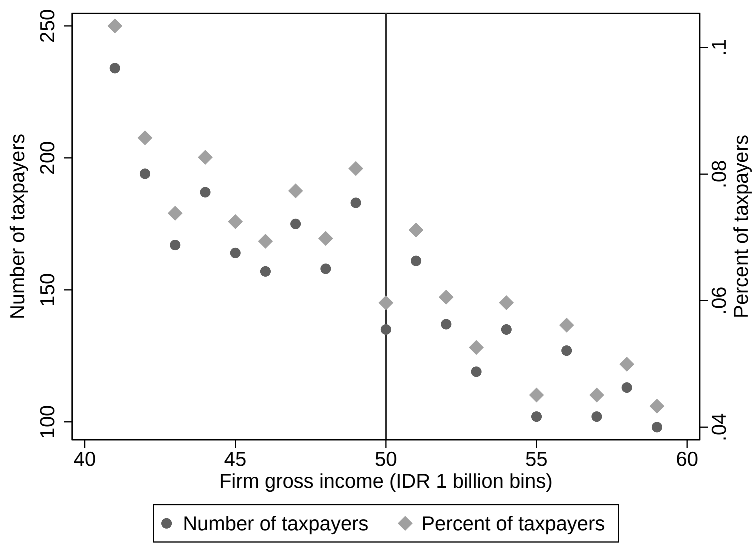

Panel B: Before 2009 MTR schedule reform

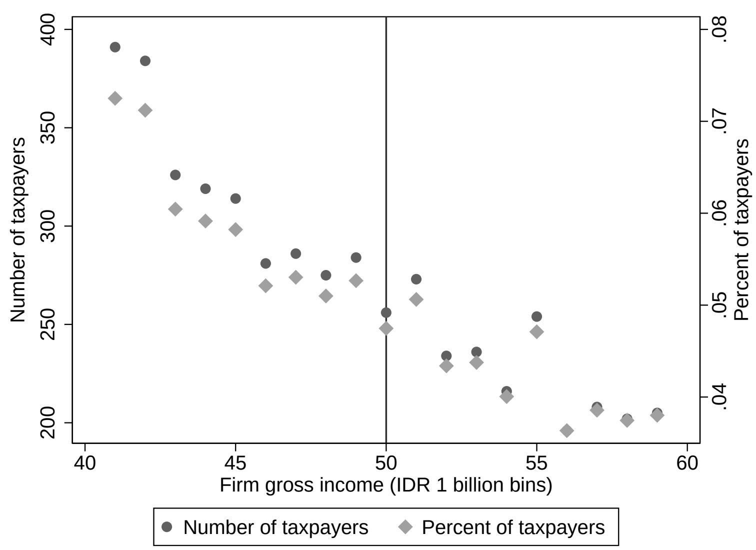

Notes: This figure shows taxpayer density at the IDR 50 billion notch introduced by the 2009 corporate income tax schedule. The sample includes data for tax years 2003-2011 for all corporate taxpayers with non-zero taxable income.

Table B.1: Tax Office Staffing

	MTO tax offices				Non-MTO tax offices			
	2008	2009	2010	2011	2008	2009	2010	2011
	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Taxpayers-to-staff ratios								
Taxpayers per Auditor	18	24	23	21	107	107	115	125
Taxpayers per AR	17	26	25	20	56	105	93	80
Taxpayers per staff	4	6	6	6	10	16	17	17
Auditors								
Total auditors	329	370	366	361	1,109	1,667	1,643	1,591
Has college degree	0.79	0.79	0.84	0.90	0.74	0.64	0.70	0.75
Female	0.07	0.07	0.07	0.06	0.09	0.09	0.09	0.09
Years in DGT	8.6	9.1	10.1	11.1	7.8	7.7	8.7	9.7
Monthly salary (2007 IDR thousands)	6,227	5,920	5,616	5,880	6,066	5,470	5,167	5,295
Account Representatives								
Total ARs	349	341	341	369	2,101	1,862	2,057	2,494
Has college degree	0.83	0.86	0.85	0.81	0.70	0.70	0.68	0.70
Female	0.16	0.17	0.23	0.23	0.27	0.32	0.31	0.32
Years in DGT	8.3	9.2	9.9	10.4	7.9	9.0	9.6	9.8
Monthly salary (2007 IDR thousands)	4,502	4,426	4,237	4,279	4,490	4,417	4,114	4,073

Notes: This table displays tax office staffing descriptive statistics for MTO vs. non-MTO tax offices. Tax office staff characteristics data are from DGT's internal human resources database.

Table B.2: Baseline (2006 calendar year) characteristics of staff assigned to MTO vs. non-MTO in 2007-2008

	Assigned to MTO (1)	Assigned to non-MTO (2)
Panel A: Auditors		
Job performance		
Total Staff DP3 Score	563.1	564.4
Performance	78.7	78.9
Initiative	78.1	78.2
Responsibility	78.4	78.6
Cooperation	78.3	78.4
Honesty	78.4	78.5
Obedience	78.3	78.4
Loyalty	91.0	91.0
Other characteristics		
Has college degree	0.78	0.79
Female	0.06	0.12
Years in DGT	6.0	5.9
Panel B: Account Representatives		
Job performance		
Total Staff DP3 Score	561.1	562.2
Performance	78.7	78.9
Initiative	78.0	78.2
Responsibility	78.4	78.6
Cooperation	78.4	78.5
Honesty	78.3	78.5
Obedience	78.3	78.4
Loyalty	91.0	91.0
Other characteristics		
Has college degree	0.99	0.91
Female	0.35	0.28
Years in DGT	5.4	5.7

Notes: This table displays baseline (calendar year 2006) descriptive statistics for auditors and account representatives assigned to MTO vs. non-MTO tax offices upon their creation in 2007-2008. Note that this sample is conditional staff already employed at DGT as of 2006, and therefore excludes any new auditors or account representatives hired in 2007-2008. Tax office staff characteristics data are from DGT's internal human resources database.

Table B.3: Indonesia's Medium Taxpayer Offices

MTO	Included in Analysis?	Creation Year	Overseen Provinces or Districts
KPP Madya Jakarta Pusat	No	2004	DKI Jakarta (Center)
KPP Madya Batam	No	2005	Riau
KPP Madya Pekanbaru	No	2006	Riau Islands
KPP Madya Denpasar	No	2006	Bali
KPP Madya Tangerang	No	2006	Banten
KPP Madya Bekasi	No	2006	West Java
KPP Madya Jakarta Barat	Yes	2007	DKI Jakarta (West)
KPP Madya Jakarta Selatan I	Yes	2007	DKI Jakarta (South)
KPP Madya Jakarta Timur	Yes	2007	DKI Jakarta (East)
KPP Madya Jakarta Utara	Yes	2007	DKI Jakarta (North)
KPP Madya Bandung	Yes	2007	West Java
KPP Madya Semarang	Yes	2007	Central Java
KPP Madya Surabaya	Yes	2007	East Java
KPP Madya Sidoarjo	Yes	2007	East Java
KPP Madya Malang	Yes	2007	East Java
KPP Madya Balikpapan	Yes	2007	East Kalimantan
KPP Madya Makassar	Yes	2007	South, Southeast, and West Sulawesi
KPP Madya Palembang	Yes	2007	South Sumatra and Bangka Belitung Islands
KPP Madya Medan	Yes	2007	North Sumatra

Notes: This table lists all 19 KPP Madya offices in Indonesia operating as of 2019, along with their respective oversight regions. Table B.6 and Figure B.11 show robustness results including the 5 MTOs created in 2005-2006. KPP Madya Jakarta Pusat could not be included because the data needed for MTO assignment as of 2004 (for tax years 2000-2002) are not available.

Table B.4: Analysis sample restrictions

Criteria (1)	Total taxpayers (2)	Assigned to MTO in 2007 (3)	Not assigned to MTO in 2007 (4)
In eligible tax office as of pre-treatment year	101,829	4,272	97,557
Baseline gross income above IDR 100 million	60,600	4,181	56,419
2005 gross income and taxes paid within common support	20,858	1,479	19,379

Notes: This table shows taxpayer counts by treatment status following each analysis sample restriction. Eligible tax offices are the origin tax offices from which MTO taxpayers were selected according to the MTO creation regulations for the 13 KPP Madya offices created in 2007. Treatment status in columns (3)-(4) are computed based on the tax office in which the taxpayer files any corporate income taxes over years 2007-2008. A treated (untreated) taxpayer is in the common support when its gross income and total taxes paid fall within the 2.5th and 97.5th percentiles of the respective distributions among untreated (treated) taxpayers. Table B.6 shows robustness results to including very small taxpayers (that is, no baseline gross income restriction), and to allowing increasing the common support cutoffs to 1st - 99th percentiles. MTO creation regulations are available in the Directorate General of Taxes website: http://www.pajak.go.id/, and they are: KEP-30-PJ-2007 (Balikpapan); KEP-25-PJ-2007 (Bandung); Nomor KEP-21-PJ-2007 (Jakarta Barat); KEP-22-PJ-2007 (Jakarta Selatan); KEP-23-PJ-2007 (Jakarta Timur); KEP-24-PJ-2007 (Jakarta Utara); KEP-31-PJ-2007 (Makassar); KEP-29-PJ-2007 (Malang); KEP-19-PJ-2007 (Medan); KEP-20-PJ-2007 (Palembang); KEP-26-PJ-2007 (Semarang); KEP-28-PJ-2007 (Sidoarjo); KEP-27-PJ-2007 (Surabaya).

Table B.5: Robustness to alternative weighting schemes

	Main specification	Unweighted	Logit IPW 2005	Entropy 2003-2005	Logit IPW 2003-2005	Random Forest IPW 2003-2005
		Robustness to weighting method and matched years				
	(1)	(2)	(3)	(4)	(5)	(6)
Observations	163,572	163,572	161,922	95,188	94,237	94,233
Treated observations	11,819	11,819	11,717	6,928	6,861	6,865
Panel A: Tax payments (2007 IDR billion)						
Total tax payments	0.525	0.508	1.115	0.579	0.685	0.538
	(0.096)	(0.075)	(0.448)	(0.132)	(0.136)	(0.108)
VAT	0.371	0.350	0.838	0.428	0.497	0.389
	(0.078)	(0.061)	(0.355)	(0.107)	(0.092)	(0.089)
Corporate Income Tax	0.074	0.072	0.093	0.073	0.055	0.073
	(0.014)	(0.011)	(0.034)	(0.020)	(0.011)	(0.015)
Other income taxes	0.080	0.085	0.183	0.078	0.133	0.077
	(0.017)	(0.012)	(0.066)	(0.020)	(0.049)	(0.014)
Panel B: Reported income (2007 IDR billion)						
Gross income	9.131	7.628	10.922	10.411	8.251	8.507
	(2.181)	(1.663)	(3.133)	(2.905)	(1.893)	(2.289)
Taxable Income	0.238	0.234	0.483	0.255	0.178	0.274
	(0.072)	(0.055)	(0.256)	(0.097)	(0.059)	(0.075)
Total corporate income tax due	0.065	0.060	0.131	0.073	0.049	0.073
	(0.020)	(0.015)	(0.069)	(0.026)	(0.015)	(0.020)
Panel C: Employment						
Total workers	12.646	5.212	33.016	19.730	50.192	14.358
	(21.865)	(16.759)	(13.251)	(21.046)	(21.868)	(16.163)
Permanent workers	10.365	13.530	17.775	15.253	19.487	17.090
	(6.009)	(3.325)	(4.777)	(7.230)	(6.568)	(4.079)
Temporary workers	2.281	-8.318	15.241	4.477	30.704	-2.732
	(21.168)	(16.574)	(12.723)	(20.091)	(22.639)	(15.938)
Total wage bill (2007)	0.330	0.267	0.506	0.524	0.576	0.381
	(0.139)	(0.087)	(0.109)	(0.169)	(0.146)	(0.113)
Permanent workers	0.193	0.264	0.442	0.238	0.414	0.282
	(0.100)	(0.052)	(0.096)	(0.113)	(0.135)	(0.062)
Temporary workers	0.136	0.002	0.064	0.286	0.162	0.098
	(0.097)	(0.067)	(0.045)	(0.121)	(0.104)	(0.093)
Average yearly wage (2007 IDR million)	2.337	2.236	3.584	2.341	2.811	0.002
	(1.002)	(0.736)	(1.258)	(1.436)	(1.699)	(0.001)

Notes: See notes to Table 2.1. This table shows MTO treatment effect robustness results to alternative weighting schemes.

