Essays on Public Finance in Developing Countries

Inauguraldissertation
zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaften
der Universität Mannheim

vorgelegt von

Daniel Overbeck

im Frühjahrs-/Sommersemester 2025

Abteilungssprecher	Prof. Dr. Thomas Tröger
Referent	Prof. Dr. Eckhard Janeba
Koreferent	Prof. Dr. Arthur Seibold

Tag der Verteidigung	26.05.2025

Acknowledgements

First and foremost, I thank my main supervisor, Eckhard Janeba, for his invaluable guidance and support, his steady confidence in my work, and the welcoming atmosphere he creates at the chair. I would also like to sincerely thank Arthur Seibold for being my advisor as well as for his encouragement and key advise at decisive crossroads throughout the process of my PhD.

Throughout my studies, I enjoyed and benefitted greatly from the companionship of many amazing and inspiring people within the University of Mannheim and beyond. I thank Johannes Gallé for many online and offline meetings which kept the spirits up and the research going. Further, I thank Eliya Lungu for paving our way in Zambia.

I would also like to thank my other co-authors Tobias Seidel, Edson Severnini, Rodrigo Oliveira, and especially Nadine Riedel (2x), from whom I have learned so much. I want to thank Lukas Hack and Felix Rusche for countless discussions about research and life and moreover, Sebastian Seitz, Karl Schulz, Felix Köhler, Paul Steger, Jakob Schmidhäuser, Henning Schatz, Laura Montenbruck, Alexander Kann, and my entire cohort for being great colleagues. I am grateful for (financial) support from the Graduate School of Economics and Social Sciences, the International Growth Center, and the UNU-WIDER.

I would like to thank my siblings Benjamin, Clara, Leonard, and most of all my parents Bettina and Ludger for their continuous support, their indestructible optimism and the intellectual curiosity they spurred in me. I am extremely grateful for having shared this journey with my two amazing children Oskar and Romi, who were born right before the start and right before the end of the PhD.

Above all, I cannot put into words the gratitude I feel towards my wonderful wife Resi. Without her unconditional love and support in every dimension none of this would have been possible. This thesis is dedicated to her.

Contents

List of Figures

List of Tables

Preface

In most of what is commonly known as the developing world, tax collections are low. While rich OECD countries collect about a third of their GDP in taxes on average, countries in Sub-Saharan Africa or South Asia – the most populous regions of the world – only collect about 16% and 12% of their GDP, respectively (ATAF, 2018; World Bank, 2025). Low tax revenues can hinder public goods provision through e.g. inadequate investments into infrastructure or funding of welfare programs. They also increase states' dependencies on external lenders and aid. The recent closure of USAID, which humanitarian organisations expect to "kill millions" (Caritas, 2025), is a timely reminder that for lower-income countries, domestic resource mobilization through taxation is imperative.

Beyond generating revenues, taxation represents a key determinant of how states and societies function and also how these two interact. Tax compliance is a regular aspect of firms' and individuals' economic lives while the ability to tax and redistribute – in other words, managing public finances – is one of the most important function of governments. This includes incentivizing certain (firm) behaviour by selectively increasing or decreasing prices through taxes or tax exemptions and subsidies, respectively.

Relative to what is known for rich countries, the economic literature has yet provided only limited evidence on the role and the effects of taxes and subsidies in lower-income countries. While this represents an emerging research area, the scarcity of evidence holds particularly true for the poorest countries.

This thesis advances our understanding of public finance along three dimensions. Chapter 1 uncovers the existence of a hybrid tax system in Zambia – one of the least developed countries in the world (United Nations, 2023) – where instead of following the legal tax schedule, tax payments are simply determined through bargaining with tax collectors. Chapter 2 shows how the implementation of Special Economic Zones in India, which are essentially place-based tax exemptions for firms, led to sectoral shifts from agriculture to manufacturing and service industries in nearby villages. Chapter 3 analyzes how firms respond to the South African carbon tax, which is the first of its kind in any developing country. Each chapter is self-contained.

In Chapter 1, which is joint work with Eliya Lungu, I study the taxation of small firms in Zambia. Their taxation is highly relevant as small firms make up 80-90% of taxpayers and account for about 40% of GDP in the average developing country (World Bank, 2011b, 2019). I apply empirical as well as theoretical methods to show that bargaining over tax payments is an important feature of tax compliance and enforcement in lower-income countries such as Zambia. The empirical analysis is twofold. First, analyzing the universe of administrative tax filings from Zambia, I document sharp bunching in (i) dominated regions above tax schedule discontinuities, inconsistent with standard models of tax compliance and (ii) at round number tax payments, implying that certain payments are being targeted. Second, I provide additional qualitative evidence from a survey I conducted of Zambian firms via an external provider. The findings suggest that discussing tax payments with tax officials before filing taxes is widespread, in line with tax payments being the outcomes of bargaining. Such bargaining over taxes is consistent with fact (ii), as bargaining outcomes are often round and salient numbers, and with fact (i), because tax schedule discontinuities restrict the set of feasible bargaining outcomes. The theoretical model I propose generalizes the conventional Allingham and Sandmo (1972) model to allow for bargaining as a mode of tax compliance. I show that as long as state capacity is low, which is the case for developing countries, bargaining over taxes leads to Pareto-improvements for both taxpayers and the state. This result rationalizes why bargaining is observed for small firms in low-income countries but not for larger firms in more advanced economies.

Chapter 2 is joint work with Johannes Gallé, Nadine Riedel and Tobias Seidel and has been published in the Journal of Public Economics (c.f. Gallé, Overbeck, Riedel, and Seidel (2024)). In this chapter, I study the impact of the implementation of Special Economic Zones (SEZs) in India. According to the UNCTAD (2019), the total number of SEZs worldwide increased from 500 in 1995 to about 5,400 in 2018 - with the vast majority of the new zones being located in developing economies. Therefore, they represent a highly prevalent policy tool. I quantify the local economic impact of SEZs established in India between 2005 and 2013 econometrically using a spatial difference-in-differences design. Relying on a novel dataset that combines census information on the universe of Indian firms with geo-referenced data on SEZs, I find that the establishment of SEZs increased local manufacturing and service employment, with positive spillovers up to 10 km from the SEZ area. The analysis shows that the gains in manufacturing and service employment were accompanied by a decline in agricultural labor, especially for women, suggesting that the policy contributed to structural change. In further analysis, I document that significant local employment effects occur across different types of SEZs: privately and publicly run zones, and SEZs with different industry designations. A back-of-the-envelope calculation shows that the costs of the program (i.e. foregone tax revenues) slightly exceed the welfare gains from additional employment. Overall, I still interpret the findings to dispel the general pessimism about zone programs in developing countries outside of China.

Chapter 3, which is coauthored with Johannes Gallé, Rodrigo Oliveira, Nadine Riedel and Edson Severnini, concerns the role of taxation for environmental policy. It provides the first comprehensive analysis of how firms in emerging economies respond to carbon taxation, leveraging detailed administrative data from South Africa – a potential trailblazer for other developing countries with limited state capacity amid the growing global push for carbon pricing. I examine the dynamic impacts of the carbon tax on firm-level outcomes – such as profits, sales, capital, and labor inputs – across manufacturing and mining firms, which are key sectors in the context of the carbon tax. Contrary to concerns that carbon taxes may hinder economic growth or reduce employment, the findings show no evidence of negative average impacts on firm performance or jobs. However, this overall result masks significant heterogeneity in the tax's effects across sectors, driven by the sector-specific design elements of the South African carbon tax. Firms expecting higher effective tax rates may have intensified their use of emission-intensive machinery and depreciated capital in anticipation of the tax. This behavior appears to stem from firms resolving regulatory uncertainty or seeking to recover costs from stranded assets.

In conclusion, this thesis make three main contributions to the literature on public finance in developing countries. First, it shows how tax compliance of small firms deviates strongly from the predictions of standard economic models. Second, it highlights how place-based tax exemptions for firms can lead to beneficial outcomes for surrounding regions. Third, it provides evidence that carbon taxation, which is high up on policy agendas, may not be as harmful to economic performance as expected. These novel findings might inform policymakers on a broad variety of public finance issues and importantly, show how optimal policy recommendations may need to be adjusted to account for the institutional realities of developing countries.

Chapter 1

Bargaining Over Taxes

Joint with Eliya Lungu.

1.1 Introduction

In low- and middle income countries, state capacity is low (Besley and Persson, 2013). As a consequence, formal institutions often struggle to execute basic governmental functions effectively. Instead, informal institutions play an essential role. For example, credit markets, insurance systems as well as public good provision at the local level are partly organised through informal interactions between citizens and local elites or bureaucrats (e.g., Angelucci and De Giorgi, 2009; Olken and Singhal, 2011; Udry, 1994). Very little is known about the relevance of informal institutions on the tax side, however. As generating revenue through tax collection becomes increasingly important for developing countries (International Monetary Fund, 2024), it is crucial to understand informal institutions' role and potential in this process. In this chapter, we study how informal institutions, and in particular, interactions between bureaucrats and citizens, shape tax compliance and enforcement *inside* the official tax system.

The setting we study is the taxation of small firms, which comprise the vast majority of taxpayers in most low-and middle income countries.[1] In particular, we study firms subject to turnover taxation in Zambia, a lower middle income and least developed country (United Nations, 2023). Analyzing administrative data on the universe of tax filings, we document two novel facts – both are inconsistent with standard models of tax compliance. First,

[1] Small firms make up 80-90% of taxpayers and account for about 40% of GDP in the average developing country (World Bank, 2011b, 2019).

firms bunch e.g. an excess number of firms locate in dominated regions above tax schedule discontinuities. Second, firms bunch at turnover amounts which imply round number tax payments (not necessarily round turnover). We further provide novel evidence from our own survey showing that discussing tax payments with tax officials before filing taxes is widespread, which is in line with tax payments being the outcomes of bargaining. Such bargaining over taxes rationalizes both empirical facts, as bargaining outcomes are often round and salient numbers, and because tax schedule discontinuities restrict the set of feasible bargaining outcomes. We provide several pieces of evidence, including a randomized survey experiment, to show that alternative explanations based on audit probabilities, optimization frictions, or mistakes cannot rationalize these bunching patterns. Finally, we generalize the conventional Allingham and Sandmo (1972) model of tax evasion to allow for bargaining as a mode of tax compliance. We show that, as long as state capacity is sufficiently low, bargaining over taxes leads to pareto-improvements for both taxpayers and the state.

Our empirical analysis builds on a novel and comprehensive dataset comprising the universe of more than 5.3 million turnover tax filings from 2015 until 2021 in Zambia. The data is at a monthly frequency and allows tracking firms over time and, importantly, across two major tax reforms in 2017 and 2019. Two key empirical patterns emerge. First, we measure bunching responses to tax schedule discontinuities. During the years 2017-2018, the Zambian turnover tax schedule features a linear tax rate and several tax brackets in which tax liability increases discretely by a fixed payment. Thus, net-of-tax turnover drops discretely at each bracket threshold. Such thresholds in the tax schedule, which induce drops in net-of-tax turnover imply that it is a strictly dominated choice to have (or report) turnover in a certain region above each threshold. Specifically, firms above tax bracket thresholds would be strictly better off when reducing turnover to below the threshold. Standard models of tax compliance would therefore predict bunching below the threshold but no mass above the threshold (Kleven and Waseem, 2013). In contrast, our data reveal strong bunching *above* each threshold. We find more than twice as many firms to report turnover just above the threshold compared to what our estimated counterfactual distribution predicts. The bunching is very sharp with most tax returns only exceeding the threshold by less than one Zambian currency unit, equivalent to around USD 0.05. Such bunching above thresholds is inconsistent with any standard preferences and thus a puzzling fact.[2]

[2]Intuitively, standard preferences are such that utility is always increasing in consumption and that gener-

Second, we find strong and sharp bunching at turnover amounts which imply round number tax liabilities, e.g., tax liabilities which are multiples of 10, 50, or 100. The Zambian setting is particularly appealing for this exercise as, before the tax brackets were introduced in 2017, the turnover tax schedule was flat at a rate of 3%. We can therefore clearly distinguish bunching at round number tax liabilities from the well-known phenomenon of bunching at round numbers of the taxable income itself (Carrillo, Donaldson, Pomeranz, and Singhal, 2022; Kleven and Waseem, 2013). To see this, note that with a 3% tax rate, round liabilities often imply particularly odd amounts of turnover. Pooling all tax returns filed between 2015-2016, we document that more than 40% of the filed returns were such that the resulting liability was a multiple of 10. Importantly, these figures are not adjusted ex-post by the tax authority, but represent the raw numbers as appearing in the tax returns. This finding is inconsistent with standard models because there is neither an incentive to report nor any other reason for having such amounts of turnover. Instead, it suggests that the tax schedule gets *inverted* and certain payment amounts are targeted.

Exploiting the panel structure of the data, we show that the two empirical patterns are strongly connected. Firms that exhibit round payments prior to the reform have a much larger probability of bunching above thresholds than those that do not (or less so). This correlation is highly significant and robust to several alternative specifications. Leveraging administrative data on tax audits, we control for audit experience as well as various other characteristics and find that they neither change the significant effect of round payments nor have a significant effect themselves.

These empirical facts are puzzling as they are inconsistent with predictions from standard models of tax compliance. Thus, we present complementary results from a survey we conducted of 517 firms registered for turnover tax in Lusaka, the capital and largest city in Zambia. On average, the sample of surveyed firms matches the administrative data in terms of size, sector and gender. Several insights can be retrieved from the survey results.

The most striking result concerns the interactions between taxpayers and the Zambia Revenue Authority (ZRA). Nearly half of all respondents report to discuss the tax payments they are going to make with officials from the ZRA *before* filing their tax returns, or that such discussions are common. Of those, again half explicitly state that these discussions serve the purpose of finding agreements with officials on what should be paid. Put differently, firms bargain with officials over their tax payments. In line with the observations outlined above, the respondents state that bargaining evolves around the payment itself

ating income is always costly (c.f. Kleven, 2016; Kleven and Waseem, 2013; Saez, 2010).

rather than the correct turnover. Furthermore, we estimate a significant and negative relationship with discussions over tax payments and the perceived probability of being formally audited for tax purposes. Bargaining over taxes could therefore serve as a preemptive measure to forego the formal procedure of filing and potentially being audited.

Motivated by these findings, we propose a theoretical framework of tax bargaining which can rationalize the empirical facts. We consider a risk averse firm choosing its tax payment conditional on its true tax liability (Allingham and Sandmo, 1972). The tax authority receives utility from tax revenues and penalty payments but incurs costs from auditing. We show that there exists a region of pareto-improving tax payments in this situation: relative to the standard non-cooperative case where audits induce risk for the taxpayer and costs for the tax authority, both parties are better off when agreeing on certain payments and not making any audits. This creates a potential surplus, which, in the model, is divided via Nash-bargaining.

The framework rationalizes both empirical facts. First, bargaining often leads to round-number outcomes, as round numbers can serve as focal points (Janssen, 2006; Pope, Pope, and Sydnor, 2015; Schelling, 1960). In the model, we also show that the set of payments which is bargained over is detached from a taxpayer's true liability, which is consistent with the notion that bargained payments often simply end up on round figures. Second, notches (i.e. discrete jumps in tax liability) effectively introduce regions of payments which cannot be reached anymore. This restricts the bargaining set and agreed upon payments accumulate just below but also just *above* the threshold. Importantly, our model also offers an explanation for why bargaining over taxes might be especially prevalent in less developed economies, with limited state capacity. A key feature of state capacity is the efficiency of tax audits. This efficiency crucially hinges on third-party reporting which is oftentimes non-existent for small businesses in less developed economies (Kleven, Kreiner, and Saez, 2016). Our model demonstrates that as countries develop and build state capacity, the increased audit efficiency ultimately removes the scope for bargaining.

Our analyses include several additional pieces of evidence that rule out competing explanations for the observed bunching behavior. One prime concern might be that, in principle, the fear of being audited by the tax authority could incentivize firms to bunch above the threshold instead of below by itself, even in the absence of bargaining. Firms might simply trade a lower audit probability above the threshold against a larger tax payment. We address this concerns in four ways. First, we find no significant relationship between the event of being audited on whether a firm bunches above a threshold or not. Second, we find no substantial differences in empirical audit probabilities above versus below thresholds. Third, simulations from a standard model of tax evasion (Allingham and Sandmo, 1972; Kleven, Knudsen, Kreiner, Pedersen, and Saez, 2011) show that even if perceived audit probabilities above- and below the threshold differ strongly, bunching above the threshold is highly unlikely to occur. We demonstrate that even under extreme assumptions where the probability of being audited jumps from 10% above to 30% below the threshold, only firms that evade at least 83% of their turnover would choose to stay above the threshold. Lastly, we run a randomized survey experiment to find that also shifting a firm's perceived audit probability upwards does not increase its stated propensity to bunch above a threshold. We provide further evidence to rule out optimization frictions or mistakes as an explanation for the observed bunching behaviour.

It is important to note, that the bargaining situations considered in this paper are distinct from mere corruption (Hindriks, Keen, and Muthoo, 1999; Khan, Khwaja, and Olken, 2016) where the tax collector receives bribes in exchange for lying. Bunching above thresholds– instead of below can hardly be rationalized with such collusion between taxpayer and tax collector. Both parties would always have a clear incentive to declare turnover below the threshold and share the difference in tax liability. Our survey also elicits that the share of firms bribing tax collectors is much smaller than the share of firms engaging in bargaining. Finally, our model suggests that the government also benefits from bargaining through saving audit costs. Instead, we argue that incentive schemes for tax collectors may increase their effort in bargaining. We find that once a tax office hits its revenue target – and bonuses are being paid out to collectors – less firms tend to bunch above thresholds.