Table B.6: Robustness to alternative sample restrictions

		Robustness to sample restriction:					
	Main specification	No common support restriction	No gross income restriction	Restrict sample to 1st-99th common support		Adding 2005 and 2006 MTOs	Restrict to years 2003-2009
				Weighted	Unweighted		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Observations	163,572	455,867	192,584	293,749	293,749	202,960	130,876
Treated observations	11,819	33,043	10,172	16,287	16,287	13,967	9,493
Panel A: Tax payments (2007 IDR billion)							
Total tax payments	0.525	1.550	0.448	0.263	0.605	0.326	0.471
	(0.096)	(0.147)	(0.111)	(0.243)	(0.066)	(0.069)	(0.077)
VAT	0.371	0.712	0.331	0.163	0.374	0.233	0.346
	(0.078)	(0.096)	(0.090)	(0.185)	(0.047)	(0.057)	(0.063)
Corporate Income Tax	0.074	0.556	0.051	0.047	0.122	0.045	0.062
	(0.014)	(0.067)	(0.013)	(0.054)	(0.025)	(0.009)	(0.011)
Other income taxes	0.080	0.283	0.066	0.053	0.109	0.049	0.064
	(0.017)	(0.030)	(0.018)	(0.042)	(0.012)	(0.010)	(0.014)
Panel B: Reported income (2007 IDR billion)							
Gross income	9.131	10.378	5.901	4.845	6.929	3.983	7.763
	(2.181)	(2.695)	(2.144)	(2.808)	(1.365)	(1.214)	(1.855)
Taxable Income	0.238	1.782	0.143	0.121	0.399	0.137	0.174
	(0.072)	(0.244)	(0.081)	(0.237)	(0.105)	(0.047)	(0.062)
Total corporate income tax due	0.065	0.480	0.041	0.042	0.107	0.033	0.058
	(0.020)	(0.071)	(0.023)	(0.063)	(0.030)	(0.013)	(0.018)
Panel C: Employment							
Total workers	12.646	-36.102	27.244	8.421	1.848	10.987	25.499
	(21.865)	(20.467)	(17.353)	(29.254)	(14.860)	(15.406)	(30.210)
Permanent workers	10.365	26.893	16.377	6.290	21.942	12.042	12.962
	(6.009)	(7.201)	(4.510)	(13.758)	(7.152)	(3.088)	(5.136)
Temporary workers	2.281	-62.994	10.867	2.130	-20.094	-1.056	12.537
	(21.168)	(19.105)	(16.981)	(23.599)	(13.310)	(14.998)	(29.778)
Total wage bill (2007 IDR billion)	0.330	-0.864	0.336	0.229	0.297	0.300	0.353
	(0.139)	(0.516)	(0.119)	(0.303)	(0.138)	(0.092)	(0.125)
Permanent workers	0.193	0.297	0.228	0.196	0.421	0.205	0.208
	(0.100)	(0.198)	(0.077)	(0.258)	(0.130)	(0.059)	(0.079)
Temporary workers	0.136	-1.162	0.108	0.034	-0.124	0.095	0.145
	(0.097)	(0.476)	(0.096)	(0.122)	(0.061)	(0.066)	(0.092)
Average yearly wage (2007 IDR million)	2.337	4.423	2.728	0.312	1.495	2.022	3.633
	(1.002)	(4.258)	(0.829)	(0.002)	(0.001)	(0.782)	(1.290)

Notes: See notes to Table 2.1. This table shows MTO treatment effect robustness results to alternative sample restrictions.

Table B.7: Robustness to alternative standard error clustering levels

		Robustness to clustering	
	Main specification	Clustering at origin tax office	Clustering at regional tax office
	(1)	(2)	(3)
Observations	163,572	163,572	163,572
Treated observations	11,819	11,819	11,819
Panel A: Tax payments (2007 IDR billion)			
Total tax payments	0.525	0.525	0.525
	(0.096)	(0.094)	(0.097)
VAT	0.371	0.371	0.371
	(0.078)	(0.077)	(0.080)
Corporate Income Tax	0.074	0.074	0.074
	(0.014)	(0.014)	(0.014)
Other income taxes	0.080	0.080	0.080
	(0.017)	(0.017)	(0.019)
Panel B: Reported income (2007 IDR billion)			
Gross income	9.131	9.131	9.131
	(2.181)	(2.222)	(2.356)
Taxable Income	0.238	0.238	0.238
	(0.072)	(0.070)	(0.054)
Total corporate income tax due	0.065	0.065	0.065
	(0.020)	(0.019)	(0.015)
Panel C: Employment			
Total workers	12.646	12.646	12.646
	(21.865)	(18.887)	(20.397)
Permanent workers	10.365	10.365	10.365
	(6.009)	(5.917)	(4.293)
Temporary workers	2.281	2.281	2.281
	(21.168)	(18.817)	(19.922)
Total wage bill (2007 IDR billion)	0.330	0.330	0.330
	(0.139)	(0.124)	(0.122)
Permanent workers	0.193	0.193	0.193
	(0.100)	(0.107)	(0.095)
Temporary workers	0.136	0.136	0.136
	(0.097)	(0.083)	(0.086)
Average yearly wage (2007 IDR million)	2.337	2.337	2.337
	(1.002)	(0.987)	(0.903)

Notes: See notes to Table 2.1. This table shows MTO treatment effect robustness results to the levels at which standard errors are clustered.

Table B.8: First stage of MTO regression

Instrument:	Treatment: Taxpayer in MTO in current year (1)
(Assigned to MTO in 2007) x (Year > 2005)	0.647
	(0.008)
F-statistic	6,412.0

Notes: This table shows first stage estimates for MTO treatment effect as defined in equation (2.17). Standard errors are clustered at the taxpayer level.

Table B.9: Administrative Costs

	MTO (1)	Not MTO (2)
Total budget (2007 IDR billion)		
Staff	85.8	908.3
Goods + Capital	55.1	1187.8
Total	140.9	2096.0
Number of corporate taxpayers	18,051	1,115,850
Cost per corporate taxpayer	0.00789	0.00095

Notes: Budget data from 2016, deflated to 2007 IDR using Indonesia's GDP deflator. Taxpayer counts based on all taxpayers who filed SPT 1771 in calendar year 2016 or paid any taxes in 2016, and reflects taxpayers across all 19 regional tax offices. Taxpayers who did not file SPT 1771 in 2016 and who were never in the MTO but paid taxes in 2016 assumed to be in PTO (Primary Tax Office). We assume half of all PTO costs are for corporate taxation.

Table B.10: Detailed effects of MTO on corporate income tax returns

	Weighted means				MTO effect (IV)	
	Pre-treatment					
	Untreated	Treated	N	Treated post-treatment counterfactual	Point estimate	Standard error
Tax Filing item (2007 IDR billion)	(1)	(2)	(3)	(4)	(5)	(6)
Gross income	13.03	13.03	136,601	12.04	9.131	(2.181)
- Cost of sales	10.37	10.17	136,023	9.35	7.636	(2.029)
- Other expenses	2.16	2.42	136,549	2.04	1.126	(0.229)
Net income from business	0.69	0.49	137,135	0.59	0.427	(0.160)
+ Net income from side business	0.04	0.01	137,118	-0.04	-0.009	(0.081)
Total domestic commercial net income	0.73	0.50	137,059	0.55	0.416	(0.144)
+ Total foreign commercial net income	0.00	0.00	137,063	0.00	0.004	(0.009)
Total commercial net income	0.73	0.50	137,190	0.56	0.404	(0.149)
- Non-taxable inc. and inc. subject to final tax	0.89	0.52	137,589	0.22	0.975	(0.473)
+ Total positive fiscal adjustment	0.55	0.42	137,586	0.16	0.843	(0.448)
- Total negative fiscal adjustment	0.03	0.03	137,584	0.22	-0.124	(0.121)
Fiscal net income	0.31	0.37	137,584	0.37	0.304	(0.092)
- Compensation for fiscal loss carried forward	0.02	0.03	137,584	0.03	-0.012	(0.020)
Taxable Income	0.39	0.46	137,585	0.50	0.238	(0.072)
Total corporate income tax due	0.09	0.12	137,586	0.13	0.065	(0.020)

Notes: See notes to Table 2.1.

Table B.11: Detailed effects of MTO on tax payments

	Weighted means				MTO effect (IV)	
	Pre-treatment					
Tax payments (2007 IDR billion)	Untreated (1)	Treated (2)	N (3)	Treated post-treatment counterfactual (4)	Point estimate (5)	Standard error (6)
Total	0.371	0.371	163,572	0.409	0.525	(0.096)
VAT	0.263	0.259	163,572	0.271	0.371	(0.078)
Domestic	0.238	0.230	163,572	0.225	0.288	(0.058)
Imported	0.024	0.027	163,572	0.045	0.082	(0.047)
Other	0.001	0.002	163,572	0.001	0.001	(0.001)
Corporate Income Tax	0.046	0.056	163,572	0.067	0.074	(0.014)
Other income taxes						
Employee income tax withholding	0.029	0.025	163,572	0.037	0.021	(0.007)
Other	0.033	0.030	163,572	0.034	0.059	(0.013)

Notes: See notes to Table 2.1.

Table B.12: MTO reduced form effect heterogeneity by taxpayer baseline size

	In first MTO cohort x post (1)	In first MTO cohort x post x base size (2)	In first MTO cohort x post (3)	In first MTO cohort x post x base size (4)
	Panel A: Base size is 2005 gross income		*Panel B: Base size is 2005 taxable income*	
Total taxes paid	0.382	0.002	0.310	0.230
	(0.106)	(0.004)	(0.101)	(0.145)
N	187,363		187,363	

Notes: See notes to Table 2.1.

Table B.13: Effect of reform-induced change in net-of-tax CIT on various outcomes

			Separate by MTO status		P-value of MTO vs. Not MTO difference
		All taxpayers	MTO	Not MTO	
		(1)	(2)	(3)	(4)
Panel A: Taxable income reporting					
Intensive margin (Elasticity of Taxable Income)					
Δ Ln(Taxable Income)		0.590	0.348	0.779	0.322
		(0.198)	(0.378)	(0.216)	
	N	12,790	730	12,060	
Extensive margin					
Δ Reports positive Taxable Income		0.425	0.429	0.421	0.953
		(0.069)	(0.133)	(0.066)	
	N	19,530	1,112	18,418	
Panel B: Other income reporting outcomes					
Δ Ln(Gross Income)		-0.011	0.154	-0.093	0.518
		(0.177)	(0.327)	(0.197)	
	N	12,628	724	11,904	
Reports higher Taxable Income than in base year		0.905	0.576	1.098	0.123
		(0.150)	(0.303)	(0.151)	
	N	12,790	730	12,060	
Panel C: Tax payments					
Δ Ln(Corporate Income Tax payments)		0.631	0.154	0.915	0.112
		(0.233)	(0.436)	(0.241)	
	N	12,524	726	11,798	
Δ Ln(Total Income Tax payments)		0.537	0.366	0.634	0.505
		(0.189)	(0.353)	(0.209)	
	N	12,576	728	11,848	
Δ Ln(Total VAT payments)		-0.127	-0.614	0.113	0.321
		(0.324)	(0.679)	(0.284)	
	N	9,748	618	9,130	
Panel D: Employment					
Δ Ln(Total workers)		0.139	0.236	0.110	0.794
		(0.215)	(0.437)	(0.202)	
	N	8,766	522	8,244	
Δ Ln(Permanent workers)		0.159	0.140	0.175	0.940
		(0.199)	(0.428)	(0.166)	
	N	8,606	518	8,088	
Taxpayer FE		Yes	Yes	Yes	
Year FE		Yes	Yes	Yes	

Notes: See notes to Table 2.6. This table shows coefficients from regressions of various outcomes (displayed on separate rows) on the change in log net-of-tax CIT marginal tax rate induced by Indonesia's Marginal Tax Rate Schedule reform. All regressions are estimated in the same sample and using the same specification as the ETI specification (2.20), except for the extensive margin regression, which includes taxpayers reporting zero taxable income. Taxpayers reporting zero taxable income are assumed to face the smallest marginal tax rate in the pre-reform MTR schedule (defined based on taxable income cutoffs), and their predicted MTR in the post-reform MTR schedule (defined based on gross income cutoffs). Following the main ETI specification, all regressions with logged outcome variables further include a base-year log outcome variable control.