This paper contributes, first and foremost, to the literature on how tax administration is shaped by the institutional context of developing countries (e.g., Besley and Persson, 2014; Gadenne and Singhal, 2014; Gordon and Li, 2009a; Okunogbe and Tourek, 2024). Among others, Olken and Singhal (2011) document that, for a large share of the population, tax collection and public service provision is organized entirely outside of formal institutions. Within formal institutions, recent studies have shown that engaging non-state actors such as local elites in tax collection (Balan, Bergeron, Tourek, and Weigel, 2022) or subsidy targeting (Basurto, Dupas, and Robinson, 2020) can overcome informational barriers and thus produce more efficient outcomes in low-income countries. Along these lines, Okunogbe and Pouliquen (2022) and Aman-Rana and Minaudier (2024) show that the digitization of tax collection and thus removal of personal interactions between taxpayers and tax collectors can partly lead to lower revenue collection. To the best of our knowledge, this paper is the first to show that such interactions can serve as a mechanism to determine tax payments through bargaining – an informal arrangement from which both parties, taxpayers as well as the tax authority, benefit. Aman-Rana, Minaudier, and Sukhtankar (2023) study another such informal arrangement, namely apparent corruption. In line with our interpretation, they argue that such arrangements constitute devices to overcome the low state capacity of governments in lower income countries. Regarding corrupt tax officials, Khan, Khwaja, and Olken (2016) and Hindriks, Keen, and Muthoo (1999) analyse situations where officials and taxpayers collude and lie about true tax liabilities. In this paper, we present empirical patterns that can hardly be rationalized by the presence of corruption. Instead, our finding of bunching above thresholds is consistent with bargaining even in the absence of corruption.

We further add to the broader literature on taxation and development from which, so far, particularly little is known about firm responses to taxation in low-income countries.[3] Most existing evidence is from Rwanda where it has been shown that a substantial share of registered firms is economically inactive (Mascagni, Santoro, Mukama, Karangwa, and Hakizimana, 2022) and many firms simply always file the same amount (Tourek, 2022). Relatedly, Almunia, Hjort, Knebelmann, and Tian (2023) document how firms in Uganda depart from alleged profit-maximizing behavior when filing taxes. We document similar albeit novel facts for the case of Zambia – another low income country – and offer bargaining over taxes as a yet underexplored explanation. Methodologically, our study contributes to and builds on work investigating how incentives are navigated by firms in low enforcement environments (Anagol, Davids, Lockwood, and Ramadorai, 2022; Bachas and Soto, 2021; Best, Brockmeyer, Kleven, Spinnewijn, and Waseem, 2015; Kleven and Waseem, 2013). In particular, by contrasting the case of small firms in a least developed country to what is known about larger firms in countries at other stages of development, we highlight that the established predictions may not hold universally across the developing world.

The remainder of the paper is structured as follows. After section 1.2 details the institutional background, section 1.3 explains the data sources and empirical methodology. Section 1.4 shows the results from the administrative data and the survey. Section 1.6 rules out other explanations than bargaining. The theoretical framework is presented in Section 1.5. Section 1.7 concludes.

[3] Zambia switched frequently between low income and lower-middle income status according to the World Bank and only moved from low income to lower-middle income status in 2023.

1.2 Background

Zambia is a lower middle-income country and classifies as a least developed country according to the United Nations (2023). In 2021, it had a population of 20 million people, a GDP-per-capita of USD 1137 PPP and a tax-to GDP ratio of about 16%. It thus closely resembles the average Sub Saharan African country along these dimensions (ATAF, 2018). Its tax-to-GDP ratio is low compared to the OECD country average of 34% and just above the ratio deemed necessary to meet basic needs of citizens and businesses (Gaspar, Jaramillo, and Wingender, 2016).

Business taxation. In principle, every business in Zambia is required to be registered with the tax authority. However, as in all lower income countries, the majority of firms in Zambia is small and informal (not registered with the tax authority). Official statistics estimate that the informal sector accounted for nearly 90% of the country's employment in 2014 (Ministry of Labour and Social Security, 2018). Among the firms that are under the tax net, there is a crucial size distinction for how the tax base is determined. Businesses with annual turnover above ZMK 800,000 ($\approx$ USD $31,000$) are liable for the corporate income tax (CIT) where taxes apply to profits. Additionally, businesses are required to register for the Value-Added-Tax (VAT). If a business falls below this threshold it is liable for turnover tax in which, akin to a pure sales tax, taxes apply on turnover – at substantially lower rates. Such systems are applied widely among lower income countries with the intention to simplify tax compliance by allowing for a simple measure of the tax base and thus tax liability.[4] More than 80% of Zambia's taxpayer population is registered only for turnover tax, highlighting that also the formal sector constitutes mainly of small firms. For firms under turnover tax, voluntary VAT registration is possible but very uncommon. Turnover taxes thus matter for a large share of the population and hence arguably for welfare. However, turnover taxes only account for less than 5% of total tax revenues.

It is common in Zambia to make tax payments in cash or per cheque. In 2016, more than 90% of tax payments were done in this way. Notably, this figure stood at about 20% in 2021, and thus, has decreased drastically throughout our study period. These numbers are aggregates for all tax types and likely to be larger for turnover tax payments (Zambia Revenue Authority, 2022).

[4]Other examples of turnover tax systems in Africa include Nigeria, Kenya, Ghana, Uganda, Rwanda, Tanzania, Cameroon. Hoy, Scot, Oguso, Custers, Zalo, Doino, Karver, and Pillai (2024) provide a recent overview of such systems in Africa.

Table 1.1: Turnover Tax Schedule 2017-2018

Turnover (in ZMK)	Tax Liability (in ZMK)
0 - 3000	0
3000 - 4,200	3% of monthly turnover above 3,000
4,200.01 - 8,300	225 per month + 3% of monthly turnover above 4,200
8,300.01 - 12,500	400 per month + 3% of monthly turnover above 8,300
12,500.01 - 16,500	575 per month + 3% of monthly turnover above 12,500
16,500.01 - 20,800	800 per month + 3% of monthly turnover above 16,500
Above 20,800	1,025 per month + 3% of monthly turnover above 20,800

Notes: This table depicts the turnover tax schedule which was in place in Zambia throughout 2017 and 2018. Tax liability increases discretely at 4200, 8300, 12500, 16500 and 20800. For example Liability is $36 = (4200 - 3000) \times 0.03$ at 4200 and 225 at 4200.01. This creates discontinuities in the budget set, referred to as notches.

Another crucial feature of tax compliance are personal interactions with tax officials. According to the World Bank Enterprise Surveys, more than 80% of small firms state to regularly visit tax officials in person. This ratio is slightly higher but comparable to other sub-saharan countries ($\sim 60\%$) but twice as large as the world average ($\sim 40\%$). On the other hand, corruption through tax collection is relatively low. Only 3.4% of small firms pay bribes to tax officials compared to the world average of 9.8%.[5]

Turnover tax reforms. The turnover tax in Zambia applies to monthly sales and therefore returns have to be filed and payments have to be made on a monthly basis. The system was introduced in 2009 and initially levied a tax rate of 3% on gross sales of firms with annual turnover below the CIT threshold. In 2017, the flat rate of 3% was replaced by a graduated bands schedule: turnover thresholds were introduced above which fixed payments had to be paid. Additionally, the turnover in excess of the threshold, was taxed at 3%. This created 7 tax brackets, described in Table 1.1.

Each threshold (except the first) implies a discrete jump in tax liability and thus in the average tax rate. For example, with a turnover of 4200, a firm had a tax liability of 36 ($= (0.03 \times (4200 - 3000))$). With 4200.01, the tax liability becomes 225. The average tax rate thus jumps from $\sim 1\%$ to $\sim 5\%$. Such discrete changes in the average tax rate are referred to as *notches* in the literature (Kleven and Waseem, 2013; Slemrod, 2013). Our empirical analysis exploits these notches to study firm responses to taxation. For this exercise, the Zambian turnover tax is especially appealing as neither the thresholds nor the payment amounts coincide with salient round numbers which could serve as psychologically

[5]The data can be retrieved from `https://www.enterprisesurveys.org/en/data/exploreeconomies/2019/zambia`. Last accessed: June 12th, 2024.

appealing focal points. The reasoning behind the tax reform was that applying the same tax rate to firms of all sizes was perceived as regressive. In contrast, the tax schedule became partly regressive only after the tax reform. Fixed payments combined with a linear tax rate imply that the average tax rate is decreasing in turnover within each bracket. The tax reform therefore represented an ill-fated attempt to alleviate the concern of regressivity. In 2019, the schedule was again replaced by a flat schedule, this time with a rate of 4%.

1.3 Data & Methodology

In this section, we first explain the different data sources from both administrative records and the survey. Then, we describe the empirical methodology.

1.3.1 Data sources

Administrative data on tax returns. The administrative data used in this study has been provided by the ZRA and comprise the universe of turnover tax returns that have been filed in Zambia between 2015-2021. Each tax return contains information on the taxpayer identification number, total turnover amount, the tax liability, the tax office responsible for the taxpayer, the date a firm has registered for turnover tax and the period to which the tax return relates. The period is usually a calendar month. However in some cases, returns are covering multiple months. In total, the dataset includes more than 5.3 million firm-month observations. Additional sector information is available for a smaller subset.

Administrative data on tax audits. We further observe audits under turnover tax which can be matched to taxpayers and the time periods to which they apply. While the reporting frequency for turnover tax is monthly, audits may apply to longer periods which means that multiple tax declarations are inspected. The audit data provides information on whether the audit resulted in payments of penalties or owed taxes.

Survey data. We complement the administrative data with a survey of Zambian firms. The survey data was collected throughout November and December of 2023.[6] A local team of six surveyors, hired through the survey provider Center for Evaluation and Development (C4ED), collected responses to 30-40 questions of 517 firms under turnover

[6] The survey was designed and contracted by Daniel Overbeck alone, independently of Eliya Lungu or ZRA.

taxation in Lusaka, the capital city of Zambia. The interviews were based on a customized questionnaire and conducted in person. Access to information on taxpayer addresses is restricted by ZRA policy. Thus surveyors randomly approached small firms in Lusaka's main business areas to find survey participants. As questions on taxation can be sensitive, we believe an in-person approach to build a more trusting environment and thus to be better suited than phone interviews in this case (Blattman, Jamison, Koroknay-Palicz, Rodrigues, and Sheridan, 2016).
As survey respondents can not be linked to the administrative data, respondents were initially asked whether their business was registered for turnover tax or not and a number of questions served as checks as to whether their response was credible. Additionally, only if firms stated to have average monthly turnover of less than ZMK 70,000 ($\approx$ turnover tax threshold/12), the interview was continued. It was emphasized that the study was conducted independently of the ZRA and in fact, not even the surveyors were aware of any connection.
After the initial entry questions, firm characteristics such as sector, gender of the owner, number of employees and average turnover among others were collected. On average, survey sample firms are comparable to firms appearing in the administrative data. Further questions regarding accounting– and tax filing practices followed. The final part of the interview consisted of a randomized lab-in-the-field experiment, specifically designed to investigate potential channels driving the empirical facts established from the administrative data.

1.3.2 Empirical approach

The empirical analyses of the administrative data largely rely on the estimation of *bunching* behavior (Kleven, 2016), a method to estimate local responses to certain thresholds. In particular, we compare the distribution of turnover tax returns observed in the data to a counterfactual distribution we estimate from a subsample of the data, omitting certain points of interest.

Bunching. We estimate two kinds of bunching behavior. The first kind we are interested in focuses on the tax bracket thresholds. More precisely, we investigate behaviour around the notches at which tax liability increases discretely. Through the lenses of conventional economic models, such discontinuities offer clear incentives for firms to behave in a certain way. Here, we focus on the simple intuition behind these incentives.[7] To understand the incentives, one should consider the budget set a firm is facing i.e. its net-of-tax turnover when the tax schedule features a notch. Panel (a) of Figure 1.1 sketches the budget set as a function of turnover around a notch akin to the ones in Table 1.1. Above the threshold, net of tax turnover drops and only exceeds net-of-tax turnover at the threshold when turnover is substantially larger. The red segment thus delineates a 'dominated range' in which a firm would be strictly better off with either lower or larger turnover. The incentives that are created by the notch are expected to result in a particular density of tax returns around the notch. Panel (b) of Figure 1.1 delineates this predicted density of tax returns (in the absence of optimization frictions). Firms are expected to be either reporting turnover just below the threshold or substantially above the threshold. As no firm would have turnover in the dominated region, we would expect a hole in the distribution just above the threshold.

To now estimate the extent of bunching below the threshold we proceed as follows. We coarsen the turnover data into bins of width 100 ZMK and draw on the approach developed by Kleven and Waseem (2013), in which we control for the *affected* turnover range, in which we would expect either excess mass (bunching) or missing mass (too little tax returns). To also account for potential bunching at round turnover as well as round liability amounts, we include dummies for both cases. The estimating equation reads

$$\begin{aligned} c_j = & \sum_{i=0}^{p} \beta_j (z_j)^i + \sum_{i=z_L}^{z_U} \gamma_i \mathbf{1}[z_j = i] + \sum_{n \in \mathbf{N}} \xi_n \mathbf{1}\left[\exists k \in j | \frac{T(k)}{50} = n\right] \\ & + \sum_{r \in \{100, 500, 1000\}} \rho_r \mathbf{1}\left[\frac{z_j}{r} \in \mathbf{N}\right] + \eta_j \end{aligned} \tag{1.1}$$

where j refers to a certain bin (i.e. $(4100, 4200]$). c_j and z_j capture the density (i.e. number of tax returns) and the upper bound of bin j (i.e. $z_j = 320$). The equation estimates a polynomial approximation of the distribution of tax returns across turnover bins. β_j denote the coefficients of the polynomial vector. The coefficients ξ_n and ρ_r account for the influence of round numbers on both the liability as well as the turnover level. Turnover (i.e. the tax base) is denoted by k and the tax schedule is denoted by $T()$. The two indicator functions represent dummies for whether a turnover amount within a given bin coincides with a tax liability which is divisible by 10 and whether an upper bound of the bin itself coincides with a round number, respectively. z_L and z_U denote the lower and upper bound

[7] Thorough theoretical treatments of bunching at notches are available in Kleven and Waseem (2013), Bachas and Soto (2021) or Slemrod (2013)

Figure 1.1: Bunching concept

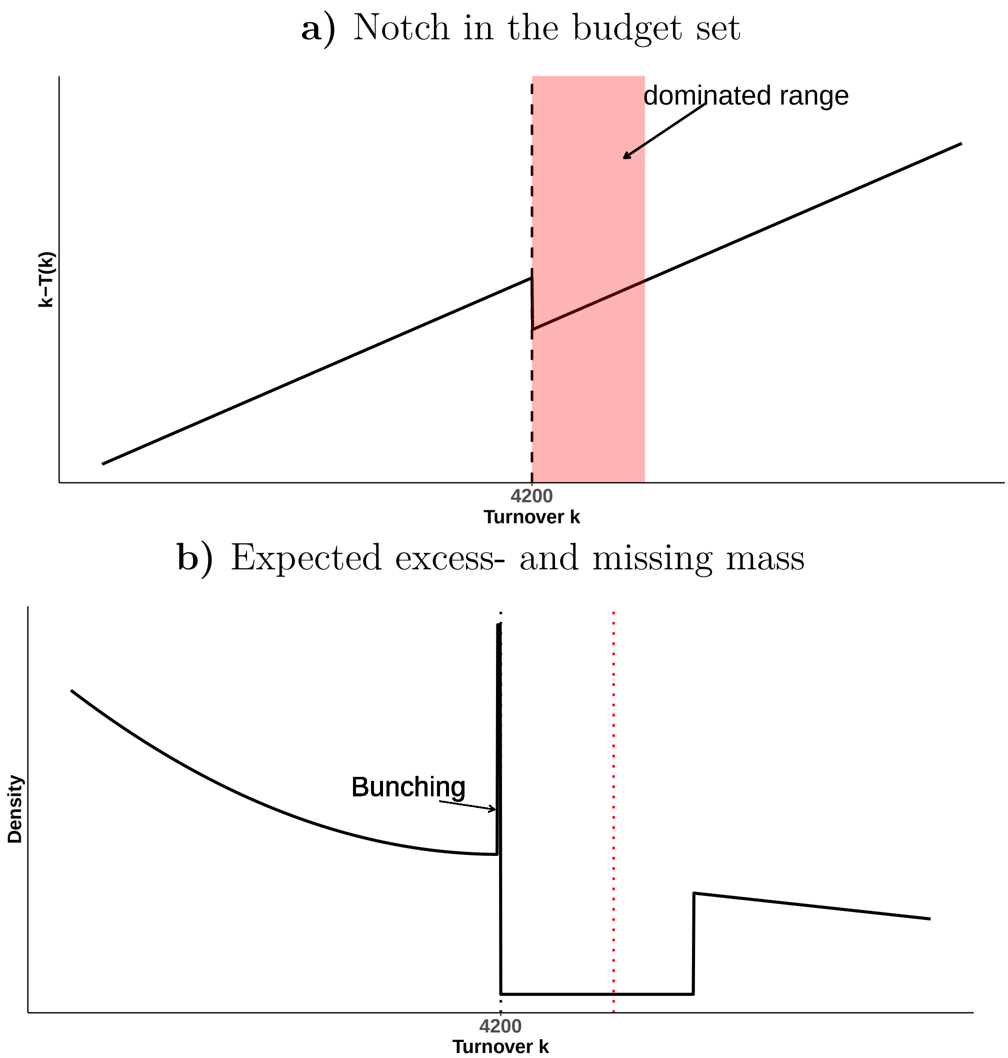

Notes: This figure illustrates the concept of bunching at notches. Panel (a) plots the budget set (i.e. net of tax turnover) on the vertical axis and turnover on the horizontal axis. The budget set is linear before the threshold, then drops discretely at the threshold and continues linearly from there on constituting. This *notch* establishes a dominated range (colored in red) above the threshold where net-of-tax turnover can be increased by either lowering or increasing turnover. Based on this tax schedule, panel (b) plots the expected distribution of firms around the threshold. One expects an accumulation i.e. *bunching* of firms just below the threshold and zero mass above the threshold within the dominated range.

of the affected region, respectively. These parameters are chosen in an iterative procedure. We start with $z_L = z_U$ and increase z_U until the area above the counterfactual (excess mass in the data) equals the are between data and counterfactual (missing mass). The counterfactual distribution is retrieved from estimating Eq. (1.1) and then fitting $\hat{c}_j$ while excluding $\hat{\gamma}_i$ (i.e. the effect of being in the affected region). The counterfactual in this case therefore estimates the density of tax returns in the hypothetical *absence of a notch*. The extent of bunching below the threshold is finally estimated by

$$B_{below} = \frac{\sum_{j=z_L}^{\bar{z}} c_j}{\frac{1}{|j \in [z_L, \bar{z}]|} \sum_{j=z_L}^{\bar{z}} \hat{c}_j} \tag{1.2}$$

To illustrate that the canonical model (and with it its predictions) might not be applicable, we estimate bunching responses also *above* the threshold. This is done by comparing the counterfactual distribution as estimated by Eq. (1.1) in the first bin above the threshold with the observed data mass in the same bin.

$$B_{above} = \frac{c_{\bar{z}+100}}{\hat{c}_{\bar{z}+100}} \tag{1.3}$$

If the canonical frictionless model is applicable, then $B_{above} = 0$. If there are optimization frictions B_{above} may be larger than 0 but less than 1 (Kleven and Waseem, 2013). We calculate standard errors on both bunching numbers as well as z_U can by bootstrapping the residuals in Eq. (1.1).

The second kind of bunching revolves around *round number tax liabilities*. As is well documented in the literature, tax data, especially from low-and middle income countries often features bunching at round numbers. The round number focus has thus far been on the tax base level and is usually attributed to poor record-keeping by taxpayers (Carrillo, Donaldson, Pomeranz, and Singhal, 2022; Kleven and Waseem, 2013). In this study, we estimate bunching at round number tax liabilities instead. With a tax rate of 3%, the Zambian setting allows us to distinguish between bunching at round liabilities and round turnover amounts. To see this, one can consider a firm desiring to pay (for some reason) 100. The amount of turnover it needs to report is $\frac{100}{0.03} = 3333.33$. Clearly, neither lazy reporting nor firms actually having such a turnover and declaring it truthfully could explain an accumulation of observations at such odd amounts.[8] To get a sense of whether firms are in fact focusing on round liability points we estimate bunching at such odd turnover amounts. We do so by binning the distribution of tax returns into bins of width 10 ZMK and estimating the following equation:

[8] If a firm, however, desires to pay 30 it would need to report then $\frac{30}{0.03} = 1000$. In this case, one could not know from the data whether the firm reported 1000 because of lazy reporting or because 30 was the desired payment amount.

$$c_j = \sum_{i=0}^{p} \beta_j (z_j)^i + \sum_{n \in \mathbf{N}} \xi_n \mathbf{1}\left[\exists k \in j | \frac{T(k)}{10} = n\right] + \sum_{r \in \{100,500,1000\}} \rho_r \mathbf{1}\left[\frac{z_j}{r} \in \mathbf{N}\right] + \eta_j \qquad (1.4)$$

where the variables are defined as in Eq. (1.1).
We estimate the extent of bunching at round number tax liabilities by first fitting $\hat{c}_j$ as estimated by Eq. (1.4) while excluding the $\hat{\xi}_n$, but including $\hat{\rho}_r$. We then compare this fitted density $\hat{c}_j$ to the observed density c_j to calculate a bunching coefficient for a round number liability r in bin j:

$$B_r = \frac{c_j}{\hat{c}_j}. \qquad (1.5)$$

As above one can calculate standard errors to these estimates by bootstrapping the residuals from Eq. (1.4).

1.4 Empirical facts

In this section, we apply our empirical methodology to the administrative tax data and establish two stylized facts, which are at odds with predictions from standard models of tax compliance. Furthermore, we present evidence from our own survey, describing how taxation works for small firms in Zambia.

1.4.1 Facts from the administrative data

Bunching below and *above* notches. We begin by estimating bunching responses to the notches in the tax schedule, which were introduced in 2017 (cf. Table 1.1). For our estimation we pool all tax returns that have been filed between January 2017 and December 2018 and thereby cover the whole time period the schedule was in place. It is worthwhile to mention that the Zambian schedule during that time is particularly suitable to estimate behavioral responses to taxes as the thresholds at 4200, 8300, 12500, 16500 and 20800 neither represent very salient numbers of turnover, nor are they associated with any other kinds of discontinuities in terms of benefits or taxes. Methodologically, we estimate bunching around thresholds as explained. In particular, we compare the distribution of tax returns around the threshold to the estimated counterfactual distribution. The latter is derived from estimating Eq. (1.1) and then fitting the density $\hat{c}_j$ while omitting the effects of the γ_i's i.e. omitting the *affected range* around the notch. As illustrated in

Figure 1.2: Bunching below and *above* the threshold

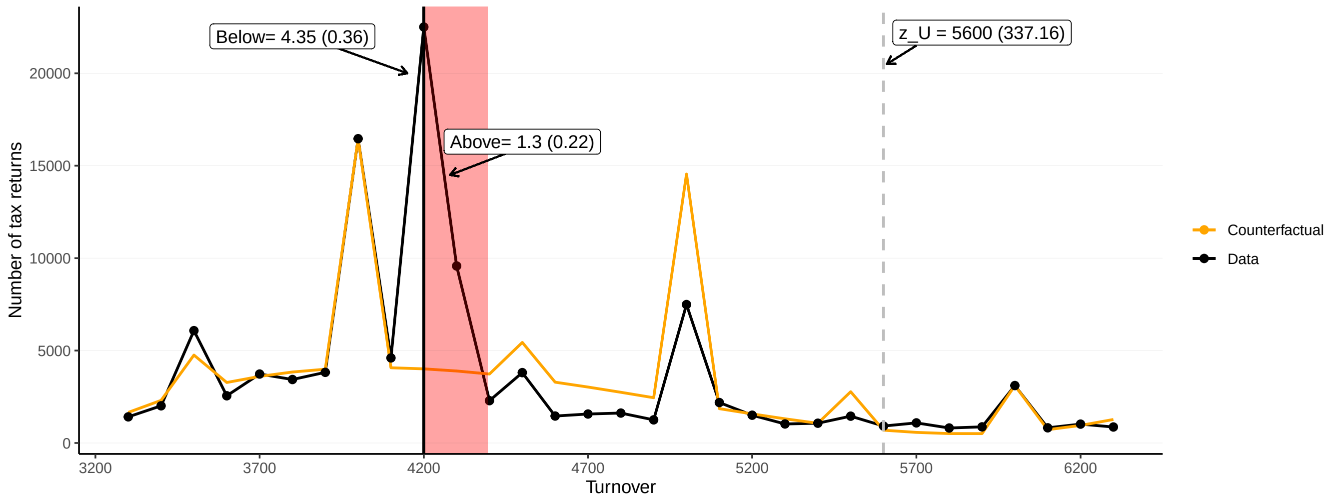

Notes: This figure plots the results of estimating bunching at the 4200 threshold. The black line depicts the empirical density. The yellow line depicts the counterfactual density as estimated by Eq. (1.1), accounting for round turnover amounts as well as round payment amounts. The black solid vertical line marks the threshold at which tax liability increases discretely i.e. the notch. The red area depicts the dominated range. The grey dashed vertical line depicts the upper bound of the omitted region z_U. Estimates of bunching below and above the threshold are derived from Eq. (1.2) & (1.3) and compare the counterfactual to the empirical density. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Years: 2017,2018.

Figure 1.1, standard models of tax compliance predict the empirical density to exceed the counterfactual density below the notch (bunching) and to be substantially below the counterfactual above the threshold.

Figure 1.2 provides graphical evidence of the bunching patterns. It contrasts the empirical density with the counterfactual density across turnover bins of width 100 ZMK. The *dominated range* is marked as red. Clearly, the counterfactual density approximates the empirical density well up to the 4100 bin. Furthermore, beyond the estimated upper bound of the affected region z_U, the counterfactual approximates the empirical density well. Focusing on the area just below the notch, in bin $(4100, 4200]$, there is a strong spike in the empirical density exceeding the counterfactual density by far. The extent of bunching is estimated by Eq. (1.2) which gives a bunching coefficient of $B_{below} = 4.35$, implying that there are more than 5 times more tax returns at the threshold than is predicted by the counterfactual. The bootstrapped standard error of 0.36 shows its statistical significance at the 99% level. Appendix 1.A shows that neither after nor before the notched schedule was in place, the distribution of tax returns featured bunching at these amounts.

When focusing on the region above the threshold, Figure 1.2 reveals another striking feature, which is at odds with the elaborated expectations. In particular, the empirical density also exceeds the counterfactual density in the first bin above the threshold. We quantify the extent of bunching *above* the threshold by following Eq. (1.3). The bunching estimate is $B_{above} = 1.3$. It is statistically significant at the 99% level with a bootstrapped standard error of 0.22. This bunching result is robust to using bins of different sizes as well as to relying on a completely non-parametric approach using the pre-reform distribution as the counterfactual (see Appendix 1.A). Recall, that the standard theory, as delineated in Figure 1.1 predicts no tax return to appear within the dominated range, because the tax schedule offers clear incentives to reduce turnover to below the notch. Instead of a sharp decline above the notch in the density above the notch, we find strong and significant bunching within this range.

To shed more light on this unexpected bunching response, Figure 1.3 zooms in on the area around the threshold and plots the density of tax returns in substantially smaller bins of size 1 ZMK instead of 100 ZMK. The figure clearly shows that the bunching response above the threshold is very sharp. While there is a strong accumulation of tax returns between 4200 and 4201, the number of tax returns drops for turnover above 4201. To contrast, this with the incentives the tax schedule is offering at the notch, we plot the average tax rate. This underlines the puzzling bunching pattern: firms are declaring exactly 1 ZMK more, even though the average tax jumps from about 1% at 4200 to more than 5% at 4201.

In addition to the threshold at 4200, we estimate bunching responses at the other thresholds. Table 1.2 reports B_{below} and B_{above} for the all thresholds. The pattern of significant bunching below as well as *above* holds across all thresholds. Additionally, also the concentration of tax returns which report turnover sharply above the threshold is consistent throughout. In Appendix 1.A, we provide the plots for these other thresholds and show that the extent of bunching patterns is very persistent over all 24 months in which the notched schedule was in place. The last row of Table 1.2 shows that a substantial fraction of tax returns within the first bins above thresholds, reported a turnover between the threshold and the threshold +1.

To the best of our knowledge, we are the first to document bunching of tax returns in strictly dominated regions. It is important to note, however, that the literature has highlighted the role of optimization frictions which may lead firms to end up in dominated regions above notches. The idea is that frictions render firms unable to bunch below a notch exactly and thus they end up above the notch. However, there is no reason to expect that frictions will lead to bunching at a specific point in the dominated region. This is also

Figure 1.3: Sharp bunching *above* the threshold

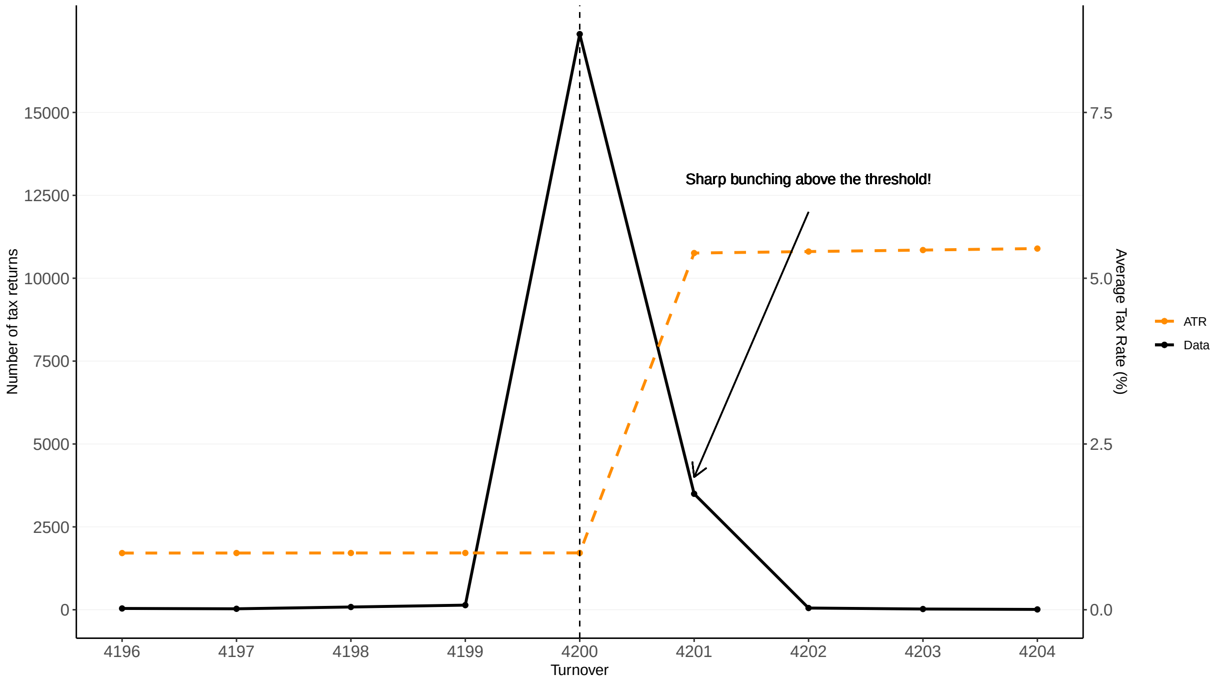

Notes: The black solid line plots the empirical distribution of tax returns around the 4200 threshold with bin size of ZMK 1. The dashed orange line plots the average tax rate that firms are facing given their turnover. Data Source: ZRA. Years: 2017,2018.

confirmed by existing empirical evidence, which shows that frictions lead to diffuse mass and not bunching above thresholds (e.g., Anagol, Davids, Lockwood, and Ramadorai, 2022; Kleven and Waseem, 2013). We thus rule out optimization frictions as the main driver of these bunching results. Further evidence and a discussion on the role of optimization frictions are provided in Section 1.6.2.

Bunching at round number tax liabilities. We also estimate bunching at *round number tax liabilities*. To do so we, we pool all tax returns filed between 2015 and 2016 and proceed as described in Section 1.3.2. During these two years, the tax schedule featured a flat rate of 3% for all businesses irrespective of their size. This allows us to cleanly distinguish between bunching at round numbers of the tax base and bunching at round number tax liabilities. Figure 1.4 contrasts the the density of tax returns as observed in the data with the counterfactual density over the binned turnover distribution. The counterfactual density resembles the fitted values $\hat{c}_j$ as estimated by Eq. (1.4) without the influence of the round number tax liabilities (i.e. omitting ξ_n) but accounting for potential bunching at round number turnover. Bunching at round number tax liabilities is present whenever the density in data exceeds the counterfactual density in a given bin.

Panels (a) to (c) of Figure 1.4 provide striking visual evidence of such bunching. The empirical density features several mass points at which it exceeds the counterfactual density by far. Further, the extent of the bunching increases with the salience of "roundness". For example, while there is a strong bunching of tax returns at 1333.33 ZMK, which implies a liability of 40 ZMK (and is not captured by the counterfactual), the extent of bunching at 1666.66 ZMK, which implies a liability of 50 ZMK is nearly four times as large. The same observation holds for other multiples of 50, as highlighted in the figures. In most cases, our estimated counterfactual captures bunching at round number turnovers well and the discrepancies between empirical and counterfactual densities are small. In a number of cases, however, round turnover amounts and round number tax liabilites coincide. This is true for e.g. all multiples of 1000. Overall, bunching at round number tax liabilities is strong across a wide range of turnover. In Appendix 1.B, we provide the corresponding bunching estimates for all liabilities which are divisible by 10, show that the patterns are persistent across time and further, that out of all tax returns filed between 2015 and 2016, 40% imply a tax liability divisible by 10.

To the best of our knowledge, we are the first to show clear evidence of bunching at round

Table 1.2: Bunching estimates

	Thresholds $\bar{k}$				
	4200	8300	12500	16500	20800
$\mathbf{B}_{below}$	**4.35**	**2.6**	**0.94**	**2.04**	**1.53**
	(0.36)	(0.33)	(0.21)	(0.43)	(0.13)
$\mathbf{B}_{above}$	**1.3**	**1.31**	**0.49**	**0.95**	**0.55**
	(0.22)	(0.21)	(0.17)	(0.42)	(0.08)
Sharp bunchers	**37%**	**37%**	**19 %**	**31 %**	**17%**

Notes: This table shows the estimated bunching coefficients for bunching above (Eq. (1.2)) and bunching *above* (Eq. (1.3)) for all notches. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. The last row depicts the share of sharp bunchers i.e. tax returns with turnover within the first bin above the threshold which report between the threshold and the threshold +1 i.e. bunch sharply above the threshold. Data source: ZRA. Years: 2017,2018.