Table B.14: Robustness of ETI estimates

										By 2008-2009 predicted tax change	
	Main specification (1)	Unweighted regressions (2)	Re-estimated weights (3)	Restricting estimation to 2007-2010 balanced sample (4)	Using lagged data for instrument and baseline controls (5)	No taxpayer fixed effect (6)	No baseline controls (7)	Use 2008-2009 change only (8)	Use 2008-2010 change only (9)	Predicted tax cut (10)	Predicted tax raise (11)
Panel A: First Stage											
Endogenous:	0.980	0.985	0.989	0.978	0.954	0.960	0.970	0.953	0.956	0.983	0.988
Δ Ln(Net-of-tax rate)	(0.010)	(0.003)	(0.014)	(0.010)	(0.013)	(0.008)	(0.012)	(0.009)	(0.010)	(0.013)	(0.054)
F-statistic	3,629.32	56,314.51	1,881.93	3,397.39	1,692.22	4,284.21	6,596.91	4,235.58	3,382.25	2,191.57	145.77
N	12,790	26,300	6,790	10,752	10,874	14,740	13,122	8,269	7,680	9,412	3,378
Panel B: IV (ETI estimate)											
Outcome:	0.590	0.661	0.610	0.411	0.466	1.036	0.470	0.977	1.132	0.625	1.277
Δ Ln(Taxable Income)	(0.198)	(0.073)	(0.261)	(0.201)	(0.375)	(0.258)	(0.356)	(0.311)	(0.354)	(0.230)	(1.338)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Taxpayer FE	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No	Yes	Yes
Sector FE	No	No	No	No	No	Yes	No	Yes	Yes	No	No
MTO dummy	No	No	No	No	No	Yes	No	Yes	Yes	No	No
Panel C: MTR raise needed to generate MTO effect on Corporate Income Tax revenues											
Taxing MTO taxpayers	xx	xx	xx	30 pp	31 pp	xx	31 pp	xx	xx	xx	xx
Taxing all taxpayers	8 pp	9 pp	9 pp	7 pp	xx	xx	xx	13 pp	16 pp	9 pp	xx
Panel D: Revenue-maximizing corporate income tax rate											
Revenue-max CIT MTR	56%	53%	55%	65%	62%	42%	62%	44%	40%	55%	37%

Notes: See notes to Table 2.6 and Section 2.5.2. Values in Panel C are calculated for statistically significant ETIs only.

Table B.15: Robustness to ETI heterogeneity by MTO assignment vs. treatment status

	Weighted by MTO balancing weights			Unweighted		
	MTO status		P-value of MTO vs. Not MTO difference	MTO status		P-value of MTO vs. Not MTO difference
	MTO	Not MTO		MTO	Not MTO	
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: MTO status indicates taxpayer was in MTO first cohort assignment						
Δ Ln(Taxable Income)	0.348	0.779	0.322	0.348	0.615	0.493
	(0.378)	(0.216)		(0.378)	(0.094)	
N	730	12,060		730	12,060	
Panel B: MTO status indicates whether taxpayer was in MTO in each outcome year						
Δ Ln(Taxable Income)	0.614	0.549	0.886	0.887	0.580	0.481
	(0.418)	(0.192)		(0.425)	(0.096)	
N	730	12,060		730	12,060	
Taxpayer FE	Yes	Yes		Yes	Yes	
Year FE	Yes	Yes		Yes	Yes	

Notes: See notes to Table 2.6. This table shows robustness estimates for the difference in ETIs between MTO and non-MTO taxpayers. Columns (1)-(3) present estimates weighted by MTO balancing weights, while columns (4)-(5) show unweighted estimates. Panel A shows replicates the ETI estimates presented in Table 2.6, for which MTO status is defined based on first MTO cohort assignment. Panel B presents estimates for which MTO status is define based on actual MTO treatment in each regression year. As regression years are 2008-2010, Panel B includes taxpayers assigned to MTO in the 2009 MTO expansion.

Table B.16: CIT income tax increases to match MTO effects: extrapolated counterfactual

	MTO IV treatment effect (IDR billion)	MTR raise needed to generate MTO effect on total revenue	
		Taxing MTO taxpayers	Taxing all taxpayers
	(1)	(2)	(3)
Panel A: Main counterfactual: tax change among analysis sample taxpayers			
Corporate Income Tax	0.091	xx	8 pp
Total Income Taxes	0.180	xx	17 pp
Panel B: Counterfactual tax change extrapolated to taxpayers in 19 regions			
Corporate Income Tax	0.091	6 pp	5 pp
Total Income Taxes	0.180	12 pp	9 pp

Notes: See notes to Table 2.7.

Table B.17: Industry and Geographic composition of all MTO and PTO taxpayers

	MTO (1)	PTO (2)	MTO share (3)	MTO share of base year total taxes paid (4)
Panel A: Sectoral composition				
Wholesale or retail trade	5,427	134,872	4%	59%
Manufacturing	2,329	13,303	15%	73%
Construction	2,167	80,582	3%	48%
Other service activities	1,573	34,600	4%	66%
Financial services, insurance, real estate	980	15,085	6%	71%
Mining and/or extraction	875	5,942	13%	60%
Transportation	717	12,194	6%	64%
Telecommunications and/or publishing	245	11,439	2%	53%
Hotels and restaurants	197	3,460	5%	54%
Healthcare and social work	161	5,175	3%	61%
Education	83	8,673	1%	44%
Total	14,671	316,652	4%	64%
Panel B: Geographic composition				
Java (except Jakarta)	5,300	189,339	3%	66%
Jakarta	4,364	117,436	4%	61%
Sumatra	1,614	31,623	5%	72%
Riau Islands	1,679	32,297	5%	67%
Kalimantan	972	21,106	4%	70%
Sulawesi	847	29,222	3%	63%
Bali	809	17,306	4%	75%
Total	15,585	438,329	4%	64%

Notes: This table presents descriptive statistics for industry and geographic location for all taxpayers in MTO by 2011. MTO status and geographic region is determined based on filing of SPT 1771 (Corporate Income Taxes). Industry code is determined based on filing of SPT 1721 (payroll taxes). Percent of total taxes paid is based on 2005 tax year. Firms that did not file any SPT 1771 or SPT 1721 forms by 2011 are excluded from these tables.

Table B.18: 2005 average tax payments for all MTO and PTO taxpayers

Tax payments (2007 IDR billion)	MTO (1)	PTO (2)
Total	3.232	0.097
VAT	2.033	0.125
Domestic	1.360	0.091
Imported	2.346	0.725
Other	0.264	0.052
Corporate Income Tax	0.794	0.028
Other income taxes		
Employee income tax withholding	0.456	0.093
Other	0.591	0.038
Taxpayers	13,838	561,767

Notes: This table presents descriptive statistics of base year tax payments for all taxpayers in MTO as of 2011. MTO status is based on filing of SPT 1771 by 2011. Taxpayers that did not file any SPT 1771 by 2011 but made tax payments are assumed to be in PTOs of regional tax offices whose MTO creation was in 2007. Figures are based on payments for tax year 2005.

Table B.19: 2005 average CIT line items for all MTO and PTO taxpayers

Tax Filing item (2007 IDR billion)	MTO (1)	PTO (2)
Gross income	80.15	4.81
- Cost of sales	70.07	5.06
- Other expenses	11.89	0.73
Net income from business	3.74	0.14
+ Net income from side business	1.13	0.04
Total domestic commercial net income	4.59	0.16
+ Total foreign commercial net income	0.02	0.01
Total commercial net income	4.82	0.29
- Non-taxable inc. and inc. subject to final tax	5.20	0.57
+ Total positive fiscal adjustment	2.74	0.28
- Total negative fiscal adjustment	2.10	0.19
Fiscal net income	2.18	0.03
- Compensation for fiscal loss carried forward	0.40	0.02
Taxable Income	2.33	0.07
Total corporate income tax due	0.85	0.03
Total Taxpayers	12,683	136,433

Notes: This table presents descriptive statistics of base year corporate income tax line items for all taxpayers in MTO as of 2011. MTO status is based on filing of SPT 1771 by 2011. Figures are based on filings for tax year 2005. Firms that did not file SPT 1771 for tax year 2005 are excluded from this table.

Table B.20: 2005 average employment for all MTO and PTO taxpayers

	MTO (1)	PTO (2)
Total workers	331.00	51.54
Permanent workers	137.44	25.03
Temporary workers	193.56	26.51
Total wage bill (2007 IDR billion)	5.37	0.88
Permanent workers	3.62	0.71
Temporary workers	1.74	0.17
Average yearly wage (2007 IDR million)	0.03	0.01
Total taxpayers	12,257	164,554

Notes: This table presents descriptive statistics of base year employment for all taxpayers in MTO as of 2011. MTO status is based on filing of SPT 1771 by 2011. Employment data is based on form SPT 1721. Taxpayers that did not file any SPT 1771 by 2011 but filed SPT 1721 are assumed to be in PTOs. Figures are based on SPT 1721 filings for tax year 2005. Firms that did not file any SPT 1771 by 2011 or did not file SPT 1721 for tax year 2005 are excluded from this table.

B.2 Chapter 2 Data Appendix

Corporate Income Tax: Form SPT 1771

Taxpayers file SPT 1771 forms at the headquarter level, reporting aggregate income across all branches. The corporate income tax filing microdata includes all non-identified line items from Form SPT 1771, and are tracked over time under consistent variable names.[1]

Each observation in the dataset is a taxpayer filing for a particular tax year at a particular date. The variables in the SPT 1771 microdata contain each line item from the main form (SPT 1771) and its Annex I (SPT 1771-I). In particular, it includes each component of the major corporate income tax line items, such as net income (gross income - cost of sales - other expenses), fiscal net income (net income +/- fiscal adjustments), taxable income (fiscal net income - compensation for fiscal loss carried forward), and the amount of tax overpaid or underpaid by the taxpayer as of the year end.

When analyzing effects on tax payments, we assume that all corporate income tax overpayments are refunded to the taxpayer, and thus subtract them from corporate income taxes paid as reported in the payments data. In practice, less than 1% of taxpayers in our analysis sample overpaid corporate income taxes.

Finally, SPT 1771 microdata includes the tax office code under which the corporate income tax form was filed, and an indicator for whether the filing is a correction filing or an original filing. We use the tax office code under which SPT 1771 was filed to define whether the taxpayer has been assigned to an MTO or not, and the correction indicator to construct variables tracking correction filing timing and content.

Employee Income Tax Withholding: Form SPT 1721

Firms are required to report the amount of personal income tax withheld from employees' paychecks on a monthly basis through Form SPT Masa 1721. The SPT 1721 microdata consists of two datasets, one covering tax years 2002-2008, and the other covering tax years 2009-2013. The split reflects a major change in form SPT 1721 that produced finer reporting by different employee categories. Because only very few observations are available for tax year 2008 (the last year under the old form), we exclude SPT 1721 records for tax year 2008 from all analyses.

[1] The forms SPT 1771 and SPT 1771-I have also remained largely unchanged over the analysis period, and are available at http://www.pajak.go.id/sites/default/files/formulir_pajak/Formulir%20SPT%201771-%24.pdf.

Each observation in the 2002-2008 dataset is a branch-level year-end reporting of cumulative income tax withholdings, reported at the branch level. The 2009-2013 data is further disaggregated by month, with cumulative totals for the year reported in the month of December. In terms of variables, the information consistently reported in both datasets includes: number of employees, wage bill, and individual income tax withheld. These data are also separated by two groups of employees: permanent and/or pensioner employees, and temporary employees.[2]

We combine the two datasets to construct a taxpayer-level annual panel dataset . Within each dataset, we aggregate the branch-level data to the taxpayer level. As the 2002-2008 data are reported in year-end totals, we use the year-end total reported in the December monthly filing for the 2009-2013 data.

Tax Payments

Detailed tax payments data are from the Treasurer's Modul Penerimaan Negara (MPN; State Revenues Module) database, and cover all types of income tax and VAT paid by corporations.

Each observation in the tax payments data is a branch-level payment made on a particular date for a particular tax type and month. The tax type variable differentiates different types of income and VAT. We break taxes down by the following major categories: corporate income taxes, VAT, and other income taxes.[3]

Tax Audits, Assessments, and Disputes

DGT may conduct a tax audit of any or all of a taxpayer's filings and payments. At the end of every audit, DGT issues a tax assessment letter and/or a tax collection letter to the taxpayer. The tax assessment letter informs the taxpayer of outstanding tax obligations

[2]Number of employees at year-end in the 2002-2008 data and in the 2009-2013 data reflect the total of unique employees employed during the respective tax year. While the 2009-2013 data distinguishes between permanent vs. pensioner employees, the 2002-2008 data does not. As a result, we sum the 2009-2013 employee numbers to construct a consistent series of permanent and/or pensioner employment.

[3]These categories are sub-divided in the data by tax articles. For income taxes: PPh Pasal 25/29 (corporate income tax monthly installments and year-end payments), PPh Pasal 21 (domestic employee withholding), PPh Pasal 26 (foreign employee withholding), PPh Pasal 22 (income tax on import transactions), PPh Pasal 23 (income tax on capital transactions), and PPh Final or Pasal 4 (2) (income tax withholding on gross payments of certain items). For VAT: PPn Domestic, PPn Import, and PPn Other. Tax payments that count towards a company corporate income tax liability include PPh Pasal 25/29, PPh Pasal 23 and PPh Pasal 22. Finally, the payment data also includes codes for administrative penalties levied on income and/or VAT taxes. These penalties account for roughly 0.08% of all tax payments in the data.