Figure 1.4: Round Liability Bunching

a) Turnover 1000-4000

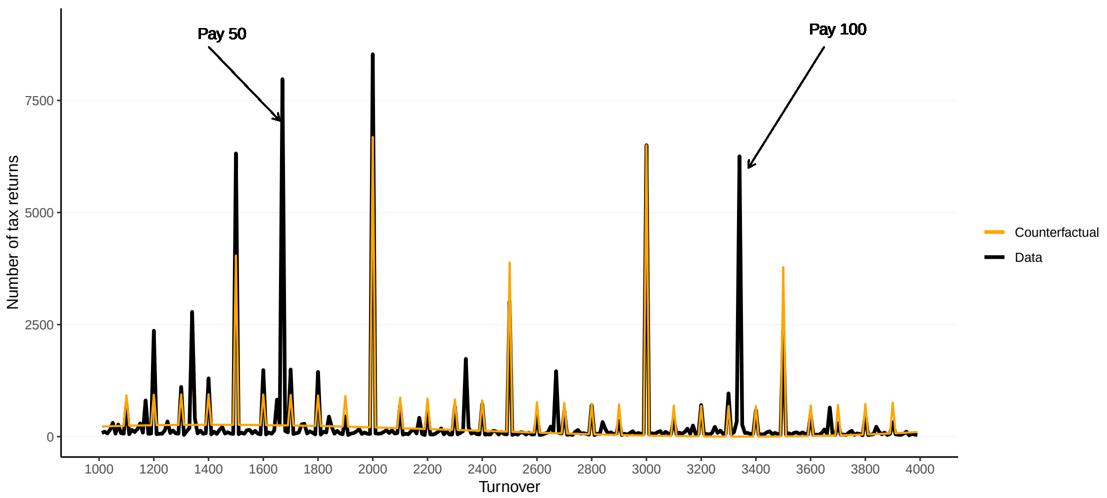

b) Turnover 4000-7000

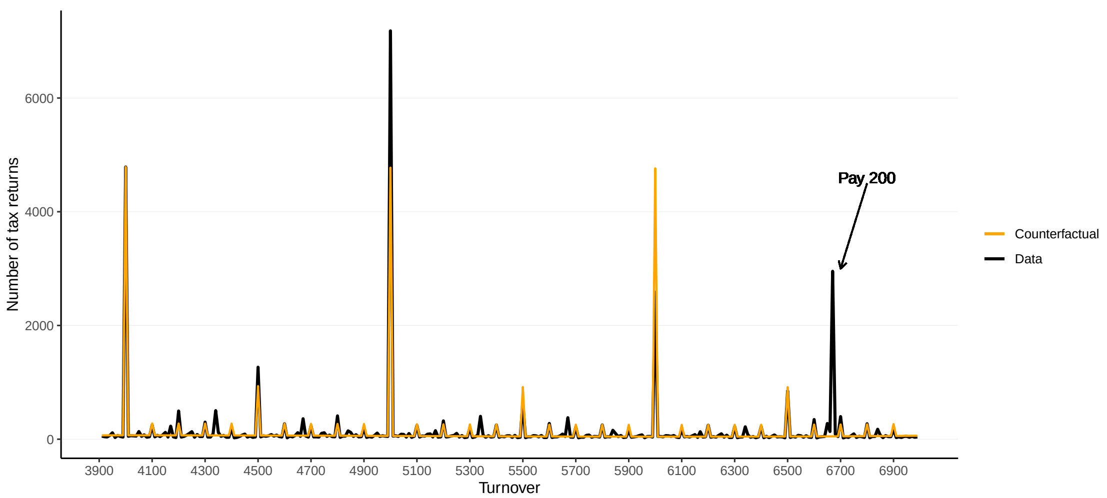

c) Turnover 7000-10000

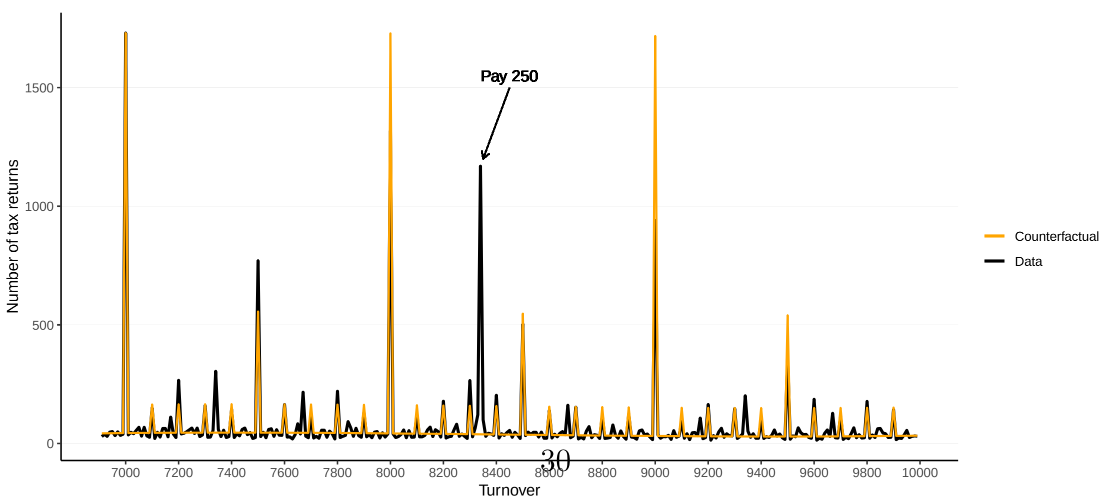

Notes: This figure plots the results of estimating bunching at round liability figures over the range 1000 to 10,000. The black line depicts the empirical density. The yellow line depicts the counterfactual density as estimated by Eq. (1.4). Bunching at payments of 50, 100, 200 and 250 are highlighted by arrows. Data source: ZRA. Years: 2015,2016.

number tax liabilities. This type of bunching is yet undocumented in the literature, but highlights that the payment amount itself rather than the declared tax base is the salient amount when tax returns are filed.

Bunching at round liabilities predicts bunching above thresholds. In a next step, we show that the behaviors underlying these two empirical facts are linked. We start by identifying firms that ever bunch *above* a threshold in the years 2017-2018. Exploiting the panel structure of the data, we then regress an indicator for whether a firm has ever bunched *above* on the share of the firm's tax returns prior to 2017 that were rounded at the liability level, i.e. imply a liability divisible by 10. To clearly distinguish from rounding at the turnover level, we do not include returns with a turnover divisible 1000, which implies both round liability and round turnover.

We include control variables across multiple dimensions.[9] First, we include several other pre-reform characteristics. Analogous to the variable capturing rounding at the liability level, we calculate the share of tax returns with round turnover amounts prior to the reform. We further calculate the share of tax returns with turnover that is exactly equal to the firm's previous month's tax return. This variable captures the phenomenon that firms in lower income countries may rely on past tax returns as a heuristic and simply file the same amount again (cf. Tourek (2022)). Finally, we account for the extent of pre-reform compliance (i.e. how often a tax return has been filed).

Second, we make use of the audit data to control for whether a firm has ever been audited and whether a payment of penalties or owed tax had to be made. We include sector- and tax office fixed effects for all regressions.

Table 1.3 presents the results. The different columns correspond to 3 different specifications, gradually including more controls. In all specifications, the strongest and only statistically significant predictor for whether a firm ever bunches above is exerting round tax liabilities. Overall, the results imply that a 1 percentage point (pp) increase in the share of tax returns with a round liability is associated with a 0.3 pp larger probability of bunching above a threshold. This implies an increase of 16% relative to the baseline. This effect is stable and significant at the 99% level throughout all three specifications. Further insights can be gained by inspecting the effects of the other regressors. First, we observe no significant effect of turnover rounding, speaking against the explanation that the observed phenomenon is a mere issue of bad record keeping. Second, targeting past payments is not

[9] As shown in Appendix 1.A, neither any specific sector nor any specific geographic location are significant predictors. Regressions on either tax-office or –sector fixed effects deliver a R^2 of about 0.001.

Table 1.3: Predictors for bunching *above*

	Dependent variable:		
	Ever bunched above (0/1)		
Firm characteristics	(1)	(2)	(3)
Liability rounding	**0.165*****	**0.168*****	**0.167*****
	(0.011)	(0.012)	(0.012)
[Liability × Turnover] rounding	**0.147*****	**0.152*****	**0.151*****
	(0.011)	(0.012)	(0.012)
Turnover rounding	**−0.006**	**−0.003**	**−0.004**
	(0.011)	(0.011)	(0.012)
Targeting past payments		**−0.014**	**−0.013**
		(0.016)	(0.016)
Audited (0/1)			**0.015**
			(0.056)
[Audited × Penalty] (0/1)			**0.097**
			(0.064)
Pre-reform compliance			**0.00003**
			(0.0003)
# Firms	18,516	18,516	18,516
Taxoffice FE	✓	✓	✓
Sector FE	✓	✓	✓
Baseline mean	**0.019**	**0.019**	**0.019**
R^2	0.023	0.023	0.023

Notes: This table shows the estimated correlations between bunching *above* thresholds and firms' behavioral characteristics prior to the notched schedule (Years 2015, 2016). 'Liability rounding' denotes the share of a firm's tax returns which are rounded at the liability level (with a round liability but odd turnover, e.g. 3333.33). 'Turnover rounding' denotes the share of tax returns rounded at the turnover level. 'Targeting past payments' is the share of tax returns with turnover that is exactly equal to the firm's previous month's tax return.'Audited' is an indicator for whether a firm has been audited. 'Penalty' is an indicator for whether a firm had to pay a penalty due to tax reasons. 'Pre-reform compliance' measures how often the firm has filed taxes prior to the reform. Data source: ZRA. Years: 2015-2018.

a significant predictor of bunching above suggesting that the heuristics on past payments plays a minor role. Third, the experience of being audited itself does not have a large influence on the outcome variable. Though, we do see a sizable effect of whether a firm has had penalty payments from tax audits, this effect is insignificant[10] and importantly its inclusion in the regression does not alter the effect of the strongest predictor: liability rounding.

1.4.2 Results from a firm survey

The empirical facts established from the administrative data show a stark contrast to the behaviour of firms under standard assumptions. However, as is typical for administrative data, it contains only limited information on firms' characteristics and no information on tax compliance behavior. Thus, we complement our analysis with findings from our own survey, which we described in Section 1.3.1.

Sample. Overall, 517 firms were surveyed in Lusaka in November and December 2023. Firm owners were quasi-randomly selected and approached in-person by the surveyors. The quasi-random selection followed a "snow-ball procedure". Once consent was given, the interview of approximately 30 questions started. A size restriction on a firm's turnover was put in place to ensure capturing only small firms. Figure 1.5 compares the distribution of sectors across firms in the survey with the administrative records. As the share of firms within most sectors aligns well, the survey can be considered representative for the firm population.

Firm characteristics. Panel A of Table 1.4 provides information on general firm characteristics. Half of all firms had not more than 2 employees (25% had no employees) and on average monthly turnover was ZMK 24,000. In 35% of cases, the owner was female and the most common sector was wholesale & trading. The results show that the vast majority of firms keeps detailed records of sales (i.e. turnover) and more than half of those keep track of those in a more proficient way than solely on paper logs. Interestingly, an even larger share states to keep detailed records of costs. As costs are irrelevant for turnover tax purposes but relevant for disposable income, this suggests that firms care more about what is paid in taxes ultimately rather than what the correct turnover is. It also suggests that firms are not generally lacking the resources to keep sales records. Further, tax literacy appears to be widespread. 96% state the correct turnover tax rate when asked and 80% are educated to at

[10] Audits have very low explanatory power, even when taken at face-value. The R^2 when only keeping audit variables as regressors is 0.003.

Figure 1.5: Industry compositions of firm samples; by data source

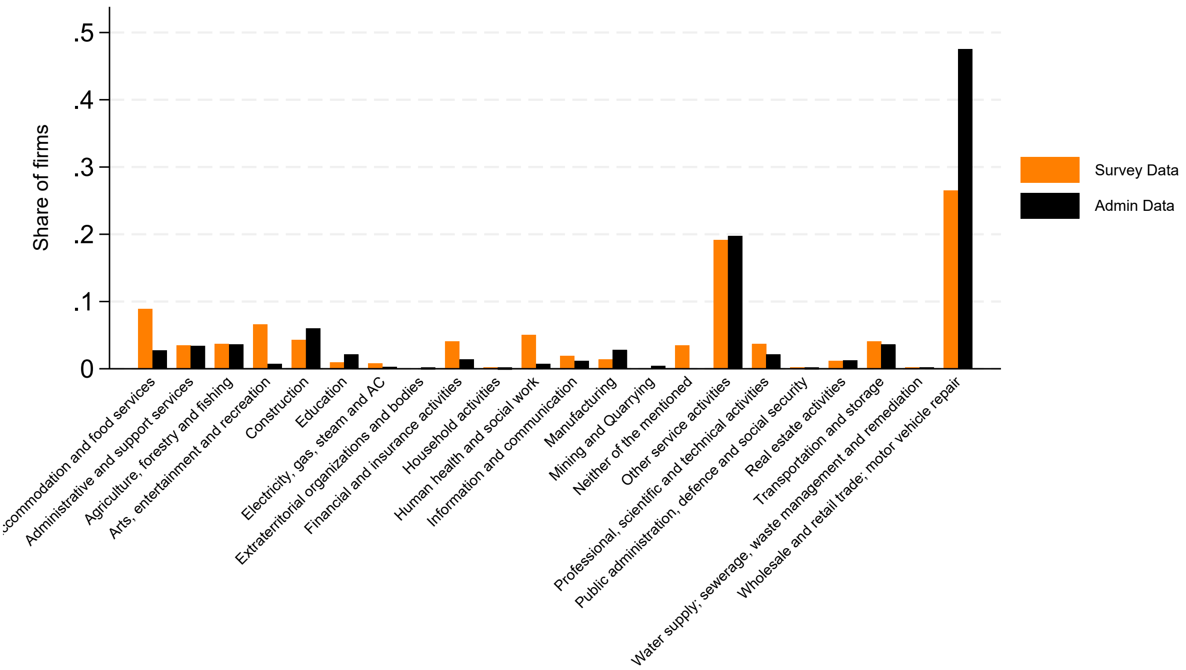

Notes: This figure plots the distribution of firms across industries in both the survey data as well as the administrative data. Data: ZRA (administrative data) and own collected data (survey data). Years: 2015-2021 (administrative data) and 2023 (survey data).

least a college degree. Finally, as opposed to a key assumption in standard models, 93% of surveyed firms state not to incorporate their tax liability when making business decisions. This implies that the marginal tax rate is irrelevant for the firm's optimization problem and thus, doing business and deciding on the reported tax base pose two separate problems.

Tax behaviour. Panel B of Table 1.4 concentrates on tax behavior. The most striking pattern which emerges is that discussions between officials and taxpayers about the tax payment is commonplace. Focusing on firms that were active during the period in which the notched schedule was in place (2017, 2018), 55% of respondents said they discuss the payment they are going to make with an official *before* filing their tax return or that such discussions are common. Overall, the share of firms stating to discuss payments with tax officials is 42%. Out of these responses, more than a third explicitly stated that during such discussions taxpayer and tax official "find an agreement" on what should be paid.

Table 1.4: Characteristics and tax behavior of surveyed firms

	Mean (1)	N (2)
Panel A: General		
Female (owner)	0.35	517
College or higher (owner)	0.8	517
# employees (median)	2	517
monthly turnover in (ZMK, average)	24,000	517
Firm does business with government	0.25	517
Whether firm keeps detailed records of sales	0.89	517
- only on paper	0.39	451
Whether firm keeps records on expenditures	0.95	517
Whether firm knows turnover tax rate	0.96	515
Whether firm thinks about tax liability when making business decisions	0.07	517
Panel B: Taxes		
Whether firms (including self) discuss payments with tax officials before filing		
All firms	0.42	517
- "to agree on payments"	0.36	216
- "to clarify tax liability"	0.57	216
Firms active during notched schedule (2017, 2018)	0.55	193
Perceived probability of being audited for tax reasons	0.19	515
	Coef.	p-value
- correlation with firm discussing payments	-0.05	0.01
Perceived percentage of firms bribing tax collectors	0.15	257

Notes: This table provides results from the survey conducted on firms under turnover tax in Lusaka between November and December 2023. Source: survey data.

In other words, firms are bargaining over tax payments with tax officials. As we show formally in Section 1.5, the prevalence of bargaining over tax payments can rationalize the seemingly irrational empirical facts in the administrative data. Interestingly, there is a strong and significant negative correlation between discussing tax payments and the perceived probability of being audited. This could hint at the fact that for the tax authority, bargaining over taxes before returns are being filed may be a substitute to formal audits after returns have been filed. Regarding corruption, we find that 15% of firms are deemed to pay bribes to tax collectors. This is much lower than the share of firms discussing payments with officials, suggesting that discussions do not serve a mere corruption purpose.

1.5 A model of tax bargaining

In this section, we show that the prevalence of bargaining rationalizes the observed bunching patterns. We begin by verbally delineating our main argument before formally introducing a theoretical model. Alternative explanations, which we rule out, are discussed in depth in Section 1.6.

1.5.1 Main idea

In many lower income countries, including Zambia, discussions with tax officials play a crucial role for tax compliance of small firms (cf. Section 1.2). Our survey has elicited that these discussions often serve the purpose of discussing intended tax payments before returns are filed and thus can be interpreted as bargaining situations. Against this background, it is clear that, what is observed as turnover in the tax data has a very different interpretation. In particular, rather than representing true turnover as a result of economic activity, tax returns indicate a turnover that results in a payment which is *agreed upon* by taxpayers and tax officials. We argue that following this interpretation is important to rationalize the observed patterns.

To see the argument, it is useful to consider bunching at round number tax liabilities first. This bunching pattern aligns well with the idea that tax returns represent the outcomes of bargaining. First of all, it highlights the focus on the actual payment as the salient amount when filing taxes. Further, a large literature has highlighted that round numbers serve as focal points in bargaining situations (e.g., Albers and Albers, 1983; Janssen, 2001, 2006; Pope, Pope, and Sydnor, 2015; Schelling, 1960). Thus, the strong accumulation of tax returns at turnover amounts which imply round payments may indicate that these are bargained and agreed upon payments. For example, when bargaining, it is more likely that two parties would settle on a payment of e.g. 100 ZMK instead of 99 ZMK.

Against this background, we turn to rationalize the fact that firms bunch above notches – a strictly dominated dominated choice through the lenses of standard models. The analysis so far suggests that firms and tax officials bargain over tax payments and the declared turnover represents the *inverted* tax schedule to arrive at the agreed upon payment. Under a flat tax schedule any such payment can be reached. When the tax schedule features notches, however, these notches effectively introduce regions of payment that cannot be reached anymore. To see this, one can consider the example of the 4200 threshold in the Zambian case (cf. Table 1.1). When declaring 4200 ZMK as turnover, the payment will be

36 ZMK. But when declaring a bit more, e.g. 4201 ZMK, the payment will be 225 ZMK. Thus, it is not possible to implement a payment between 36 and 225. When firms and tax officials are bargaining and would have settled on a payment within this interval, the agreed upon payment will end up either below or above that interval. In the tax data, this will create bunching on both sides of the notch. In what follows, we provide a theoretical framework to formalize this idea and offer a microfoundation for bargaining.

1.5.2 Theoretical model

To study how and when bargaining over taxes between taxpayers- and collectors evolves, we begin by establishing equilibrium tax payments in a non-cooperative setting (without bargaining). Then, we proceed to characterize pareto-improving payments in a cooperative setting (with bargaining). We show how the bargained payments relate to true tax liability and how the scope for bargaining vanishes along the path of economic development. Finally, we demonstrate how bargaining rationalizes the empirical facts described in Section 1.4.1.