(none, underpaid, overpaid), while the tax collection letter is typically used to levy administrative tax sanctions resulting from the audit.[4] Our tax audit microdata consists of two datasets covering this audit process.

The first dataset covers all audits since 2009, and documents what was audited and why (that is, the audit triggers). Each observation in this dataset is an audit occurrence, and it includes the following main variables: the taxpayer anonymized ID, the audit date, the object audited (e.g., CIT, VAT, location changes), the tax period audited (e.g., a particular month or range of months), and the audit trigger (e.g., risk analysis, office routine, etc.).

The second dataset is specific to VAT audits, and covers the audit result process for all audits since 2002. Each observation in this dataset is either the issuance of a tax collection letter or of a VAT underpayment tax assessment letter. The available variables are: the taxpayer anonymized ID; the issuance type (collection or underpayment assessment) and date; and the total underpaid amount (or administrative penalty) found in the audit.

In addition, because either a tax collection letter or an underpayment letter is a legal instrument with which DGT may confiscate the owed amount/levied penalties, this dataset further includes as variables the issuance dates of all subsequent letters exchanged between DGT and the taxpayer during the tax dispute process. Specifically, these are: a warning letter (issued if the amount/penalty is not paid by its deadline), a distress warrant (issued if the underpaid tax is not settled within 21 days of the warning letter), and a confiscation letter (issued if the underpaid amount is not settled within 48 hours of the distress warrant).

Finally, because, by law, taxpayers are only required to pay the amount of taxes they agree to have underpaid (so long as the amount to which the taxpayer disagrees is formally disputed through an objection letter), the data further includes: the amount of taxes the taxpayer disagrees to have underpaid; the date in which the taxpayer filed an objection letter concerning the disagreed amount; and lastly, in case the objection is denied, the date in which the taxpayer filed an appeal to the Tax Court requesting further review of the case.

Tax Office staffing

We compute staff descriptive statistics for MTO and PTO offices using anonymized staff-level panel data provided by DGT. These data include basic staff demographic characteristics, as well as information on staff position (i.e., auditor or AR) at different points in time.

[4]For a more detailed description of Indonesia's tax audit and assessment process, see, for example, `https://www.pwc.com/id/en/pocket-tax-book/english/pocket-tax-book-2019.pdf`.

Information on staff position and years of experience are then matched with position-specific and experience-specific wage schedules to compute average salary statistics.

Sample Restrictions for Matching

When constructing our analysis sample and computing balancing weights, we attempt to mimic the MTO assignment process conducted by DGT as closely as possible. Appendix Table B.4 outlines each sample restriction step. First, we focus on taxpayers who were registered as of 2006 in a tax office from which MTO taxpayers were sourced (that is, in an "eligible" tax office for MTO selection). The list of tax offices from which MTO taxpayers were sourced can be obtained separately for each MTO from its creation regulation.[5] This bring us to 101,829 corporate taxpayers registered in an eligible tax office as of 2006, of which 4,272 were assigned to an MTO in 2007.

Second, a large number of the taxpayers registered in eligible tax offices are small microbusinesses that would not have been shortlisted for MTO assignment. We therefore exclude taxpayers with gross income below IDR 100 million (roughly USD 10,000 at the 2007 exchange rate) during baseline years 2003-2005, bringing the shortlisted sample to 60,600 taxpayers, 4,181 of which were assigned to an MTO in 2007.

Finally, as recommended in the propensity score and matching literature (Dehejia and Wahba, 1999; Heckman et al., 1997; Stuart, 2010), we focus on taxpayers whose baseline MTO assignment inputs share common support. We define common support based on the 2.5th and 97.5th percentiles of each MTO assignment input. For example, in our main specification the matched variables are the 2005 gross income and the 2005 total taxes paid. The treated (untreated) taxpayers in the common support are those whose 2005 gross income and total taxes paid fall within the 2.5th and 97.5th percentiles of the 2005 gross income and total taxes paid distributions of the untreated (treated) taxpayers.

With this final restriction in place, we arrive at our analysis sample of 20,858 taxpayers, 1,479 of which are assigned to an MTO. Appendix Table B.6 presents robustness results to the gross income and common support restrictions.

[5] In particular, each regulation lists in an attachment all the NPWPs (Tax IDs) assigned to its respective newly created MTO. NPWPs are composed of 15 digits. The first 9 digits uniquely identify the firm, the next three identify the tax office in which the NPWP is registered, and the last 3 identify the branch (e.g., 000 indicates headquarters). While we cannot directly match these IDs to our data as our data are anonymized, we can extract from each NPWP in the regulation the origin tax office from which it came as the NPWP's middle 10th-12th digits.

B.3 Chapter 2 Model Appendix: Adding an evasion margin on the cost dimension

Suppose that in addition to the model outlined in Section 2.3 above, we add, as in Best et al. (2015), for lines that are not hidden (i.e., for which firms pay taxes), firms have another margin of evasion: they may misreport costs $\hat{c} \neq c(y)$ at a cost $\alpha g(\hat{c} - c(y))$, with $g(0) = 0$ and g convex, such that $0 \leq \hat{c} \leq y$. We assume that some fraction of reported costs μ are not deductible from taxes.[6]

For business lines on which firms that evade taxes entirely, the decision remains unchanged from the model above.

For business lines on which firms do not evade entirely, instead of the maximization problem in equation (2.1), these firms solve the following:

$$\max_{y,\hat{c}}(1-\tau)y - c(y) + \tau\mu\hat{c} - \alpha g(\hat{c} - c(y)) \tag{B.1}$$

which for an interior solution (i.e., $0 < \hat{c} < y$) yields the optimum conditions

$$c'(y^t) = 1 - \tau\frac{1-\mu}{1-\tau\mu} \tag{B.2}$$

and

$$\alpha g'(\hat{c} - c(y^t)) = \tau\mu \tag{B.3}$$

where y^t is the optimal level of production y for firms that pay tax.

In this model, an increase in the cost of evasion α reduces evasion for these business lines, as they will increase their reported costs (see equation (B.3)). It does not, however, affect real output choices for business lines paying taxes as long as we are at an interior solution where $\hat{c} < y$), which remained governed by equation (B.2).

To understand the net effect of an increase in enforcement α, we need to reconsider the indifference condition for which business lines will evade or not (i.e., equation (2.5)). This is now given by:

$$y^e_{l^*}(\alpha) - c(y^e_{l^*}(\alpha)) - \alpha b(y^e_{l^*}(\alpha))h(l^*) = (1-\tau)y^p - c(y^p) + \tau\mu\hat{c} - \alpha g(\hat{c} - c(y^p)) \tag{B.4}$$

An increase in α now has an ambiguous effect on the extensive margin decision, because

[6] Best et al. (2015) endogenize μ. We take it as a fixed parameter for our purposes.

there are now two effects. As in equation (2.5), increasing α decreases the profits from an evaded business line (the left-hand side of equation (B.4)), because it increases evasion costs (by $b(y_l^e(\alpha))h(l^*)$, from the envelope theorem).[7] On the other hand, it also reduces the after-tax profits from a non-evaded business line, because the intensive margin cost of evasion has also increased (by $g(\hat{c} - c(y^p))$, again by the envelope theorem). Which of these dominates is ambiguous.

This implies that there is a possibility that, when there is both intensive margin and extensive margin evasion, increasing enforcement costs could actually backfire, i.e., could lead to a decrease in total tax revenue. For this to happen, however, two conditions would need to hold. First, one would need that $g(\hat{c}-c(y^p)) > b(y_l^e(\alpha))h(l^*)$, so increasing α leads to more extensive margin evasion, rather than less. Second, the lost tax revenue from marginal lines that are induced to evade entirely (given by $\tau[y^p - \tau\mu\hat{c}]\frac{\partial l^*}{\partial \alpha}$), would have to offset the increase in tax revenue for all infra-marginal lines (given by $-\int_{l^*}^{L} \tau\mu\frac{\partial \hat{c}}{\partial \alpha}$).

The welfare analysis in Section 2.3.3, however, is unaffected by considering the possibility of intensive margin evasion as well. The only difference is that (unobserved) private compliance costs in equation (2.8), γ, now need to include private compliance costs for both fully evading and partially evading business lines. This can be written as

$$\gamma = \int_0^{l^*} \alpha b(y_l^e(\alpha))h(l) + \int_{l^*}^{L} \alpha g(\hat{c} - c(y^p))$$

Otherwise, the key expressions for calculating the effect of administrative and tax changes on welfare in equations (2.8), (2.9), (2.12), and (2.14) that guide our empirical analysis remain unchanged.

B.4 Chapter 2 Tax formulas Appendix

We slightly modify the notation in Section 2.3 to account for the fact that we have a progressive tax schedule, and so we consider changes to the top marginal rate. Using the notation from Saez et al. (see 2012), we note that a change in marginal tax rates has two components, a mechanical effect (dM) and a behavioral effect (dB). Under the assumption of a constant ETI ε and that all taxpayers above $\bar{z}$ face a single marginal tax rate, the

[7] Note that this would imply a reduction in reported revenues, since some lines would now evade entirely; the increase in reported costs would be ambiguous, since lines evading entirely would cease reporting costs, but those only partially evading would report higher costs.

mechanical effect of a tax change $d\tau$ is given by:

$$dM \equiv N \cdot (z^m - \bar{z})\, d\tau > 0 \tag{B.5}$$

while the behavioral effect is:

$$dB \equiv -N \cdot \varepsilon \cdot z^m \left(\frac{\tau}{1-\tau} \right) d\tau < 0 \tag{B.6}$$

where z^m is the average taxable income among those taxpayers, and τ is the top marginal tax rate. In other words, dM is the total revenue that would be raised for a percentage point change $d\tau$ to the top marginal tax rate τ absent any behavioral responses, whereas dB captures the behavioral reduction in total taxable income reported for that same change. The change in revenue is the difference between the mechanical effect and the behavioral effect, i.e., $dR = dM + dB$.

Combining these terms yields the expression for the marginal excess burden of taxation:

$$-\frac{dB}{dR} = \frac{\varepsilon \tau \rho}{1 - \tau - \varepsilon \tau \rho} \tag{B.7}$$

where $\rho = \left(\frac{z^m}{z^m - \bar{z}} \right)$ is the Pareto parameter.

We can also use the estimated ETI to compute optimal marginal tax rates as a function of v, the marginal cost of public funds. Modifying equation (2.9) to take into account the fact that we are considering a top marginal tax rate change, the top optimal tax rate is given by $\tau^* = \frac{1}{1+\rho\varepsilon\frac{v}{v-1}}$. This is given by rewriting equation (2.9) as $W_\tau = (v-1)dM + vdB$, and using equations (B.5) and (B.6).

To compare the administration reform with the tax change, recall that in Section 2.3.1, we derived in equation (2.12) the relationship between marginal tax rate changes and changes in administration. This is given by:

$$\frac{d\tau}{d\alpha}|_R = -\frac{\tau \frac{dz}{d\alpha} - \frac{da}{d\alpha}}{z\left(1 - \frac{\tau}{1-\tau}\varepsilon_{1-\tau}\right)} \tag{B.8}$$

where $\tau \frac{dz}{d\alpha} - \frac{da}{d\alpha}$ is the empirically estimated change in tax revenue (net of administration costs) from the introduction of the MTO estimated in Section 2.4, $\varepsilon_{1-\tau}$ is the estimated elasticity of taxable income with respect to the net of tax rate estimated in Section 2.5.2, and τ is the marginal tax rate from which we are starting. To take this to the data, we modify this equation slightly to account for the fact that we have a progressive tax schedule, and therefore are considering changes to the top rates. Modifying equation (B.8) to consider the effect of an increase in the top marginal rate yields:

$$\frac{d\tau}{d\alpha}|_R = -\frac{\overbrace{\tau\frac{dz}{d\alpha} - \frac{da}{d\alpha}}^{\text{Total MTO effect}}}{\underbrace{N\left(z^m - \bar{z}\right)}_{\text{Total income subject to raise}}\left[1 - \underbrace{\left(\frac{\tau}{1-\tau}\right)\varepsilon_{1-\tau}\left(\frac{z^m}{z^m - \bar{z}}\right)}_{\text{Behavioral effect}}\right]} \tag{B.9}$$

where N is the number of taxpayers above the 2006 top rate taxable income threshold $\bar{z}$ of IDR 100 million (i.e., those already paying the top marginal rate), $\rho = \left(\frac{z^m}{z^m - \bar{z}}\right)$ can be computed from the tax data, and τ is the pre-period, top marginal tax rate (30 percent).