Environment and agents. We consider a taxpayer with turnover z. The taxpayer's disposable income is given by $y = z - T - \pi$, where T denotes the tax payment and π any potential penalties. We will refer to the firm owner as the taxpayer in the following. Her preferences over consumption are described by $U_F = v(y)$ with $v'() > 0$ and $v''() < 0$. We further assume that the inverse of both $v()$ as well as $v'()$ exist. Given turnover z, the taxpayer needs to decide on the tax payment T. We assume that the firm will never pay more than it legally owes and thus $T \leq T(z)$. If T deviates from the the true liability $T(z)$, such that $T < T(z)$, the taxpayer risks, that, when being audited, the tax authority notices the discrepancy. In that case, she has to pay a fine (then $\pi > 0$). We denote the joint probability of being audited and caught by p. The fine that has to be paid if caught is proportional to the evaded amount by factor ξ (cf. Allingham and Sandmo (1972)). The taxpayer's maximization problem thus reads:

$$\max_T \mathbf{E}[U_F] = (1-p)v(z-T) + pv(z-T-(T(z)-T)(1+\xi)) \tag{1.6}$$

Denoting the solution to (1.6) as T^*, we can write the expected tax payments and fines in the *non-cooperative* setting as

$$T^{nc} \equiv (1-p)T^* + p(T^* + (T(z) - T^*)(1+\xi)). \tag{1.7}$$

On the receiving end of the tax payment is the tax authority. We assume that it has linear utility U_G in tax payments and fines and costs of engaging in audits. The latter is described as $c(p)$ with $c'() > 0$ and $c''() > 0$, accomodating the notion that increasing the probability of detecting tax evasion is increasingly costly. Let κ be the elasticity of costs with respect to the implemented audit and detection probability p: $\kappa = \frac{\partial c(p)}{\partial p}\frac{p}{c(p)}$. We assume that $\kappa \geq 1$.[11] The decision problem of the tax authority reads:

$$\max_p \mathbf{E}[U_G] = T^{nc}(p) - c(p). \tag{1.8}$$

Figure 1.6 depicts a numerical example of both utility functions graphically as a function of the tax payment. The expected utilities in the non-cooperative equilibrium are denoted by $E[U_F|T^{nc}]$ and $E[U_G|T^{nc}]$ respectively.

Pareto-improvements through bargaining. We now consider the option of taxpayer and –authority bargaining over tax payments. As shown in our survey results (cf. Section 1.4.2), such behaviour is a common feature of how tax payments are determined. We thus extend the model as follows. If both taxpayer and tax authority can agree on a tax payment in advance, then no audits will take place. Naturally, an agreement may only be reached if both parties are better off than without bargaining. If there is no agreement, the outside option realizes which is the non-cooperative equilibrium as described above. We therefore consider the set of payments which are pareto-improving relative to the non-cooperative equilibrium. Clearly, the taxpayer would agree on any payment T such that $v(z-T) \geq E[U_F|T^{nc}]$ i.e. she is as least as well off as in the non-cooperative setting. We define the maximum amount of certain payment the taxpayer would be willing to make as T_F, derived from the following inequality:

$$v(z-T) \geq E[U_F|T^{nc}] \iff T \leq z - v^{-1}(E[U_F|T^{nc}]) \equiv \mathbf{T_F}. \tag{1.9}$$

For the tax authority, agreeing on a certain tax payment saves the audit costs $c(p)$. On the other hand, costs related to bargaining with the taxpayer may occur. We denote the

[11] Technically, this implies that increasing the audit probability by 1%, increases the cost by at least 1%.

Figure 1.6: Utility functions and outside options in the bargaining model

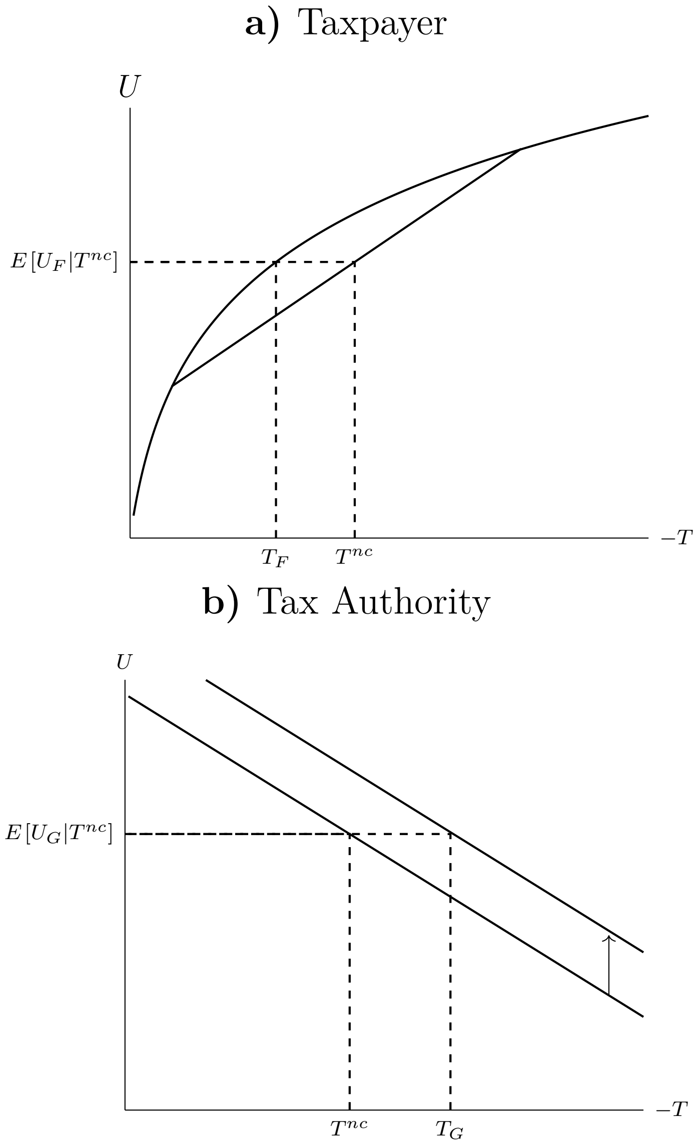

Notes: This figure illustrates the utilities of taxpayer and tax authority as a function of the tax payment T. T^{nc} denotes the equilibrium payment in the setting without bargaining. T_F and T_G depict the outside option payments and thus determine the set of pareto-improving payments.

costs of audits net of bargaining costs as $\tilde{c}(p)$.[12] Note that if $\tilde{c}(p) > 0$, bargaining shifts the utility function of the tax authority upwards as depicted in panel (b) of Figure 1.6. Thus, the minimum tax payment the tax authority would accept is given by T_G which can be

[12]For simplicity, we set bargaining costs for firms to 0. However, this assumption is not essential to our results.

derived from the following inequality:

$$T \geq T^{nc} - \tilde{c}(p) \equiv \mathbf{T_G}. \tag{1.10}$$

Both T_F and T_G are depicted in Figure 1.6.

Equilibrium. An equilibrium with bargaining can only exist if taxpayer and tax authority are at least as well off as in the non-cooperative setting. Thus it needs to hold that:

$$T_F \geq T_G. \tag{1.11}$$

Intuitively, this means that bargaining is only beneficial if the maximum amount the taxpayer is willing to pay is larger than the minimum amount the tax authority is willing to accept. The bargaining set is therefore given by the interval (T_G, T_F). In Proposition 1, we formalize how the true tax liability $T(z)$ relates to the set of possible bargaining outcomes.

Proposition 1. *The true tax liability $T(z)$ is outside the bargaining set. In particular $T(z) \geq T_F$.*
Proof: See Appendix 1.F.

Two insights can be derived from Proposition 1. First, taxpayers engaging in bargaining, can reduce their de-facto tax payment without inducing the risk of penalties through audits. Second, the payment resulting from bargaining is detached from a taxpayer's true liability, which is consistent with the notion that bargained payments often simply end up on round figures.

We turn to ask under which circumstances the bargaining set is non-empty and scope for bargaining exists. In particular, why we would not expect such bargaining to happen in more developed countries. One key feature of our model that may change along the path of development is the costliness of doing successful audits $c(p)$. In fact, the turnover tax for small firms outside the VAT net is difficult to assess and therefore to audit. This is because most transactions are made in cash and leave no paper trail. There is therefore virtually no third party reporting in most cases. In more developed countries, third-party reporting and withholding is more common and successfully assessing and auditing a large share of a firm's income is easy for tax authorities (Kleven, Kreiner, and Saez, 2016). In our model, the costliness of doing successful audits can be characterized as the elasticity of the cost function with respect to audits. A lower κ is therefore associated with larger state capacity. The role of κ for the bargaining outcome is summarized in Proposition 2.

Proposition 2. *With $\kappa \to 1$, the bargaining set collapses and $T_F = T_G = T(z)$.*
Proof: See Appendix 1.F.

The intuition behind Proposition 2 is that as countries develop and their state capacity grows, they are able to enforce taxes more efficiently. In its limit, our model predicts that the set of possible bargaining outcomes breaks down.

Dividing the surplus. If condition (1.11) holds, then there exists a surplus from bargaining which is given by $T_F - T_G$. We assume that this surplus is divided among taxpayer and tax authority via Nash-bargaining.[13] Clearly, the most desirable outcome for the taxpayer would be to pay only T_G, i.e. the minimum amount the tax authority would accept. The most desirable outcome for the tax authority is the extract the maximum amount the taxpayer is willing to pay, which is T_F. Denoting the bargaining power (this can be thought of e.g. as enforcement capacity) of the tax authority as α, the following Nash-product is maximized:

$$\max_T (T_F - T)^{1-\alpha}(T - T_G)^{\alpha}. \tag{1.12}$$

Rewriting the first-order condition of Eq. (1.12), we arrive at the *cooperative* solution

$$T^c = \alpha T_F + (1-\alpha)T_G.$$

Bunching in the bargaining model. We now turn to illustrate how the bargaining framework can rationalize the empirical bunching patterns. First, we consider the linear tax rate setting: Once the firm and the tax authority have agreed on a payment $x = T^c$, the firm needs to report $T^{-1}_{schedule}(x)$. As the tax rate is at 3%, this is simply $\frac{x}{0.03}$. A linear tax schedule is bijective and therefore, every payment that was agreed upon could be reached. The behavorial literature has shown that bargaining outcomes are strongly concentrated on round figures, supporting the notion that bunching at round tax payments stems from bargaining (Albers and Albers, 1983; Pope, Pope, and Sydnor, 2015; Schelling, 1960).

Next, we consider the notched tax schedule (cf. Table 1.1). As notches imply jumps in

[13] Nash-bargaining offers a convenient way to ensure uniqueness of the bargaining outcome. Other bargaining models incorporating e.g. focal points (Janssen, 2001) may be considered in future work.

the average tax rates, they effectively introduce regions of tax liability that can not be implemented. We illustrate this at the example of the 4200 threshold:

$$T_{schedule}(4200) = 36; T_{schedule}(4200.01) = 225$$
$$\Longleftrightarrow$$
$$T^{-1}_{schedule}(x) = \emptyset, x \in (36, 225)$$

Now, for all firms with $T^c \in (36, 225)$ the bargaining problem becomes a case distinction of whether the Nash product in Eq. (1.12) is larger above or below the threshold.[14]
Thus, firms within this range report turnover either above or below the threshold. Tax returns will accumulate on both sides which generates the observed bunching. If we consider T_F and T_G as functions of true turnover z, the cutoff for bunching *above* is at $\tilde{z}$ such that

$$(T_F(\tilde{z}) - 36)^{(1-\alpha)}(36 - T_G(\tilde{z}))^{\alpha} = (T_F(\tilde{z}) - 225)^{(1-\alpha)}(225 - T_G(\tilde{z}))^{\alpha}.$$

The case of Nash-bargaining over tax payments in the presence of notch is depicted in Figure 1.7. In this graphical example, the notch leads to the bunching *above* the threshold. Finally, let true turnover z be continuously distributed with density f. The total bunching mass on both sides of cutoff is given by

$$\underbrace{\int_{z_B}^{\tilde{z}} f(z)dz}_{\textbf{Bunching below}} \quad + \quad \underbrace{\int_{\tilde{z}}^{z_A} f(z)dz}_{\textbf{Bunching above}} \tag{1.13}$$

where $\tilde{z}$ is the cutoff value from above, z_A is such that $T^c(z_A) = 225$ and z_B such that $T^c(z_B) = 36$. Eq. (1.13) therefore provides a mapping of the distribution of true turnover to the bunching we see in the data.

Insights. Our model shows that bargaining over taxes is pareto-improving over the non-cooperative standard setting under reasonable assumptions for lower income countries. This holds because firms are risk averse and/or the tax authority has higher costs of auditing ex-post than discussing payments with firms ex-ante. Note that each condition implies room for bargaining independently of each other. We demonstrated how the model ratio-

[14] Eq. (1.12) is strictly concave in T under the assumption that $T_F \geq T_G$. Therefore the maximization problem turns into a case distinction between the two options if the unconstrained optimum is within $(36, 225)$.

Figure 1.7: Nash-Product with notch

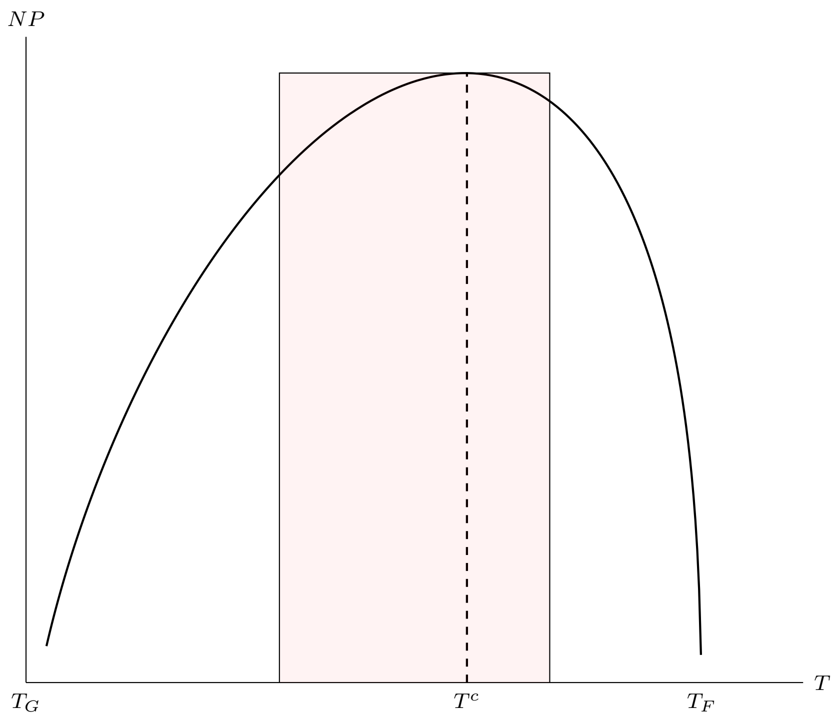

Notes: This figure illustrates the Nash-product (NP), which the tax authority and taxpayer maximize when bargaining over the payment (concave curve). T^c denotes the unconstrained solution. The red rectangle illustrates a region of payments which become unreachable due to a notch in the tax schedule. In the illustrated case, the NP is larger *above* the notch than below. Therefore, the declared turnover needs to be just above the threshold.

nalizes the two empirical facts, which are otherwise puzzling. First, bargained payments are detached from a firm's true liability, corroborating the notion that payments end up on round numbers. Second, notches restrict the feasible bargaining set, leading to payments bunching above the threshold.

1.5.3 Discussion

We will now discuss the role of corruption and the incentives of tax collectors in the process of bargaining.

The role of corruption. Interactions between taxpayers and tax officials are inherently connected to concerns about corruption. One might worry that a system, which allows for such interactions supports corruption, where firms can bribe tax officials and make lower tax payments. While we cannot rule out that such activities are taking place, we argue that corruption is unlikely to be the main driver. First, from our survey as well as the World Bank Enterprise Surveys, we see that meetings and discussions with tax officials are much more common than bribery of tax officials. The World Bank estimates that in Zambia, more than 80% of small firms regularly visit tax officials, but less than 4% bribe tax officials during such visits (cf. sections 1.2 and 1.4.2). This suggests that the discussions generally serve another purpose than exchanging bribes. Bargaining and eventually agreeing on a tax payment as outlined above is one such purpose, supported by the results of our own survey. One should also note, that bunching above a threshold is hardly consistent with corruption. If firms and tax officials would collude, they would always have the clear incentive to agree on a tax payment below the threshold and share the difference to the payment that would have been due above the threshold. Hindriks, Keen, and Muthoo (1999) provide a theoretical framework for corruption which formalizes this notion.

Incentives of tax officials. If corruption does not play a major role, one could ask, what the incentives of tax officials are besides that. First, it is to note that bribery is obviously a misconduct and tax officials could lose their (relatively well paid) job.[15] Second, the ZRA has a system of performance benefits in place which rewards tax offices with good collection performance. These benefits, however, can only be paid out to individual tax officials when the tax office reaches its prescribed revenue target. Thus, thanks to the incentive scheme, incentives of officials and tax authority are aligned. In a bargaining situation, this incentive could lead tax officials to push for firms to bunch above the threshold. This would further mean that once the revenue target is reached, one would observe less bunching above. We confirm this channel empirically and estimate a significant drop in the likelihood of firms bunching above once a tax office has hit its revenue target (cf. Appendix 1.D).

1.6 Alternative explanations

This section points out several alternative explanations for the empirical facts and shows that each of them are unlikely to be sufficient explanations.

[15]The median annual salary for tax collectors in Zambia is more than twice as much as the country's GDP per capita. Source: `https://worldsalaries.com/average-tax-officer-salary-in-zambia/`. Last accessed: June 21st, 2024.

1.6.1 Audits

One alternative explanation might be that even without bargaining, the threat of being audited might motivate firms to pay more and therefore bunch above the threshold. We provide four pieces of evidence to rebut this concern.

Correlational Analysis. As seen in Table 1.3, we do not find a significant relationship between a firm ever being audited and a firm ever bunching above. While there is indeed a sizable point estimate for whether a firm has ever paid a fine, it is imprecisely estimated. Note that even if we would take this estimate at face-value it does not explain a lot of the variation in the explanatory variable. This is because the number of firms that ever bunch above by far exceeds the number of firms that have ever paid a fine after being audited.

Empirical audit probabilities. If firms would strategically bunch above thresholds to avoid audits, then the probability of being audited should be substantially lower above than below the threshold for it to be worthwhile to make higher tax payments. In Appendix 1.C, we estimate empirical audit probabilites conditional on reported turnover and show that they do not differ substantially above vs. below a threshold.