Appendix C

Appendices for Chapter 3

C.1 Chapter 3 Appendix Figures and Tables

Figure C.1: Percent of Boston students suspended out-of-school

Panel A: Charters vs. Other

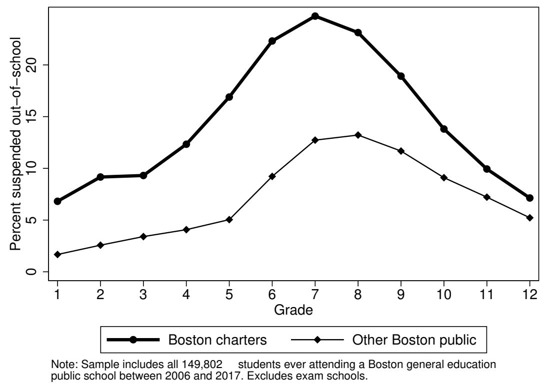

Panel B: Black or Hispanic vs. Other

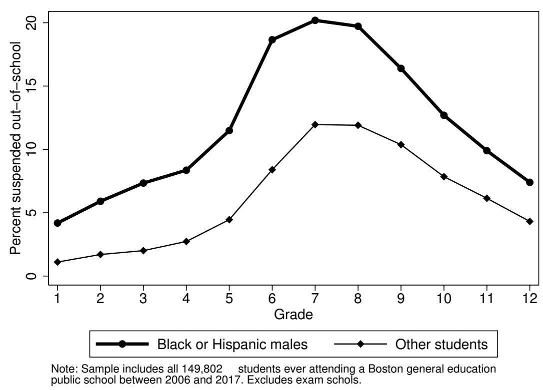

Note: The sample includes all students ever attending a Boston general education public school between 2006 and 2017, excluding exam schools.

Figure C.2: Percent of Boston students suspended in-school

Panel A: Charters vs. Other

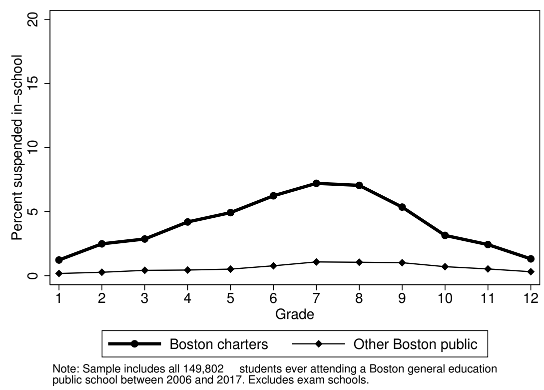

Panel B: Black or Hispanic vs. Other

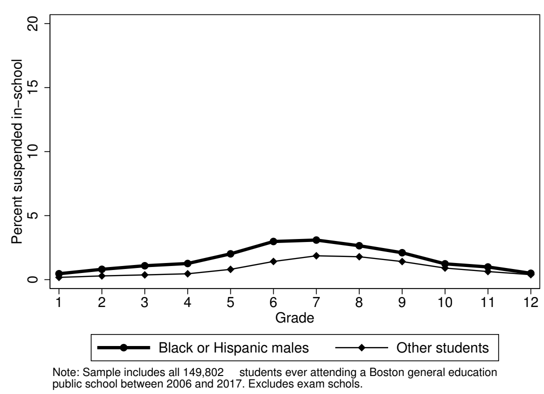

Note: The sample includes all students ever attending a Boston general education public school between 2006 and 2017, excluding exam schools.

Figure C.3: Year-by-year charter attendance 2SLS treatment effects on test scores

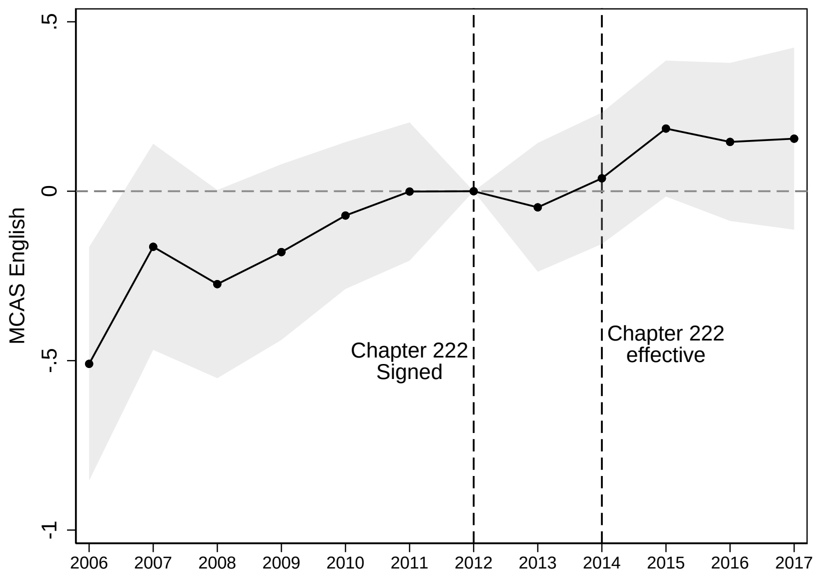

Note: See notes to Figure 14.

Figure C.4: Applicant suspension propensities

Panel A: Suspended vs. Not suspended

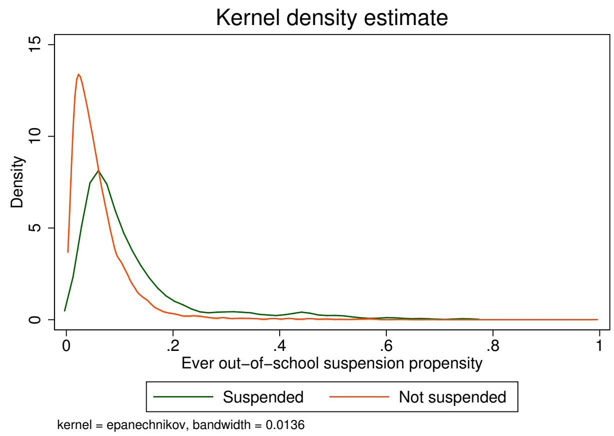

Panel B: Charter lottery winner vs. loser

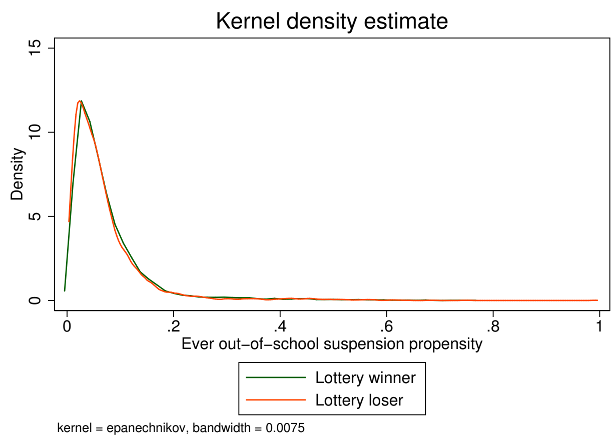

Note: This figure displays the distribution of student suspension propensity scores by suspension status and charter offer status. In Figure A4a, suspended students are those ever suspended, whether in-school or out-of-school, in the first year after the charter lottery. Applicant suspension propensities are estimated in two steps. First, a logit regression of a dummy for whether a student is ever suspended in academic year y is regressed on a rich set of predictors measured as of year y-1 on the sample of Boston students in grades 3-8 who never apply to charter schools. Table A5 displays the list of predictors along with their odds ratio coefficients. Second, the covariance structure estimated in this first step is used to predict suspension propensities in the sample of charter applicants, using each applicant's baseline grade measures as predictors for the applicant-specific suspension propensity.

Table C.1: IV Charter attendance effect on first year post-lottery outcomes

	Lottery losers mean (1)	Charter attendance effect (2SLS) (2)
First stage		
Instrument: any charter offer		0.501***
		(0.012)
F-statistic		1,821
Discipline outcomes		
Suspended out-of-school	0.091	0.173***
		(0.018)
Suspended in-school	0.017	0.081***
		(0.010)
Days suspended out-of-school	0.296	0.824***
		(0.113)
Days suspended in-school	0.037	0.179***
		(0.030)
Expelled	0.001	0.002
		(0.002)
MCAS test scores		
Math	-0.365	0.400***
		(0.033)
English	-0.457	0.227***
		(0.034)
N	4,054	8,206

Notes: This table displays 2SLS estimates of charter attendance for Boston charter middle school applicants. The first stage estimate is the regression coefficient of the any-charter attendance dummy on an any- charter lottery offer dummy, controlling for fully-saturated charter application risk sets, and a set of baseline covariate controls. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.2: Charter attendance effect for post-lottery outcomes

		Charter attendance effect			
	Lottery losers mean	1 year after lottery	2 years after lottery	3 years after lottery	4 years after lottery
	(1)	(2)	(3)	(4)	(5)
First stage					
Instrument: any charter offer		0.501***	0.353***	0.293***	0.223***
		(0.012)	(0.013)	(0.013)	(0.014)
F-statistic		1,821	762	494	246
Discipline outcomes					
Suspended out-of-school	0.091	0.173***	0.153***	0.187***	0.050
		(0.018)	(0.027)	(0.034)	(0.043)
Days suspended out-of-school	0.296	0.824***	0.957***	0.863***	-0.633
		(0.113)	(0.204)	(0.240)	(0.715)
Suspended in-school	0.017	0.081***	0.053***	0.068***	0.034
		(0.010)	(0.016)	(0.019)	(0.023)
Days suspended in-school	0.037	0.179***	0.240***	0.198***	0.068
		(0.030)	(0.060)	(0.054)	(0.061)
Expelled	0.001	0.002	-0.002	0.000	0.000
		(0.002)	(0.003)	(0.003)	(0.006)
MCAS test scores					
Math	-0.365	0.400***	0.754***	0.761***	0.814***
		(0.033)	(0.049)	(0.061)	(0.100)
English	-0.457	0.227***	0.486***	0.511***	0.716***
		(0.034)	(0.050)	(0.063)	(0.104)
N	4,054	8,206	7,886	7,548	3,657

Notes: This table displays 2SLS estimates of charter attendance for Boston charter middle school applicants, separately estimated for each year since the charter lottery application. The first stage estimate is the regression coefficient of the any-charter attendance dummy on an any- charter lottery offer dummy, controlling for fully-saturated charter application risk sets, and a set of baseline covariate controls. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.3: Heterogeneity in charter attendance effect on first year post-lottery outcomes by baseline suspension

	Attended any charter	Attended any charter x Suspended at baseline
	(1)	(2)
Test scores		
MCAS Math	0.402***	-0.038
	(0.033)	(0.107)
MCAS English	0.226***	0.032
	(0.034)	(0.122)
First stage		
Excluded instruments		
Any charter offer	0.505***	-0.003*
	(0.012)	(0.001)
Any charter offer x Suspended at baseline	0.044	0.598***
	(0.045)	(0.043)
F-statistic	1,792	626
p-value	0.000	0.000
Degrees of freedom		
df1	1	1
df2	7,779	7,779
N	8,206	

Notes: This table displays 2SLS estimates of heterogeneity in charter attendance effects in the first year after lottery by applicant's baseline grade suspension status. All regressions control for fully-saturated charter application risk sets and non- disciplinary baseline covariate controls. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.4: Year-by-year 2SLS estimates of charter attendance treatment effects

	Treatment: ever attended any charter											
	Before Chapter 222 Signing							After Chapter 222 Signing				
										Chapter 222 effective		
Post-lottery calendar year:	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Discipline outcomes												
Suspended	0.344***	0.181***	0.308***	0.127**	0.207***	0.176***	0.243***	0.172***	0.158***	0.101**	0.160**	0.064
	(0.080)	(0.067)	(0.061)	(0.052)	(0.041)	(0.037)	(0.034)	(0.038)	(0.040)	(0.045)	(0.064)	(0.073)
Suspended out-of-school	0.314***	0.170**	0.244***	0.050	0.167***	0.142***	0.235***	0.171***	0.167***	0.092**	0.096	0.065
	(0.077)	(0.066)	(0.057)	(0.049)	(0.039)	(0.035)	(0.033)	(0.037)	(0.039)	(0.043)	(0.061)	(0.070)
N	549	875	1,205	1,545	1,768	1,916	3,145	3,979	4,921	4,738	3,576	2,147
Test score outcomes												
MCAS Math	0.705***	0.851***	0.607***	0.571***	0.529***	0.774***	0.566***	0.691***	0.676***	0.566***	0.665***	0.588***
	(0.157)	(0.146)	(0.136)	(0.120)	(0.095)	(0.092)	(0.083)	(0.089)	(0.099)	(0.111)	(0.159)	(0.176)
MCAS English	0.038	0.393***	0.157	0.129	0.387***	0.429***	0.450***	0.421***	0.525***	0.593***	0.430***	0.621***
	(0.149)	(0.140)	(0.136)	(0.124)	(0.093)	(0.087)	(0.081)	(0.087)	(0.099)	(0.110)	(0.159)	(0.189)
N	545	855	1,170	1,460	1,714	1,863	3,055	3,866	4,736	4,400	3,374	2,058
First stage												
Instrument: any lottery offer	0.475***	0.455***	0.426***	0.437***	0.495***	0.488***	0.419***	0.369***	0.320***	0.307***	0.257***	0.289***
	(0.039)	(0.034)	(0.030)	(0.028)	(0.024)	(0.023)	(0.018)	(0.017)	(0.017)	(0.019)	(0.024)	(0.031)
F-statistic	155	183	205	259	429	454	540	453	342	241	120	92
N	545	855	1,170	1,460	1,714	1,863	3,055	3,866	4,736	4,400	3,374	2,058