Simulations from tax evasion model. We further address the concern that these empirical audit probabilities might not be known to taxpayers and differences might be perceived as larger. In Appendix 1.C, we simulate a model of tax evasion along the lines of Kleven, Knudsen, Kreiner, Pedersen, and Saez (2011) with endogenous audit probability for the case of the Zambian turnover tax. We demonstrate that even under extreme assumptions where the perceived probability of being audited jumps from 10% above to 30% below the threshold, only firms that evade at least 83% of their turnover would choose to stay above the threshold.[16]

Randomized survey experiment. Finally, we took an experimental approach to test whether perceived audit probabilities could drive firms to bunch above thresholds. At the end of our survey, we included a randomized experiment with the aim to test the audit channel and also a government contracting channel. The treatment consisted of providing

[16] This result is for the standard model which assumes quasi-linear utility (Bachas and Soto, 2021; Kleven, Knudsen, Kreiner, Pedersen, and Saez, 2011; Kleven and Waseem, 2013). Accounting for risk aversion, we conclude that only firms evading at least 58% or 50% of their turnover would choose to bunch above, when imposing CRRA utility and a risk-aversion parameter of 2 or 3, respectively.

information and thereby shifting beliefs of respondents.

The experimental setup was as follows. We randomly assigned each respondent to one of three groups: control group or one of two treatment groups. We thereby followed the standard methodology of information experiments. First, it was tested whether the threat of audits might motivate firms to bunch *above* the threshold instead of below. For example, if audit probability increases discretely when moving below the threshold, firms might trade-off higher tax payments with lower audit probability below the threshold. The first treatment group therefore received information on the amount of audits and money seized therein by the ZRA. Second, as 75% of respondents stated that securing a business contract with the government was very competitive, we investigated whether this could be driving firms to put themselves in a good position with the government and bunch *above* the threshold rather than below. The second treatment group therefore received information on the amount of business the government is doing with SMEs and a reminder that one needs a tax clearance certificate to engage in such contracts. The control group only received information about the total number of firms registered under turnover tax.

In order to measure a firm's propensity to bunch above the threshold, the survey followed a specific procedure. In the beginning of the survey, the average turnover of the firm was inquired. Later, the firm was asked which tax payment the respondent thinks is appropriate for a firm with the previously stated turnover. Then the randomized information treatment occurred. Once the information was given, respondents were confronted with a hypothetical notch: they faced a situation in which their stated appropriate payment is not feasible anymore, instead they only had the option to either pay 10% more or 10% less. Opting for the larger payment would then be associated with bunching above.[17]

We estimate treatment effects of of our two information treatments by running the following regression.

$$y_i = \alpha + \beta_1 treat_1 + \beta_2 treat_2 + \mathbf{X}_i + \gamma_e + \epsilon_i \tag{1.14}$$

where y_i denotes the response of respondent i after being treated, $\mathbf{X_i}$ are controls, γ_e are enumerator fixed effects and ϵ denotes the error term. $treat_1$ and $treat_2$ correspond to indicator for whether the respondent belongs to treatment group 1 or 2 respectively. We then test the null hypothesis $H_0 : \beta_1 = 0$ or $H_0 : \beta_2 = 0$ to investigate whether the treatments have a causal effect on the response y.

[17] As we are unable to link respondents' identities to actual tax records, we had to rely on this stated measure of the propensity to bunch above threshold. A more detailed description of the survey and examples are provided in Appendix 1.E.

Table 1.5 shows the treatment effects of both information messages as estimated from Eq. (1.14). The last two columns of the table show results for specification where the the average turnover is fixed at 10,000 and the appropriate payment is exogenously shifted to 500 (column (4)) such that both feasible payments (450 or 550) are legal under the contemporary tax rate of 4% as well as to 300 (column (3)) where both options would be illegal. This serves as a check that responses are not driven by such legal considerations.
The results are striking in the sense that we estimate no significant effect of any treatment throughout. While the baseline shows no significant effect, also adding controls for turnover does not change this result in column (2). Further, neither columns (3) or (4) show a statistically significant impact on of the treatment on propensity to bunch *above*, suggesting that the result in columns (1) and (2) is not driven by legal considerations. Note that our confidence intervals also exclude existing estimates on the effect of deterrence messages on tax compliance.[18] Overall, this null-result rejects the hypothesis that standard explanations such as the threat of being audited can rationalize the empirical facts.
Clearly, as in any survey experiment, the outcomes we measure are only stated and do not necessarily coincide with real actions. However, being unable to match survey respondents to their administrative tax records, we consider our approach to be second-best.

1.6.2 Frictions

We will now discuss the role of frictions which may hinder firms to respond to incentives optimally. This could lead firms that actually wanted to bunch below to bunch above a threshold, instead.

Optimization frictions. Potentially, taxpayers locating above the threshold, were actually aiming to have a turnover just below the threshold but simply failed to do so. This idea has been formalized in the literature already and in the following, we will consider two

[18] A comparison to existing estimates of the effect of deterrence messages (such as increasing the salience of audit probabilities) is not straightforward for two reasons. First, the measure of increasing salience is heterogeneous across studies. Second, most studies measure effects on taxes paid. We are interested in bunching above thresholds. To still offer a comparison, we can translate the outcome variable of *bunching above* into an equivalent increase of tax paid by 10%. A positive effect of e.g. 5.6% (the upper bound on our confidence interval in column 1, row 1) would thus imply an increase in taxes paid by 0.56%. This is substantially lower than existing estimates from lower income countries (e.g., Mascagni and Nell, 2022; Shimeles, Gurara, and Woldeyes, 2017) and also richer countries (e.g., Fellner, Sausgruber, and Traxler, 2013; Slemrod, Blumenthal, and Christian, 2001)

Table 1.5: Informational treatment effects on *bunching above*

	Dependent variable				
	Bunch above threshold Baseline Liability:				
		Stated		ZMK 300	ZMK 500
	(1)	(2)	(3)	(4)	(5)
Audit Treatment	**0.0261**	**0.0287**	**0.0263**	**0.0124**	**0.0198**
	(0.0299)	(0.0304)	(0.0305)	(0.0325)	(0.0446)
Contract Treatment	**0.0110**	**0.0158**	**0.0149**	**0.0140**	**0.00952**
	(0.0294)	(0.0301)	(0.0297)	(0.0312)	(0.0426)
# Firms	517	510	510	251	259
Size control		✓	✓	✓	✓
Enumerator FE			✓	✓	✓

Notes: This table shows the results from the randomized survey experiment. The coefficients in columns (1), (2), and (3) represent the estimated effect of the two information treatments on an indicator for whether survey respondents would choose a larger tax payment if their favored option was no longer feasible. We interpret this as the propensity to bunch above. Additionally, column (2) controls for turnover. Columns (4) and (5) represent the same coefficients only that the initial tax payment option was not chosen by the respondents but fixed at ZMK 300 or ZMK 500 respectively. Robust standard errors are in parentheses. Data source: survey data. Year: 2023.

such approaches and show that optimization frictions are unlikely to drive the observed bunching pattern.

Kleven and Waseem (2013) explain the presence of mass in the dominated region by optimization frictions. In their application, sharp bunching *below* the threshold suggests an extensive margin of frictions: either a taxpayer is able to manipulate turnover to lie below the threshold (in that case, exactly) or the taxpayer is unable to adjust at all. In the case of Zambia, the bunching *above* the threshold is also sharp. Bunching above therefore seems to be a strategic choice rather than driven by frictions to bunch below the threshold.

A more recent paper by Anagol, Davids, Lockwood, and Ramadorai (2022) explains excess mass in dominated regions by the probability distribution of opportunities around the threshold. The idea is that given the threshold as a target, taxpayers draw from a discrete set of opportunities around the target. This might result in some taxpayers "overshooting" the target and ending up above the notch. However, as in Kleven and Waseem (2013), this model also predicts a hole in the distribution above the notch, rendering it clearly inconsistent with the excess mass in this area. We conclude that neither type of optimization frictions discussed in the literature is able to explain the observed bunching pattern.

Record keeping. Another form of frictions could be a firm's low capacity to keep detailed records of turnover and therefore file correctly. This could also lead firms to rely on heuristics such as previous months' returns and file the same amount again. If unable to file same amount again, due to a notch, firms could deviate either upwards or downwards (Tourek, 2022). We argue that record keeping and this form of heuristics play a minor role in the observed bunching pattern for two main reasons. First, the share of *targeted past payments* is uncorrelated with the phenomenon of bunching above thresholds (cf. Table 1.3). Second, our survey shows that 89% firms are able to keep records of at least some form (cf. Table 1.4).

1.6.3 Mistakes

One could also consider the possibility that firms simply make mistakes. This could happen when firms are inattentive to the tax schedule or do not understand it correctly. Almunia, Hjort, Knebelmann, and Tian (2023) for example argue that in Uganda, a substantial share of VAT returns are wrongly filed because firms are confused.

We provide five pieces of evidence to show that such an explanation is insufficient to rationalize bunching above thresholds. First, we see the cross sectional bunching pattern for all periods from January 2017 to December 2018 (see Figures 1.A.7 and 1.A.8). One would assume that if firms were simply confused, there would be at least some learning over 24 filing periods, which we do not find. Second, we do not observe firms to directly 'correct' their turnover from above the threshold to below the threshold in the next filing period, which would be the natural reaction if bunching above would have constituted a mistake (see Figure 1.A.9). Third, for the sample of firms which are observed both above a threshold as well as below a threshold at different points in time, we document that 58% of these firms bunch above a threshold after they have already bunched below a threshold in a previous tax period. This is clearly inconsistent with the idea that bunching above a threshold simply constitutes a mistake. Fourth, in our survey, we elicited that most firms are aware of the correct tax rate, speaking against the channel of mere tax illiteracy. Finally, the focus on round number tax liabilities suggest that firms are aware of the implications of filing certain turnover for the tax liability.

1.7 Conclusion

This chapter shows that bargaining over tax payments is an important feature of tax collection in lower income countries. The empirical setting is Zambia but we argue that the factors driving this behaviour are similar in many other low and lower-middle income countries and thus, our results apply more generally.

We study firms subject to turnover taxation by analyzing administrative data on the universe of turnover tax filings in Zambia and establishing two novel empirical facts. First, we find strong and sharp bunching *above* tax schedule discontinuities which is a strictly dominated choice in standard models. Second, we find strong bunching at odd turnover amounts, which imply round number liabilities. These observations are at odds with predictions from standard models of tax compliance, but can be rationalized when interpreting tax payments as bargaining outcomes between taxpayers and tax officials.

We gather several pieces of evidence from data on tax audits, our own survey of more than 500 firms as well as a randomized survey experiment which reject competing explanations for the observed bunching patterns. Finally, we propose a simple theoretical framework of tax compliance, which explains how and when bargaining occurs and rationalizes both empirical facts. It shows that bargaining over taxes leads to pareto-improvements for both taxpayers and the state as long as state capacity is sufficiently low. As countries develop state capacity, bargaining over taxes becomes obsolete.

Overall, our results inform the debate of how the circumstances of lower income countries shape the way (tax-)administration works. We show how puzzling facts in the tax data can be rationalized when accounting for the presence of bargaining as a mode of tax compliance and –enforcement. While the fear might be that such a system generously invites corruption, we argue that bargaining over taxes benefits both taxpayers and tax authority. Hence, shutting it off has welfare implications, the quantification of which we consider an important avenue for future research. Rather, our study suggests that considering this mechanism yields important insights when formulating policy recommendations. In particular, the official tax schedule is less important than one might think. Reforms of the latter may therefore have little impact on revenue. Instead, other policies which could increase bargaining outcomes may be more promising in this regard. As such, lowering the VAT threshold and thereby exposing more firms to third-party reporting constitutes one example.

Appendix

1.A Anatomy of bunching *above*

This section complements section 1.4 by further investigating the phenomenon of bunching *above* the threshold. Figures 1.A.1 and 1.A.2 plot the bunching figures for the other thresholds. Figures 1.A.3 and 1.A.4 plots the distribution of tax returns before and after the notched schedule has been in place, respectively. Figures 1.A.5 and 1.A.6 plot bunching responses with different counterfactual distributions. Figures 1.A.7 and 1.A.8 show that the cross sectional distribution is consistent across months. Figure 1.A.9 shows that bunching firms do not immediately bunch below in the following periods. In Figure 1.A.11, we show that most firms only bunch *above* once. Tables 1.A.1 and 1.A.2 show the heterogeneity of bunching above by sector and tax office, respectively.

Figure 1.A.1: Bunching around thresholds 8300 and 12500

a) Bunching at 8300

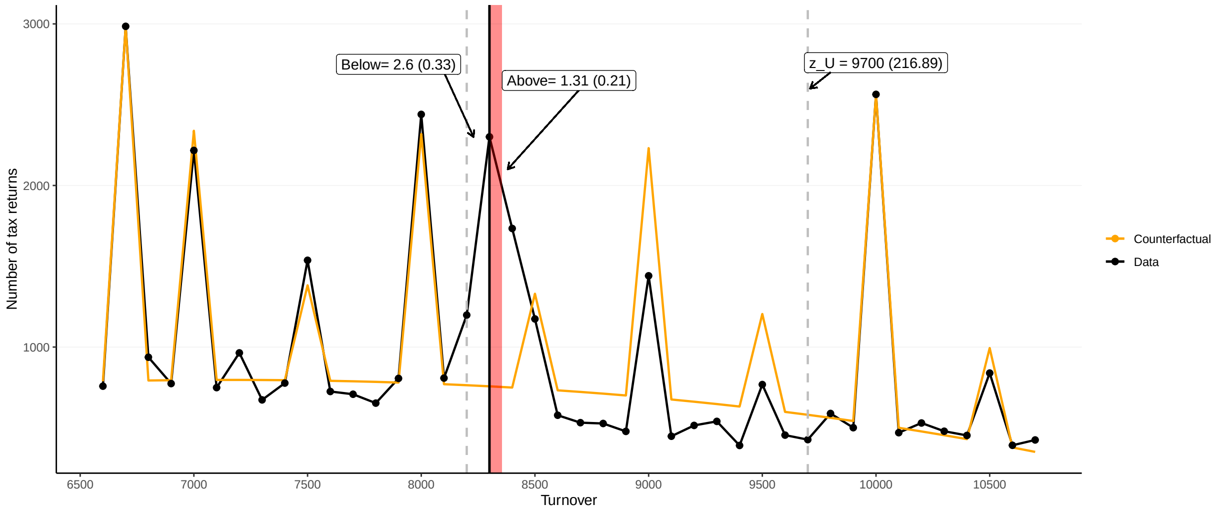

b) Bunching at 12500

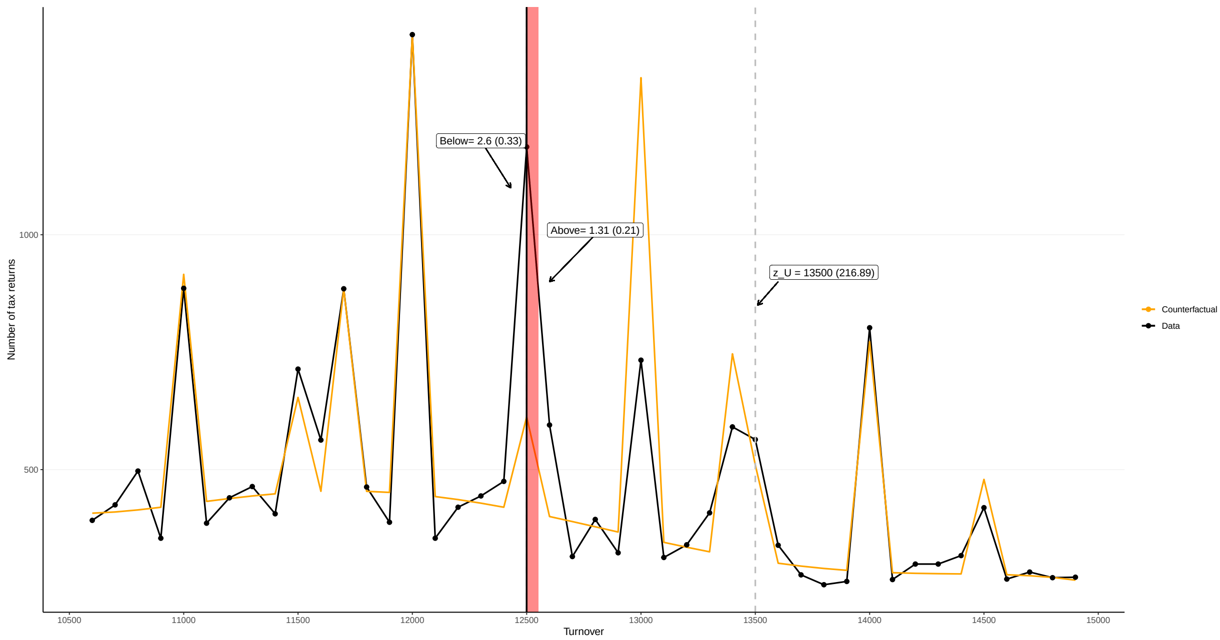

Notes: This figure plots the results of estimating bunching at the 8300 and 12500 thresholds (panel (a) and (b), respectively). The black lines depict the empirical densities. The yellow lines depict the counterfactual densities as estimated by Eq. (1.1), accounting for round turnover amounts as well as round payment amounts.The black solid vertical lines mark the threshold at which tax liability increases discretely i.e. the notch. The red area depicts the dominated range. The grey dashed vertical lines depict the lower and upper bounds of the omitted region z_L, z_U. Estimates of bunching below and above the threshold are derived from Eq. (1.2) & (1.3) and compare the counterfactual to the empirical density. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Years: 2017, 2018

Table 1.A.1: Bunching above by sector

	Dependent variable:	
	Ever bunched above	
Sector (reference: financial activities)	(1)	(2)
Accommodation and food service activities	**−0.020**	**−0.024**
	(0.035)	(0.035)
Administrative and support service activities	**−0.024**	**−0.027**
	(0.031)	(0.032)
Agriculture, forestry and fishing	**−0.014**	**−0.018**
	(0.031)	(0.032)
Construction	**0.003**	**0.0002**
	(0.032)	(0.032)
Education	**−0.011**	**−0.013**
	(0.037)	(0.038)
Electricity, gas, steam and air conditioning supply	**−0.043**	**−0.031**
	(0.052)	(0.054)
Human health and social work activities	**0.009**	**0.006**
	(0.042)	(0.042)
Information and communication	**−0.043**	**−0.046**
	(0.046)	(0.046)
Manufacturing	**−0.029**	**−0.032**
	(0.033)	(0.033)
Mining and Quarrying	**0.007**	**−0.044**
	(0.042)	(0.084)
Other service activities	**−0.021**	**−0.023**
	(0.029)	(0.029)
Professional, scientific and technical activities	**−0.029**	**−0.032**
	(0.033)	(0.033)
Real estate activities	**−0.023**	**−0.029**
	(0.034)	(0.035)
Transportation and storage	**−0.019**	**−0.023**
	(0.031)	(0.031)
Water supply; sewerage, waste management and remediation	**−0.043**	**−0.046**
	(0.084)	(0.084)
Wholesale and retail trade; repair of motor vehicles and motorcycles	**−0.015**	**−0.018**
	(0.029)	(0.029)
Taxoffice FE		✓
# Firms	22,361	22,361
R^2	0.001	0.002

Notes: This table shows the results from regressing an indicator variable of whether a firm has ever bunched *above* a threshold on sector indicator variables. The reference sector is "Financial and insurance activities". Standard errors are in parantheses. Not shown, but small and insignificant are the estimated coefficients for the sectors: "Activities of extraterritorial organizations and bodies", "Activities of households as employers; undifferentiated goods- and services- producing activities of households for own use", "Arts, entertainment and recreation", "Public administration and defence; compulsory social security" and non-classified sectors. Data source: ZRA. Years: 2017-2018.