Notes: This table displays 2SLS regression coefficients displayed in Figures 2 and 3, which are estimated from year-by-year regressions of the outcomes listed on the left on an ever-attended-charter dummy. The instrument in each regression is an any-charter lottery offer. All regressions control for fully-saturated charter application risk sets and baseline grade covariates. Since charter applicants enter the sample in different years and at different grades, all regressions include outcome year, grade, and years-since-charter-lottery fixed effects. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.5: OLS effect of suspensions on charter applicant test scores

	Suspension OLS in BPS		Suspension OLS in Charter	
	Not suspended mean	Suspended	Not suspended mean	Suspended
	(1)	(2)	(3)	(4)
MCAS Test Scores				
Math	-0.342	-0.163***	0.040	-0.110***
		(0.035)		(0.026)
English	-0.419	-0.150***	-0.173	-0.096***
		(0.037)		(0.027)
N		4,582		3,619

Notes: This table reports OLS estimates of the effect of being suspended on a student's test score outcomes conditional on the school type that the student attends (Charter or Boston Public Schools). Regressions control for the student's propensity to be suspended and for all baseline covariates listed in Appendix Table C.6. See Table C.6 and Figure C.4 for details on the estimation of the student suspension propensity. The sample is applicants to charter schools offering seats for entry grades 5 or 6 in academic years 2004-2005 through 2014-2015. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.6: Predictors used in estimating charter applicant suspension propensity scores

Variables commonly used as covariate controls in charter lottery studies	Odds ratio coefficient (1)	*Additional predictors from disciplinary and enrollment recods*	Odds ratio coefficient (2)
Demographics		*Baseline grade suspension*	
Female	0.571***	Ever suspended out-of-school	4.049***
	(0.012)		(0.138)
Asian	0.648***	Ever suspended in-school	2.326***
	(0.045)		(0.265)
Black	2.036***	Days suspended out-of-school	1.098***
	(0.082)		(0.011)
Hispanic	1.434***	Days suspended in-school	0.978***
	(0.061)		(0.052)
Other non-white	1.858***		
	(0.089)		
Baseline grade measures		*Baseline grade enrollment*	
Free or reduced price lunch	1.477***	Days attended school	0.994***
	(0.046)		(0.000)
English Language Learner	0.769***	Transferred to another school	1.164***
	(0.020)		(0.083)
English MCAS	0.803***	Repeated baseline grade	1.066***
	(0.012)		(0.058)
Math MCAS	0.781***	Immigrant	0.652***
	(0.013)		(0.036)
Special education	1.059***	Age	1.124***
	(0.024)		(0.016)

Note: This table reports odds ratio coefficients from a school or out-of-school) status on the listed variables plus grade fixed effects. The logistic regression is estimated on a sample of Boston students who never applied to charter schools. The sample contains students in grades 3-8 between between years 2004 and 2017. Suspension propensity scores are then predicted for charter applicants using applicants' baseline grade measures as predictors.

Table C.7: Suspension effect in charters vs. Charter effect on suspended: robustness to excluded instruments

	Instruments: individual charter offers					Instruments: individual charter offers plus interactions with dummy for baseline out-of-school suspension					Instruments: individual charter offers plus interactions with applicant suspension propensity score				
	Treatments					Treatments					Treatments				
	Suspended	Attended charter	Attended charter x Suspended	Suspension effect in charters (1)+(3)	Charter effect on suspended (2)+(3)	Suspended	Attended charter	Attended charter x Suspended	Suspension effect in charters (6)+(8)	Charter effect on suspended (7)+(8)	Suspended	Attended charter	Attended charter x Suspended	Suspension effect in charters (11)+(13)	Charter effect on suspended (12)+(13)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
MCAS test scores															
Math	-3.824***	0.073	3.813***	-0.011	3.886***	-1.336***	0.239***	1.318***	-0.018	1.557***	-1.096***	0.304***	0.876***	-0.219	1.181***
	(1.239)	(0.092)	(1.175)	(0.264)	(1.115)	(0.481)	(0.054)	(0.425)	(0.172)	(0.403)	(0.412)	(0.048)	(0.340)	(0.157)	(0.325)
English	-2.890***	-0.106	3.097***	0.207	2.991***	-1.506***	-0.001	1.678***	0.172	1.677***	-1.210***	0.078	1.160***	-0.050	1.238***
	(1.037)	(0.085)	(0.995)	(0.224)	(0.937)	(0.503)	(0.057)	(0.458)	(0.172)	(0.432)	(0.423)	(0.050)	(0.354)	(0.159)	(0.337)
First stage F-statistics															
F-statistic	2.497	7.846	2.754	--	--	2.548	17.929	3.076	--	--	2.867	32.648	3.557	--	--
Degrees of freedom	8	8	8	--	--	18	18	18	--	--	18	18	18	--	--
N			8,149					8,149					8,149		

Notes: This table displays robustness to the set of excluded instruments for the estimates in Table 3.3. Instruments in Columns (1)-(5) are individual charter offers only; whereas Columns (6)-(10) and (11)-(15) present estimates interacting individual charter offers with a dummy indicating if the applicant was suspended out-of-school in the baseline grade, or the applicant suspension propensity score, respectively.

Table C.8: Covariate balance for charter middle school lottery applicants

	Any-charter lottery losers mean	Offered any charter seat
	(1)	(2)
Discipline baseline		
Suspended out-of-school	0.037	0.001
		(0.005)
Suspended in-school	0.003	0.000
		(0.001)
Expelled	0.000	0.000
		(0.000)
N	9,646	
Academic achievement baseline		
MCAS Math	-0.408	-0.019
		(0.024)
MCAS English	-0.482	-0.002
		(0.025)
N	8,906	
Time-varying demographics		
Low income	0.744	0.000
		(0.011)
Special education	0.201	-0.014
		(0.010)
Limited English Proficient	0.257	-0.003
		(0.011)
N	9,646	
Gender and race		
Female	0.488	0.000
		(0.013)
Race		
Black	0.437	-0.013
		(0.012)
Hispanic	0.248	0.020*
		(0.011)
White	0.170	-0.005
		(0.008)
Asian	0.033	0.003
		(0.005)
N	9,646	
Balance joint F-statistic p-value		0.456

Notes: This table displays covariate balance on baseline characteristics of charter lottery winners and losers. Column (2) displays OLS regression coefficients from regressions of each baseline characteristic on an any-charter offer dummy. All regressions control for fully-saturated charter application risk sets. Means for losers of all charter lotteries are displayed in Column (1) for reference. The joint F-statistic corresponds to the t-statistic of the any-charter offer dummy coefficient from a stacked regression of all baseline characteristics on the any-charter offer dummy. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors at the attended school level are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.9: Year-by-year charter attendance covariate balance

	Instrument: any lottery offer											
Post-lottery calendar year:	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Baseline discipline												
Suspended out-of-school	-0.007	-0.038**	-0.009	-0.009	0.017	0.007	0.001	-0.003	-0.001	0.002	-0.004	0.006
	(0.018)	(0.017)	(0.016)	(0.014)	(0.011)	(0.009)	(0.007)	(0.007)	(0.007)	(0.008)	(0.009)	(0.013)
Suspended in-school	--	--	--	--	0.002	0.001	-0.001	-0.003	-0.002	0.001	-0.001	0.004
	--	--	--	--	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)	(0.003)	(0.004)	(0.006)
Expelled	--	--	--	--	-0.003	-0.002	-0.001	--	--	--	--	--
	--	--	--	--	(0.002)	(0.002)	(0.001)	--	--	--	--	--
N	549	875	1,205	1,545	1,768	1,916	3,145	3,979	4,921	4,738	3,576	2,147
Baseline test scores												
MCAS Math	0.081	0.005	-0.038	-0.024	-0.052	-0.035	-0.024	0.010	0.049	0.058	0.062	0.068
	(0.135)	(0.095)	(0.069)	(0.059)	(0.052)	(0.049)	(0.038)	(0.035)	(0.034)	(0.037)	(0.044)	(0.060)
MCAS English	0.051	0.028	-0.018	-0.044	-0.008	0.006	0.012	0.025	0.040	0.066*	0.074	0.060
	(0.144)	(0.100)	(0.075)	(0.062)	(0.053)	(0.050)	(0.038)	(0.036)	(0.034)	(0.037)	(0.045)	(0.060)
Baseline demographics												
Low income	0.036	0.031	0.029	0.019	0.018	0.032	-0.033**	-0.030*	-0.015	-0.007	0.009	0.012
	(0.044)	(0.037)	(0.033)	(0.030)	(0.026)	(0.024)	(0.016)	(0.016)	(0.014)	(0.016)	(0.019)	(0.025)
Special education	0.002	-0.013	-0.009	0.006	-0.035	-0.026	-0.016	-0.006	-0.013	-0.044***	-0.025	-0.045*
	(0.038)	(0.032)	(0.028)	(0.026)	(0.022)	(0.020)	(0.016)	(0.015)	(0.014)	(0.016)	(0.019)	(0.026)
Limited English Proficient	-0.015	-0.028*	-0.005	0.020	0.009	0.036*	0.014	-0.003	-0.005	-0.031*	-0.012	-0.008
	(0.018)	(0.016)	(0.018)	(0.020)	(0.018)	(0.019)	(0.017)	(0.017)	(0.017)	(0.019)	(0.022)	(0.029)
N	549	875	1,205	1,545	1,768	1,916	3,145	3,979	4,921	4,738	3,576	2,147
Gender and race												
Female	0.083*	0.053	0.006	-0.014	-0.002	-0.002	0.001	0.010	0.002	0.014	0.025	0.025
	(0.047)	(0.040)	(0.035)	(0.032)	(0.029)	(0.027)	(0.021)	(0.019)	(0.019)	(0.020)	(0.025)	(0.033)
Race												
Black	-0.036	-0.017	-0.010	-0.029	0.012	-0.014	-0.020	-0.034*	-0.017	-0.013	-0.015	-0.027
	(0.045)	(0.038)	(0.033)	(0.030)	(0.027)	(0.025)	(0.019)	(0.018)	(0.017)	(0.019)	(0.022)	(0.028)
Hispanic	0.018	0.008	-0.008	0.013	0.013	0.030	0.035**	0.025	0.015	0.006	0.011	0.024
	(0.038)	(0.031)	(0.025)	(0.025)	(0.023)	(0.021)	(0.017)	(0.016)	(0.016)	(0.018)	(0.021)	(0.028)
White	0.023	0.015	0.022	0.033	-0.017	-0.023	-0.017	0.000	-0.006	0.009	0.007	-0.006
	(0.035)	(0.031)	(0.028)	(0.025)	(0.021)	(0.020)	(0.013)	(0.012)	(0.011)	(0.012)	(0.015)	(0.022)
Asian	-0.002	-0.005	0.001	0.004	-0.009	-0.002	0.008	0.005	0.008	0.002	0.007	0.005
	(0.016)	(0.012)	(0.009)	(0.009)	(0.008)	(0.008)	(0.007)	(0.007)	(0.007)	(0.008)	(0.009)	(0.011)
N	549	875	1,205	1,545	1,768	1,916	3,145	3,979	4,921	4,738	3,576	2,147
Joint F-statistic p-value	0.455	0.176	0.769	0.669	0.607	0.556	0.320	0.420	0.823	0.194	0.740	0.505

Notes: This table displays covariate balance on baseline characteristics of charter lottery winners and losers for each outcome year. Columns (1)-(12) display OLS regression coefficients from regressions of each baseline characteristic on an any-charter offer dummy. All regressions control for fully-saturated charter application risk sets, grade, and years-since-lottery fixed effects. The joint F-statistic corresponds to the t-statistic of the any-charter offer dummy coefficient from a stacked regression of all baseline characteristics on the any-charter offer dummy. Test scores are standardized by grade and year to have mean zero and unit standard deviation at the state level. Robust standard errors at the attended school level are displayed in parentheses. *** significant at 1% level; ** significant at 5% level; * significant at 10% level.