Figure 1.A.2: Bunching around thresholds 16500 and 20800

a) Bunching at 16500

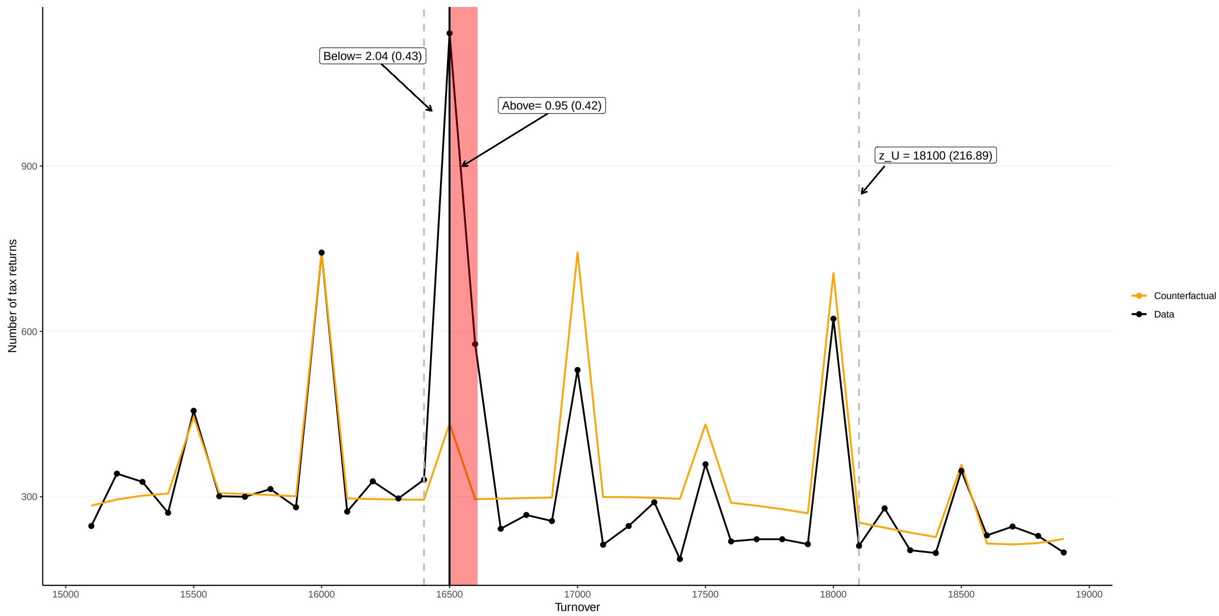

b) Bunching at 20800

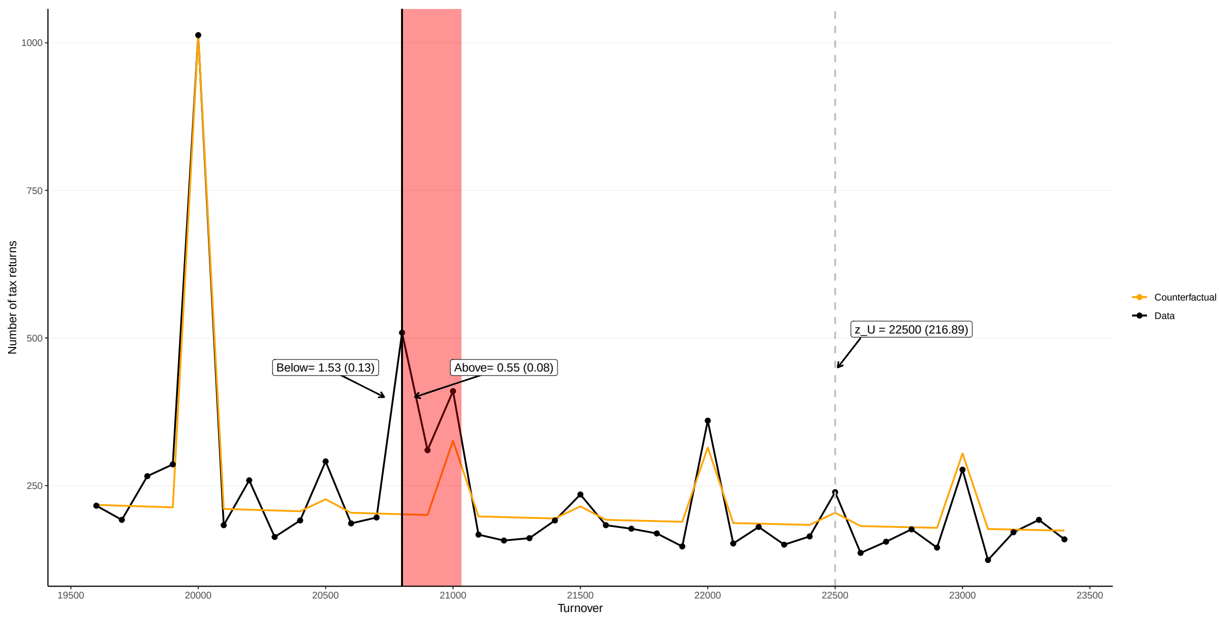

Notes: This figure plots the results of estimating bunching at the 16500 and 20800 thresholds (panel (a) and (b), respectively). The black lines depict the empirical densities. The yellow lines depict the counterfactual densities as estimated by Eq. (1.1), accounting for round turnover amounts as well as round payment amounts. The black solid vertical lines mark the threshold at which tax liability increases discretely i.e. the notch. The red area depicts the dominated range. The grey dashed vertical lines depict the lower and upper bounds of the omitted region z_L, z_U. Estimates of bunching below and above the threshold are derived from Eq. (1.2) & (1.3) and compare the counterfactual to the empirical density. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Years: 2017, 2018

Figure 1.A.3: Distribution before notched schedule

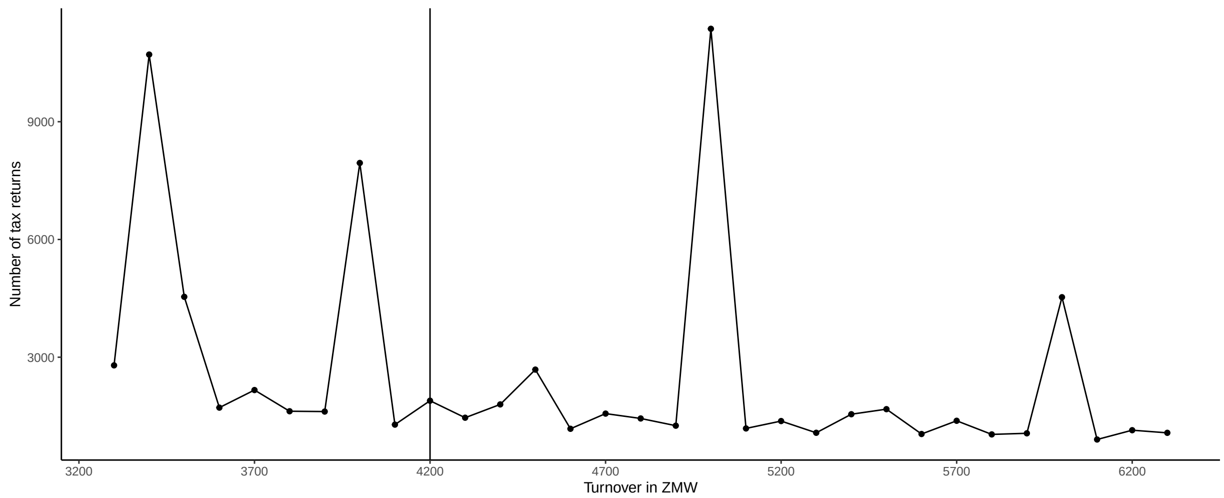

Notes: This figure plots the distribution of tax returns around the threshold before the notched schedule was in place. Data source: ZRA. Years: 2015, 2016

Figure 1.A.4: Distribution after notched schedule

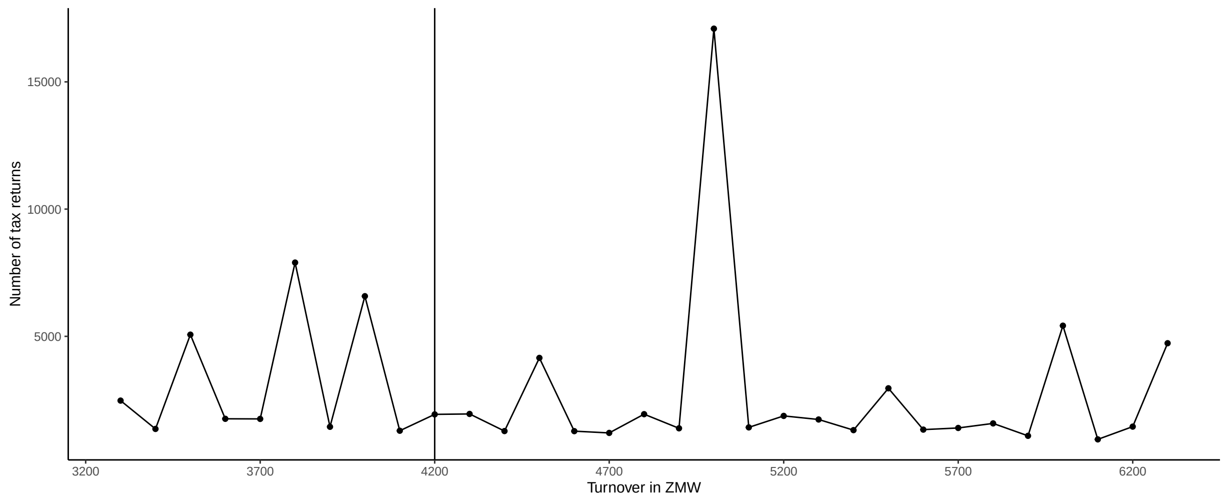

Notes: This figure plots the distribution of tax returns around the threshold after the notched schedule was in place. Data source: ZRA. Years: 2019-2021

Figure 1.A.5: Bunching with binsize of 50

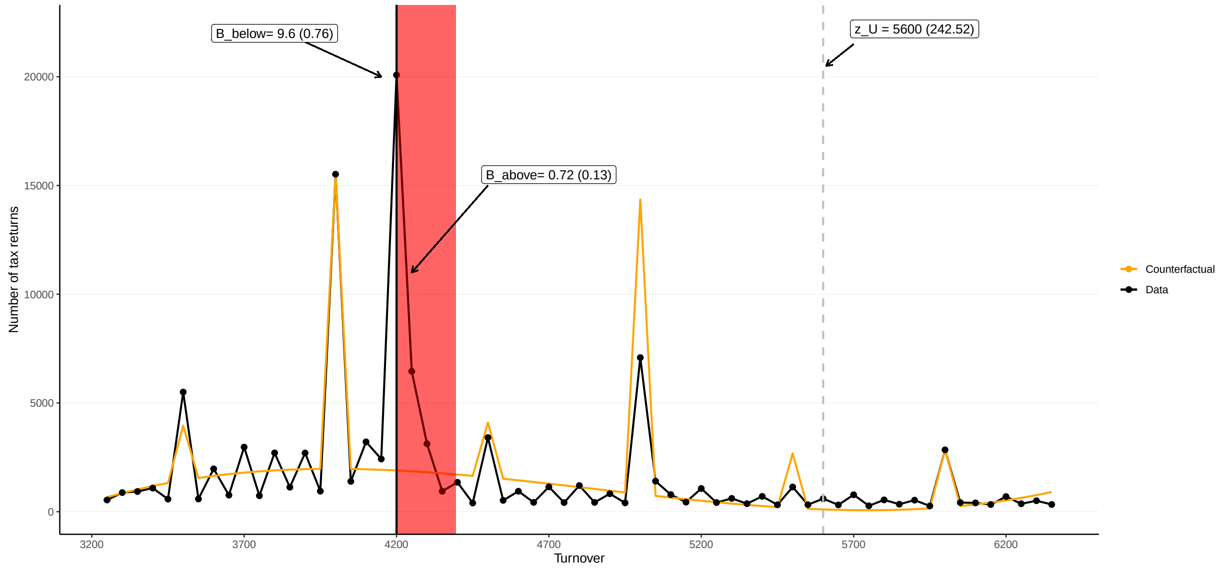

Notes: This figure plots the bunching response using a binsize of 50. Data source: ZRA. Years: 2017-2018

Figure 1.A.6: Bunching with pre-reform data as counterfactual

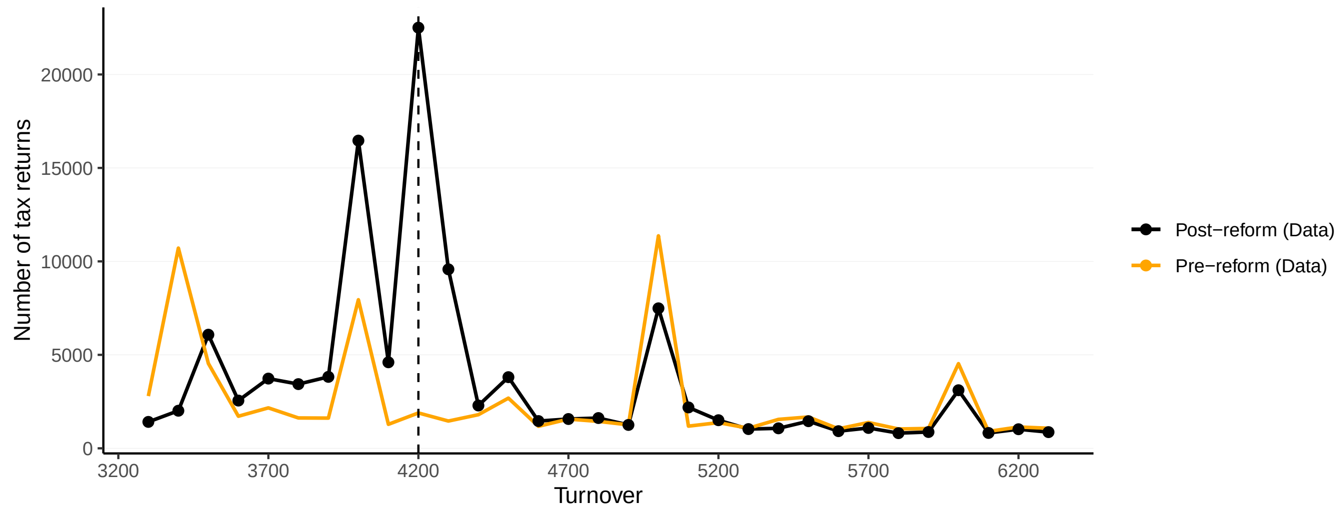

Notes: This figure plots the distribution of tax returns around the threshold before and during the notched schedule was in place. Data source: ZRA. Years: 2015-2018

Figure 1.A.7: Bunching patterns over time in 2017

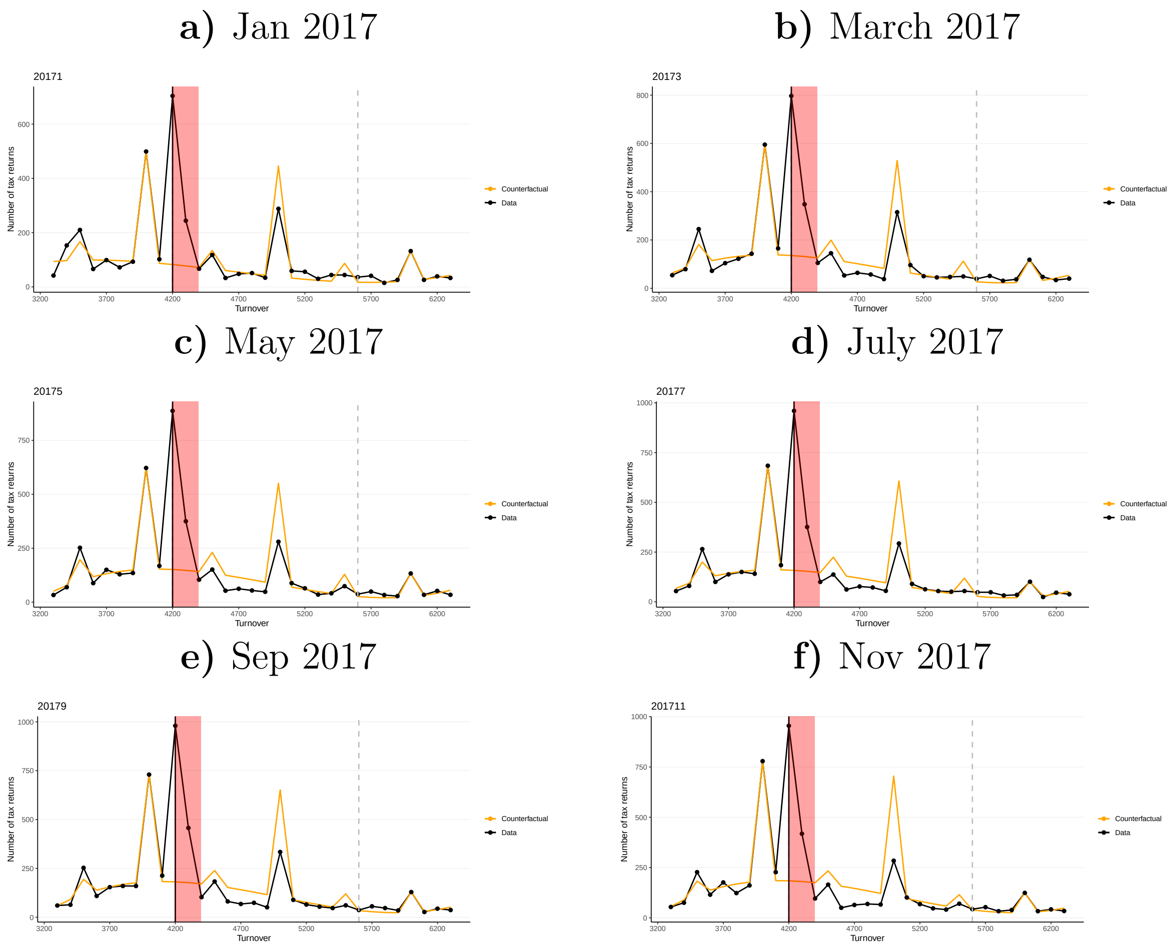

Notes: This figure plots the results of estimating bunching at the 4200 threshold in different months throughout 2017. The black line depicts the empirical density. The yellow line depicts the counterfactual density as estimated by Eq. (1.1), accounting for round turnover amounts as well as round payment amounts. The black solid vertical line marks the threshold at which tax liability increases discretely i.e. the notch. The red area depicts the dominated range. The grey dashed vertical line depicts the upper bound of the omitted region z_U. Estimates of bunching below and above the threshold are derived from Eq. (1.2) & (1.3) and compare the counterfactual to the empirical density. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Years: 2017