Table C.10: Charter lottery winners vs. losers: covariate balance

	Enrolled in MA Public School		Has English MCAS		Has Math MCAS	
	Any-charter lottery losers mean	Offered any charter seat	Any-charter lottery losers mean	Offered any charter seat	Any-charter lottery losers mean	Offered any charter seat
	(1)	(2)	(3)	(4)	(5)	(6)
Outcome year 1	0.950	0.014	0.896	0.013	0.906	0.014
		(0.005)		(0.007)		(0.007)
N	9,646		9,646		9,646	
Outcome year 2	0.909	0.025	0.872	0.022	0.870	0.025
		(0.007)		(0.008)		(0.008)
N	9,646		9,646		9,646	
Outcome year 3	0.885	0.015	0.842	0.011	0.839	0.015
		(0.008)		(0.009)		(0.009)
N	9,646		9,646		9,646	
Outcome year 4	0.774	-0.003	0.732	0.005	0.732	0.009
		(0.009)		(0.013)		(0.013)
N	9,646		5,385		5,385	

Notes: This table displays differential attrition between charter lottery winners and losers. Columns (2), (4), and (6) display OLS regression coefficients from regressions of dummies indicating enrollment in a MA public school, availability of English MCAS test score, and availability of math MCAS test score, respectively, on an any-charter offer dummy. Since MCAS is not administered for grade 9, differential attrition estimates for MCAS test scores in outcome year 4 excludes 6th grade applicants, for which grade 9 is the expected grade in the 4th outcome year. All regressions control for fully-saturated charter application risk sets. Means of lottery losers' attrition indicators are displayed in Columns (1), (3), and (5) for reference.

Table C.11: Charter middle school lotteries: analysis sample applicant counts

	Application year										
Application year:	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014
School	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Academy of the Pacific Rim Charter		139	166	292	116	172	145	222	420	467	
Boston Collegiate Charter	155	201	197	210	233	282	264	559	552	625	406
Boston Preparatory Charter		145	206	242	177	192	182	206	209	236	118
Codman											92
Brooke Charter School Roslindale			66	85	79	93					
Brooke Charter School Mattapan								182	103	273	241
Brooke Charter School East Boston									118	217	185
Excel Academy Charter					52	130	118	129	271	318	
Excel Academy Charter School - Boston II									172	235	
KIPP Academy Boston Charter School									104	132	209
MATCH Charter School					295	262	219	490	350	459	238
Uncommon Schools - Roxbury Prep	111	131	132	132	141	151	104	537	451	338	337
Uncommon Schools - Grove Hall								429	451	338	337
Uncommon Schools - Dorchester Prep									451	338	337
UP Academy Charter School of Boston								551	209	173	152

Note: This table displays the number of charter applicants in the analysis sample, by school lottery entered and by application year. Applications are submitted in indicated year for entry into grades 5 or 6 in the Fall semester of the following academic year.

Chapter 3 Data Appendix

To estimate causal effects of suspensions and disciplinary environments on learning, I linked charter lotteries data collected by researchers at MIT's School Effectiveness and Inequality Initiative (SEII) to three administrative datasets provided by the Massachusetts Department of Elementary and Secondary Education (MA DESE): SIMS, SSDR, and MCAS.

SIMS

The SIMS dataset includes demographic information and student-level enrollment records for all MA public schools between school years 2002-3 and 2016-17. SIMS is used to compute three charter enrollment status by grade (charter treatment) and demographic controls.[1] These variables are coded as follows:

- *Charter treatment.* School codes and/or names are used to identify charters. To determine enrollment, since students may switch schools or grades in the same school year, some discretion is needed to determine in which school the student is enrolled. Since charter attendance is a treatment of interest, if in a given school year a student is enrolled for even one day at a charter school, the student is considered enrolled in a charter for that school year. Otherwise, I determine the school in which a student is enrolled based on the maximum number of days attended. This widely adopted definition of treatment is considered conservative because it counts towards treatment lower-achieving students who might leave charters mid-year.[2]

- *Demographic controls.* Dummy variables are created to indicate various demographic characteristics. While SIMS includes a wide number of interesting characteristics to be explored in further work – such as immigrant status and home language – the share students belonging to several of these characteristics is very low. Thus, I focus on more commonly explored demographic characteristics in the literature, such as gender, race, special education status, English language learner status, and low income status. Importantly, the last three characteristics may change over time. As a result, I define

[1]SIMS also includes total suspensions out-of-school, in-school, and an expulsion flag for school years 2003-04 and 2011-12. However, since the SIMS data is aggregated at the enrollment record level, incident dates are not available with which to compute all suspensions occurring prior to test-taking. As a result, I use SDDR as the primary source for data on disciplinary actions. Aggregating all incidents from SSDR at the student level gives similar figures to those reported in SIMS.

[2]Other papers implementing the same or similar strategies include Setren (2017); Abdulkadiroğlu et al. (2017, 2016); Angrist et al. (2016, 2013).

time-varying demographic controls using baseline data (that is, data prior to charter attendance) when analyzing charter and suspensions treatment effects. [3]

SSDR

The SSDR dataset includes student-level information on any drug, violent or criminal-related offenses, as well as any non-drug, non-violent or non-criminal-related offenses committed by the student on school property between school years 2003-04 and 2016-17. Incident dates, offense types, and disciplinary measures taken are available. Students are identified by the same unique identifier as in SIMS. SSDR is used to compute out-of-school and in-school suspensions, and an expulsion dummy for each student. Ever suspended dummies and total days suspended are computed. Three points must be highlighted:

- *Consider all offense types.* While rich data on offense types are available, most incidents in MA public schools entail "non-drug, non-violent or non-criminal-related" offenses only. As a result, a more detailed look into suspension effects by offense type would be limited in power, and is thus deferred to future work.

- *Timing of suspensions.* When analyzing suspensions and expulsions as outcomes, I consider incidents throughout the academic year. However, when estimating the treatment effect of suspensions on test scores and grade progression outcomes, I limit the SSDR data to incidents occurring prior to April of each school year, when the MCAS math and English test season commences.

- *Missing data.* I assume that students not cited in any SSDR incidents were not suspended. For the purposes of estimating unbiased and consistent charter attendance effects on suspensions, and suspension treatment effects on outcomes, this assumption requires no differential SSDR reporting between charters and other MA public schools. If charters are on average better reporters, charter attendance effects on suspensions will be overestimated.

While a thorough investigation of schools' reporting habits is beyond the scope of this paper, it is unlikely that differential SSDR would drive the results in this paper. If anything, since

[3] Furthermore, it is important to note that a student's classification as special education status is a function of the school in which the student is enrolled, and could therefore change if the student enrolls at a charter. In fact, Setren (2017) finds large causal effects of charter enrollment on special education declassification, as charters move special education students into more inclusive classrooms. While the study of school discipline is particularly relevant for special education populations, assessing how declassification and suspensions interact in producing aggregate charter attendance effects is beyond the scope of this project.

charters are consistently under the criticism of over-suspending students, one might expect charters to under-report rather than over-report suspensions. Moreover, higher prevalence of reported suspensions among charters is consistent throughout many US public school districts with varying degrees of data quality collection and reporting standards.

MCAS

The MCAS dataset includes annual MCAS math and English test scores for MA public school students in grades 3 through 8, and 10. Since students may retake the test, I follow the literature in considering test results for the first attempt only.[4] I then standardize the test scores for each subject by grade and year to have mean zero and unit standard deviation at the state level.

Charter lotteries

I use Boston charter middle school lottery records collected by researchers at SEII for charter seats in school years 2004-05 through 2016-17. This sample includes 12 of 17 Boston charters offering middle school grades throughout the sample period.[5] Two points on sample selection are worth emphasizing:

- *Focus on Boston.* Focusing on Boston allows me to use of multiple charter lottery offers as instruments for suspensions and charter attendance in investigating the mechanisms behind Chapter 222's effect.

- *Focus on middle school.* I focus on lotteries for middle school entry (grades 5 and 6) for three reasons. First, as I show in Appendix Figures C.1-C.2, suspensions are primarily a middle school phenomenon in Massachusetts. Second, test scores are available for grades 3-8, allowing for analysis of estimation of test score treatment effects for 1 to 4 years following charter treatment, which is not possible for high school applicants. Finally, middle school applicants have 2-3 baseline grades with test score and discipline histories with which suspension propensities can be computed.

Lotteries take place in the Spring semester for entrance in the following Fall. Charters typically make initial offers and include several other students on a waitlist. When students

[4] In school years 2014 and 2015, Massachusetts experimented switching the standardized test to PARCC exams instead of MCAS. I use the MCAS-corresponding scores provided by MA DESE in the PARCC test score datasets for all PARCC scores.

[5] These figures exclude five charters that specialize in alternative and special education, for which there are no oversubscribed lotteries.

initially offered seats decline attendance, offers are made to waitlisted students. For the purposes of this paper, an applicant is considered a lottery winner if he or she receives either an initial or an off-waitlist offer. (Angrist et al., 2016) presents charter attendance effects on test scores for initially and waitlisted applicants separately.

Importantly, some lottery applicants may be guaranteed a seat at the charter if she/he either has a sibling in the school or fills any special school priorities. These applicants are excluded from analyses as they are not subject to randomization.

Linking datasets

Lottery records and administrative datasets contain identifiable information, such as names and dates of birth, and are thus stored in a restricted access facility at the National Bureau of Economic Research, in accordance with this project's Memorandum of Understanding with MA DESE. Once lottery records are matched to SIMS on names and date of birth, identifiable information are discarded from analyses files. Unique identifiers, available in all administrative datasets, are used to construct a panel dataset tracking applicants across time. This panel dataset includes demographic controls, baseline variables, treatment variables, and outcome variables.

Bibliography

Abadie, A. and M. D. Cattaneo (2018): "Econometric methods for program evaluation," *Annual Review of Economics*, 10, 465–503.

Abdulkadiroğlu, A., J. D. Angrist, S. M. Dynarski, T. J. Kane, and P. A. Pathak (2011): "Accountability and flexibility in public schools: Evidence from Boston's charters and pilots," *The Quarterly Journal of Economics*, 126, 699–748.

Abdulkadiroğlu, A., J. D. Angrist, P. D. Hull, and P. A. Pathak (2016): "Charters without lotteries: Testing takeovers in New Orleans and Boston," *The American Economic Review*, 106, 1878–1920.

Abdulkadiroğlu, A., J. D. Angrist, Y. Narita, and P. A. Pathak (2017): "Research design meets market design: Using centralized assignment for impact evaluation," *Econometrica*, 85, 1373–1432.

Adao, R., M. Kolesár, and E. Morales (2019): "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

Almunia, M., F. Gerard, J. Hjort, J. Knebelmann, D. Nakyambadde, C. Raisaro, and L. Tian (2017): "An analysis of discrepancies in tax declarations submitted under value-added tax in Uganda," *International Growth Centre Project Report*.

Almunia, M. and D. Lopez-Rodriguez (2018): "Under the radar: The effects of monitoring firms on tax compliance," *American Economic Journal: Economic Policy*, 10, 1–38.

Angrist, J. D., S. R. Cohodes, S. M. Dynarski, P. A. Pathak, and C. R. Walters (2016): "Stand and deliver: Effects of Boston?s charter high schools on college preparation, entry, and choice," *Journal of Labor Economics*, 34, 275–318.

Angrist, J. D., S. M. Dynarski, T. J. Kane, P. A. Pathak, and C. R. Walters (2012): "Who benefits from KIPP?" *Journal of policy Analysis and Management*, 31, 837–860.

Angrist, J. D., S. M. Dynarski, T. J. Kane, P. A. Pathak, C. R. Walters, et al. (2010): "Inputs and impacts in charter schools: KIPP Lynn," *American Economic Review*, 100, 239–243.

Angrist, J. D., P. A. Pathak, and C. R. Walters (2013): "Explaining Charter School Effectiveness," *American Economic Journal: Applied Economics*, 5, 1–27.

Ashenfelter, O. (2010): "Modern Models of Monopsony in Labor Markets: a Brief Survey," *IZA Discussion Paper*.

Athey, S. and G. W. Imbens (2017): "The state of applied econometrics: Causality and policy evaluation," *Journal of Economic Perspectives*, 31, 3–32.

Athey, S., G. W. Imbens, and S. Wager (2018): "Approximate residual balancing: debiased inference of average treatment effects in high dimensions," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80, 597–623.

Atkeson, A. and A. Burstein (2008): "Pricing-to-market, trade costs, and international relative prices," *American Economic Review*, 98, 1998–2031.

Auten, G. and R. Carroll (1999): "The effect of income taxes on household income," *Review of economics and statistics*, 81, 681–693.

Autor, D., D. Dorn, and G. Hanson (2013): "The China syndrome: Local labor market effects of import competition in the United States," *The American Economic Review*, 103, 2121–2168.

Azar, J., I. Marinescu, and M. I. Steinbaum (2017): "Labor market concentration," Tech. rep., National Bureau of Economic Research.

Azar, J. A., I. Marinescu, M. I. Steinbaum, and B. Taska (2018): "Concentration in US Labor Markets: Evidence From Online Vacancy Data," Tech. rep., National Bureau of Economic Research.

Bachas, P. and M. Soto (2018): *Not (ch) your average tax system: corporate taxation under weak enforcement*, The World Bank.

Benmelech, E., N. Bergman, and H. Kim (2018): "Strong employers and weak employees: How does employer concentration affect wages?" Tech. rep., National Bureau of Economic Research.

Berger, D. W., K. F. Herkenhoff, and S. Mongey (2019): "Labor Market Power," Tech. rep., National Bureau of Economic Research.

Besley, T. and T. Persson (2014): "Why do developing countries tax so little?" *Journal of Economic Perspectives*, 28, 99–120.

Best, M. C., A. Brockmeyer, H. J. Kleven, J. Spinnewijn, and M. Waseem (2015): "Production versus revenue efficiency with limited tax capacity: theory and evidence from Pakistan," *Journal of Political Economy*, 123, 1311–1355.

Bigio, S. and E. Zilberman (2011): "Optimal self-employment income tax enforcement," *Journal of Public Economics*, 95, 1021–1035.

Blomquist, S. and W. Newey (2017): "The bunching estimator cannot identify the taxable income elasticity," Tech. rep., National Bureau of Economic Research.

Boonzaaier, W., J. Harju, T. Matikka, and J. Pirttilä (2019): "How do small firms respond to tax schedule discontinuities? Evidence from South African tax registers," *International Tax Public Finance*, 26, 1104–1136.

Borusyak, K., P. Hull, and X. Jaravel (2018): "Quasi-experimental shift-share research designs," Tech. rep., National Bureau of Economic Research.

Brockmeyer, A., A. Estefan, J. C. Suárez Serrato, and K. Ramírez (2020): "Taxing Property in Developing Countries: Theory and Evidence from Mexico," Working paper, Duke University.

Brockmeyer, A., S. Smith, M. Hernandez, and S. Kettle (2019): "Casting a wider tax net: Experimental evidence from Costa Rica," *American Economic Journal: Economic Policy*, 11, 55–87.

Brondolo, J., F. Bosch, M. E. Le Borgne, and M. C. Silvani (2008): *Tax administration reform and fiscal adjustment: the case of Indonesia (2001-07)*, 8-129, International Monetary Fund.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and labor market inequality: Evidence and some theory," *Journal of Labor Economics*, 36, S13–S70.

Carrillo, P., D. Pomeranz, and M. Singhal (2017): "Dodging the taxman: Firm misreporting and limits to tax enforcement," *American Economic Journal: Applied Economics*, 9, 144–64.

Chabrier, J., S. Cohodes, and P. Oreopoulos (2016): "What can we learn from charter school lotteries?" *Journal of Economic Perspectives*, 30, 57–84.

Chetty, R. (2009): "Is the taxable income elasticity sufficient to calculate deadweight loss? The implications of evasion and avoidance," *American Economic Journal: Economic Policy*, 1, 31–52.

Costinot, A., D. Donaldson, and C. Smith (2016): "Evolving comparative advantage and the impact of climate change in agricultural markets: Evidence from 1.7 million fields around the world," *Journal of Political Economy*, 124, 205–248.

Crandall, W., E. Gavin, and A. Masters (2019): "ISORA 2016: Understanding Revenue Administration," Tech. rep., Fiscal Affairs Department, International Monetary Fund Paper No 19/05.

Dehejia, R. H. and S. Wahba (1999): "Causal effects in nonexperimental studies: Reevaluating the evaluation of training programs," *Journal of the American Statistical Association*, 94, 1053–1062.

Devereux, M. P., L. Liu, and S. Loretz (2014): "The elasticity of corporate taxable income: New evidence from UK tax records," *American Economic Journal: Economic Policy*, 6, 19–53.

Dix-Carneiro, R. and B. K. Kovak (2017): "Trade liberalization and regional dynamics," *American Economic Review*, 107, 2908–46.

Dobbie, W. and R. G. Fryer Jr (2011): "Are high-quality schools enough to increase achievement among the poor? Evidence from the Harlem Children's Zone," *American Economic Journal: Applied Economics*, 3, 158–87.

Dube, A., J. Jacobs, S. Naidu, and S. Suri (2020): "Monopsony in online labor markets," *American Economic Review: Insights*, 2, 33–46.

Dwenger, N. and V. Steiner (2012): "Profit Taxation and the Elasticity of the Corporate Income Tax Base: Evidence from German Corporate Tax Return Data," *National Tax Journal*, 65, 117.

Fan, H., Y. Liu, N. Qian, and J. Wen (2018): "The Dynamic Effects of Computerized VAT Invoices on Chinese Manufacturing Firms," Tech. rep., National Bureau of Economic Research.

Feldstein, M. (1999): "Tax avoidance and the deadweight loss of the income tax," *Review of Economics and Statistics*, 81, 674–680.

Felix, M. and S. Wang (2021): "A Quasi-Experimental Approach to Identifying Labor Market Boundaries: Evidence from Brazilian Workers," Tech. rep., Working paper.

Fuest, C., A. Peichl, and S. Siegloch (2018): "Do higher corporate taxes reduce wages? Micro evidence from Germany," *American Economic Review*, 108, 393–418.

Gadenne, L. (2017): "Tax Me, but Spend Wisely? Sources of Public Finance and Government Accountability," *American Economic Journal: Applied Economics*, 9, 274–314.

Gordon, R. and W. Li (2009): "Tax structures in developing countries: Many puzzles and a possible explanation," *Journal of Public Economics*, 93, 855–866.

Gruber, J. and J. Rauh (2007): "How elastic is the corporate income tax base," *Taxing corporate income in the 21st century*, 140–163.

Gruber, J. and E. Saez (2002): "The elasticity of taxable income: evidence and implications," *Journal of Public Economics*, 84, 1–32.

Grubert, H. and J. Slemrod (1998): "The effect of taxes on investment and income shifting to Puerto Rico," *Review of Economics and Statistics*, 80, 365–373.

Hainmueller, J. (2012): "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies," *Political Analysis*, 20, 25–46.

Heckman, J. J., H. Ichimura, and P. E. Todd (1997): "Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme," *The Review of Economic Studies*, 64, 605–654.

Hershbein, B., C. Macaluso, and C. Yeh (2019): "Monopsony in US Labor Markets," Tech. rep., Working paper.

Hoang, P. (2019): "International Trade and the Labor Market Power of Firms: Theory and Evidence," Tech. rep., Working paper.

Hoxby, C. M. and S. Murarka (2009): "Charter schools in New York City: Who enrolls and how they affect their students' achievement," Tech. rep., National Bureau of Economic Research.

Hsieh, C.-T. and B. A. Olken (2014): "The missing "missing middle"," *Journal of Economic Perspectives*, 28, 89–108.

Jensen, A. (2019): "Employment Structure and the Rise of the Modern Tax system," Tech. rep., National Bureau of Economic Research.

Katz, L. F., J. R. Kling, and J. B. Liebman (2001): "Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment*," *The Quarterly Journal of Economics*, 116, 607–654.

Kawano, L. and J. Slemrod (2016): "How do corporate tax bases change when corporate tax rates change? With implications for the tax rate elasticity of corporate tax revenues," *International Tax and Public Finance*, 23, 401–433.

Keen, M. and J. Slemrod (2017): "Optimal tax administration," *Journal of Public Economics*, 152, 133–142.

Khan, A. Q., A. I. Khwaja, and B. A. Olken (2016): "Tax farming redux: Experimental evidence on performance pay for tax collectors," *The Quarterly Journal of Economics*, 131, 219–271.

Kleven, H. J., M. B. Knudsen, C. T. Kreiner, S. Pedersen, and E. Saez (2011): "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark," *Econometrica*, 79, 651–692.

Kleven, H. J., C. T. Kreiner, and E. Saez (2016): "Why can modern governments tax so much? An agency model of firms as fiscal intermediaries," *Economica*, 83, 219–246.

Kovak, B. K. (2013): "Regional effects of trade reform: What is the correct measure of liberalization?" *The American Economic Review*, 103, 1960–1976.

Kume, H., G. Piani, and C. Souza (2003): "A política de importação no período 1987-1998: Descrição e avaliação," *Abertura Comercial Brasileira nos Anos Noventa: Impacto sobre Emprego e Salário. MTB/IPEA, Rio de Janeiro.*

Lamadon, T., M. Mogstad, and B. Setzler (2019): "Imperfect Competition, Compensating Differentials and Rent Sharing in the US Labor Market," Tech. rep., National Bureau of Economic Research.

Lemgruber, M. A., M. A. Masters, and M. D. Cleary (2015): *Understanding revenue administration: an initial data analysis using the revenue administration fiscal information tool*, International Monetary Fund.

MacKenzie, G. (2018): "Trade and Market Power in Product and Labor Markets," Tech. rep., Working paper.

Manning, A. (2003): *Monopsony in motion: Imperfect competition in labor markets*, Princeton University Press.

McFadden, D. (1978): "Modeling the choice of residential location," *Transportation Research Record*.

——— (1981): "Econometric models of probabilistic choice," *Structural analysis of discrete data with econometric applications*, 198272.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica: Journal of the Econometric Society*, 71, 1695–1725.

Menezes-Filho, N. A. and M.-A. Muendler (2011): "Labor reallocation in response to trade reform," .

Naritomi, J. (2019): "Consumers as tax auditors," *American Economic Review*, 109, 3031–72.

Nimczik, J. S. (2017): "Job mobility networks and endogenous labor markets," .

Pomeranz, D. (2015): "No taxation without information: Deterrence and self-enforcement in the value added tax," *American Economic Review*, 105, 2539–69.

Saez, E., J. Slemrod, and S. H. Giertz (2012): "The elasticity of taxable income with respect to marginal tax rates: A critical review," *Journal of Economic Literature*, 50, 3–50.

Schmutte, I. M. (2014): "Free to move? A network analytic approach for learning the limits to job mobility," *Labour Economics*, 29, 49–61.

Schubert, G., A. Stansbury, and B. Taska (2019): "Getting labor markets right: Outside options and occupational mobility," Tech. rep., Working paper.

Setren, E. (2017): "The Impact of Specialized Services vs. High Quality General Education on Special Education and English Language Learner Students," .

Slemrod, J. (2001): "A general model of the behavioral response to taxation," *International Tax and Public Finance*, 8, 119–128.

——— (2019): "Tax compliance and enforcement," *Journal of Economic Literature*, 57, 904–54.

Slemrod, J. and W. Kopczuk (2002): "The optimal elasticity of taxable income," *Journal of Public Economics*, 84, 91–112.

Steinberg, M. P. and J. Lacoe (2017): "What do we know about school discipline reform? Assessing the alternatives to suspensions and expulsions," *Education Next*, 17, 44–53.

Stuart, E. A. (2010): "Matching methods for causal inference: A review and a look forward," *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25, 1.

Suárez Serrato, J. C. and O. Zidar (2016): "Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms," *American Economic Review*, 106, 2582–2624.

——— (2018): "The structure of state corporate taxation and its impact on state tax revenues and economic activity," *Journal of Public Economics*, 167, 158–176.

Topalova, P. (2007): "Trade liberalization, poverty and inequality: Evidence from Indian districts," in *Globalization and Poverty*, University of Chicago Press, 291–336.

——— (2010): "Factor Immobility and Regional Impacts of Trade Liberalization: Evidence on Poverty from India," *American Economic Journal: Applied Economics*, 2, 1–41.

Tortarolo, D. and R. D. Zarate (2018): "Measuring Imperfect Competition in Product and Labor Markets. An Empirical Analysis using Firm-level Production Data," .

Tucker, L. (2017): "Monopsony for Whom? Evidence from Brazilian Administrative Data," .

Ulyssea, G. and V. Ponczek (2018): "Enforcement of Labor Regulation and the Labor Market Effects of Trade: Evidence from Brazil. IZA Discussion Papers 11783," *Institute for the Study of Labor (IZA).*

UNU-WIDER (2021): *Total tax revenue, including social contributions as a share of national GDP*.

Weber, C. E. (2014): "Toward obtaining a consistent estimate of the elasticity of taxable income using difference-in-differences," *Journal of Public Economics*, 117, 90–103.

Zarate, R. D. (2016): "Chinese Import Competition effects on Domestic Firms: Disentangling the role of firms on labor markets," .