Figure 1.A.8: Bunching patterns over time in 2018

a) Jan 2018

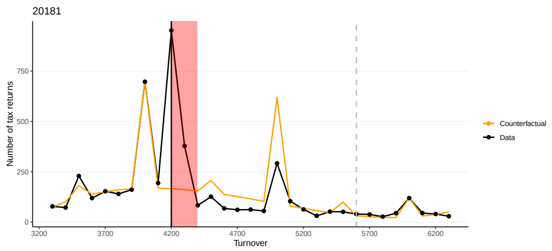

b) March 2018

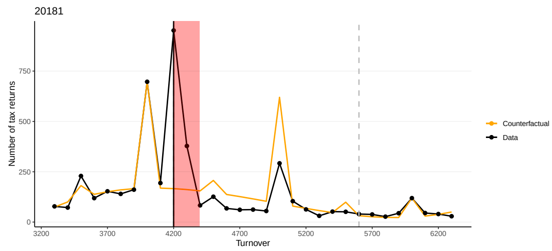

c) May 2018

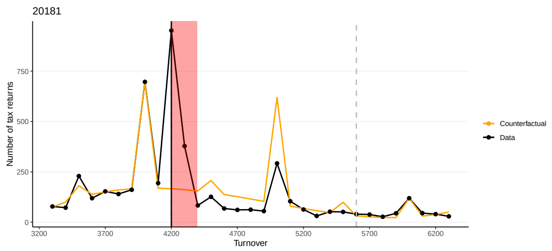

d) July 2018

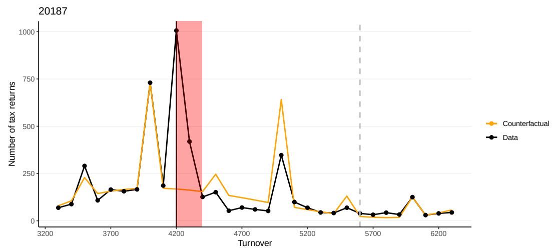

e) Sep 2018

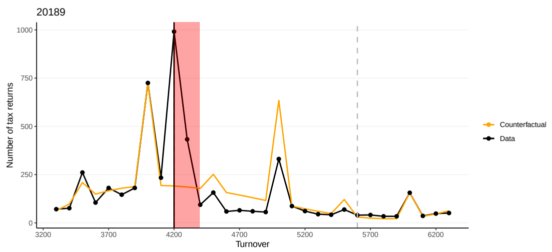

f) Nov 2018

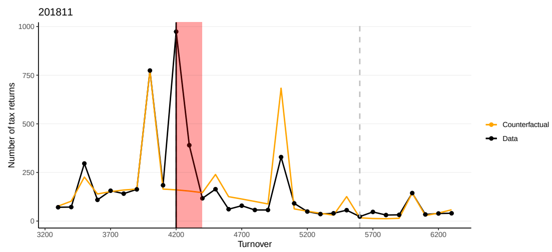

Notes: This figure plots the results of estimating bunching at the 4200 threshold in different months throughout 2018. The black line depicts the empirical density. The yellow line depicts the counterfactual density as estimated by Eq. (1.1), accounting for round turnover amounts as well as round payment amounts. The black solid vertical line marks the threshold at which tax liability increases discretely i.e. the notch. The red area depicts the dominated range. The grey dashed vertical line depicts the upper bound of the omitted region z_U. Estimates of bunching below and above the threshold are derived from Eq. (1.2) & (1.3) and compare the counterfactual to the empirical density. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Years: 2018

Figure 1.A.9: Distribution of firms after bunching *above*

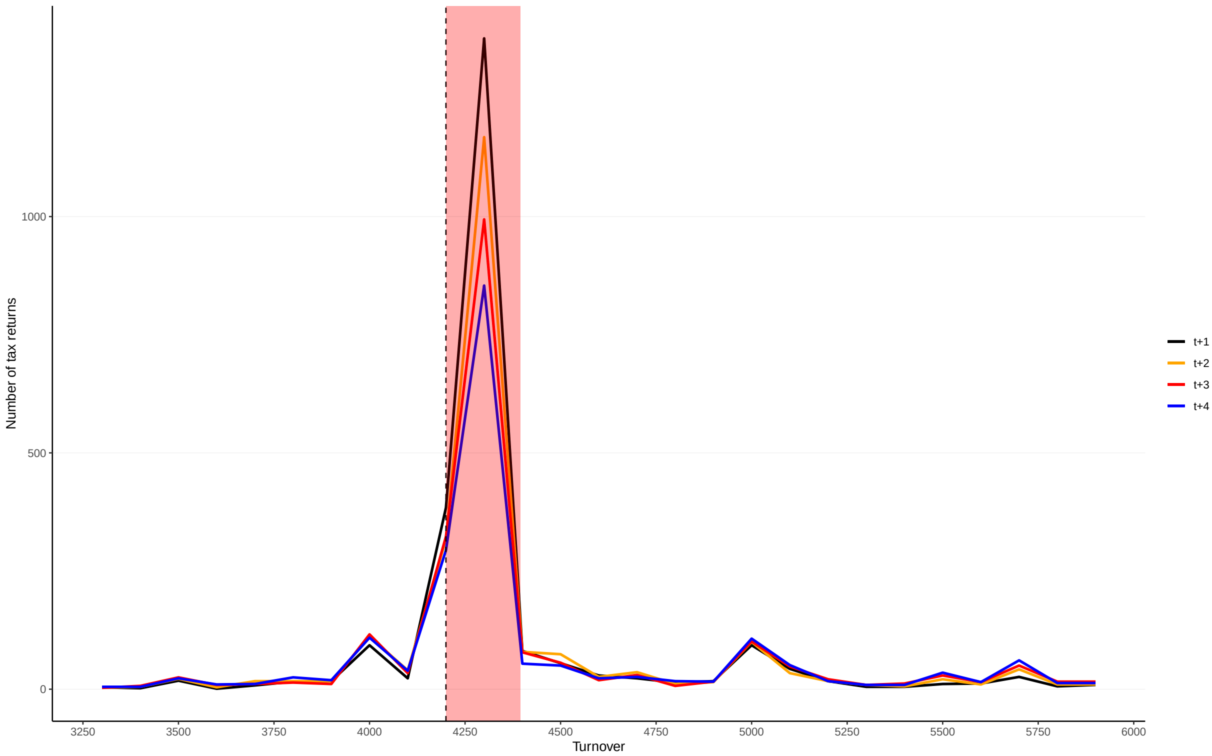

Notes: This figure plots the distribution of firms that bunched above the threshold at 4200 in a any month in the periods thereafter. The red area depicts the dominated range. Panel (a) shows the distribution one month after, panels (b) and (c) two and three months after having bunched *above*.

Figure 1.A.10: Distribution of firms after bunching *above* (only once)

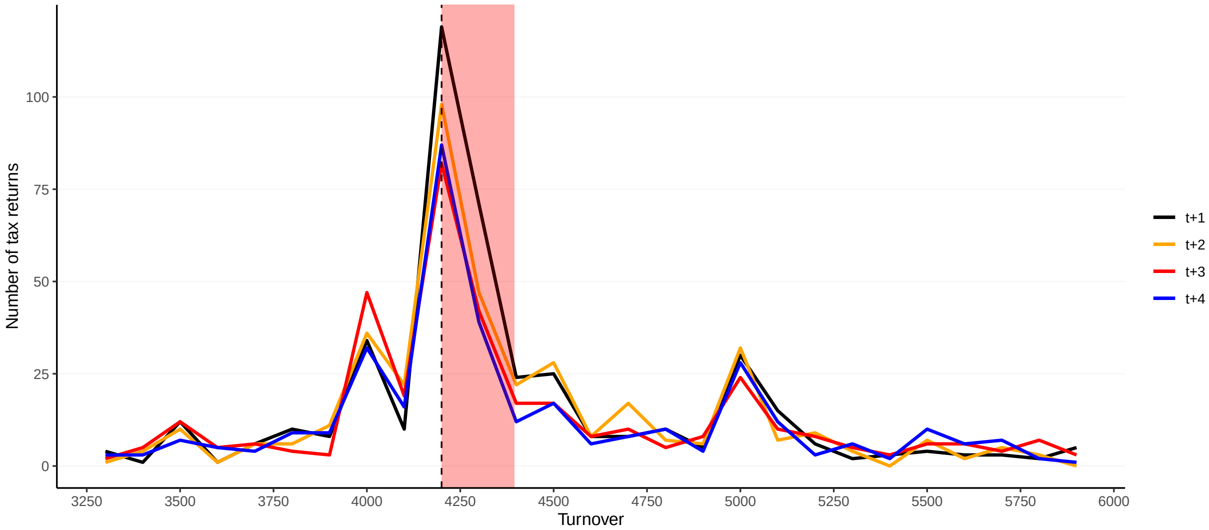

Notes: This figure plots the distribution of firms that bunched above the threshold at 4200 in any month month in the period thereafter. The red area depicts the dominated range.

Figure 1.A.11: Number of times firms bunch *above* the thresholds

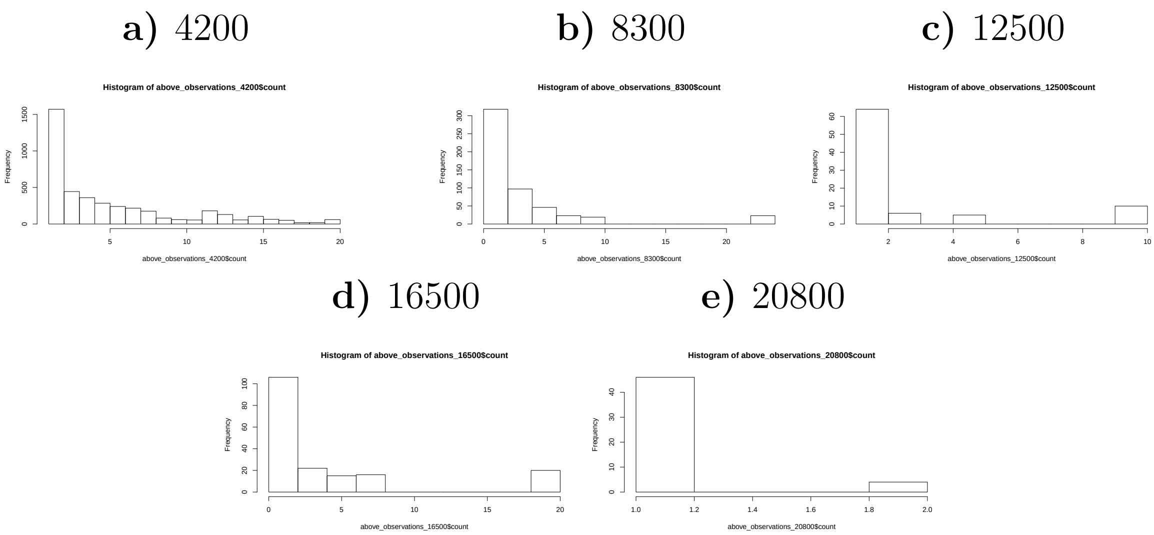

Notes: This figure plots the distribution of firms bunching above the thresholds according to the frequency they do so. For example, panel (a) shows that over 1500 firms which bunch above the 4200 threshold do so only once. Data source: ZRA. Years: 2017, 2018

Table 1.A.2: Bunching above by taxoffice

	Dependent variable:	
	Ever bunched above	
Tax office (reference: Lusaka)	(1)	(2)
Central Province	**0.0004**	**0.001**
	(0.006)	(0.006)
Chingola	**0.003**	**0.004**
	(0.008)	(0.008)
Choma	**−0.010**	**−0.009**
	(0.007)	(0.007)
Eastern Province	**0.015****	**0.015****
	(0.007)	(0.007)
Kitwe	**−0.008**	**−0.008**
	(0.005)	(0.005)
Livingstone	**0.003**	**0.002**
	(0.009)	(0.009)
Luapula Province	**0.011**	**0.011**
	(0.010)	(0.010)
Muchinga Province	**−0.004**	**−0.004**
	(0.011)	(0.011)
Ndola	**0.006**	**0.007**
	(0.005)	(0.005)
Northern Province	**0.003**	**0.003**
	(0.009)	(0.009)
Northwestern Province	**−0.005**	**−0.005**
	(0.008)	(0.008)
Western Province	**0.003**	**0.004**
	(0.009)	(0.009)
Small Taxpayer Offices (Mining– , Non Mining– , VAT North combined)	**0.036**	**0.057**
	(0.033)	(0.085)
Sector FE		✓
# Firms	22,361	22,361
R^2	0.001	0.002

Notes: This table shows the results from regressing an indicator variable of whether a firm has ever bunched *above* a threshold on tax office indicator variables. The reference tax office is "Lusaka". Standard errors are in parantheses. Data source: ZRA. Years: 2017-2018.

1.B Bunching at round number tax liabilities.

This Section relates to the analysis in Section 1.4.1 and provides further evidence on bunching at round number tax liabilities. Table 1.B.1 provides the results from estimating Eq. (1.4) (i.e. round number tax liabilites which *do not coincide* with round number turnover amounts, as graphically depicted in Figure 1.4), pooling all tax returns from 2015-2016. Bootstrapped standard errors are in parentheses. Tables 1.B.2 and 1.B.3 show results for the same exercise in only 2015 and 2016, respectively. Further, in Figure 1.B.1 we plot the distribution of deviations from tax liabilities which imply multiples of 10. We do so by creating tax liability bins of size 10, with the multiple of 10 as the mid point, e.g., [5,15]. Then, we calculate how far away each tax return is from its own bin's midpoints. We do so for all bins, for bins with round turnover amounts as the mid point and for those with only round liability amounts as the mid point.

Table 1.B.1: Bunching at round number tax liabilities 2015-2016

	Liabilities					
	40	50	70	80	100	110
B	**9.54**	**30.59**	**11.21**	**18.15**	**10208.6**	**26.79**
	(0.44)	(0.01)	(1.26)	(5.51)	(11.87)	(49.49)
	130	140	160	170	190	200
B	**6.08**	**4.46**	**6.66**	**7.03**	**3.83**	**56.69**
	(0.38)	(0.31)	(0.45)	(0.58)	(0.41)	(5.05)
	220	230	250	260	280	290
B	**5.75**	**3.72**	**28.86**	**3.65**	**5.77**	**2.95**
	(0.3)	(0.21)	(1.3)	(0.24)	(0.46)	(0.33)

Notes: This table shows the estimated bunching coefficients for bunching at round number tax liabilities which *do not coincide* with round turnovers (Eq. (1.4)) and bunching for turnover interval plotted in Figure 1.4. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Years: 2015, 2016.

Table 1.B.2: Bunching at round number tax liabilities in 2015

	Liabilities					
	40	50	70	80	100	110
B	**8.10**	**28.16**	**10.10**	**17.1**	**2806.06**	**24.67**
	(0.61)	(1.79)	(1.07)	(9.74)	(2677.99)	(707.57)
	130	140	160	170	190	200
B	**5.82**	**4.90**	**6.57**	**5.89**	**3.26**	**52.51**
	(0.36)	(0.33)	(0.46)	(0.53)	(0.38)	(5.02)
	220	230	250	260	280	290
B	**5.08**	**3.05**	**26.63**	**4.13**	**5.36**	**2.81**
	(0.28)	(0.21)	(1.63)	(0.33)	(0.53)	(0.35)

Notes: This table shows the estimated bunching coefficients for bunching at round number tax liabilities which *do not coincide* with round turnovers (Eq. (1.4)) and bunching for turnover interval plotted in Figure 1.4. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Year: 2015

Table 1.B.3: Bunching at round number tax liabilities in 2016

	Liabilities					
	40	50	70	80	100	110
B	**8.95**	**24.13**	**10.14**	**28.9**	**3446.01**	**357.01**
	(0.73)	(1.75)	(1.06)	(8.27)	(224.93)	(558.99)
	130	140	160	170	190	200
B	**7.32**	**4.63**	**7.15**	**8.36**	**4.46**	**63.33**
	(0.42)	(0.29)	(0.46)	(0.64)	(0.43)	(5.26)
	220	230	250	260	280	290
B	**6.31**	**4.62**	**31.44**	**3.08**	**6.02**	**3.61**
	(0.34)	(0.28)	(1.57)	(0.21)	(0.51)	(0.35)

Notes: This table shows the estimated bunching coefficients for bunching at round number tax liabilities which *do not coincide* with round turnovers (Eq. (1.4)) and bunching for turnover interval plotted in Figure 1.4. Standard errors are derived from bootstrapping the residuals of the counterfactual density estimation and shown in parentheses. Data source: ZRA. Year: 2016

1.B. Bunching at round number tax liabilities.

Figure 1.B.1: Deviations from 10x

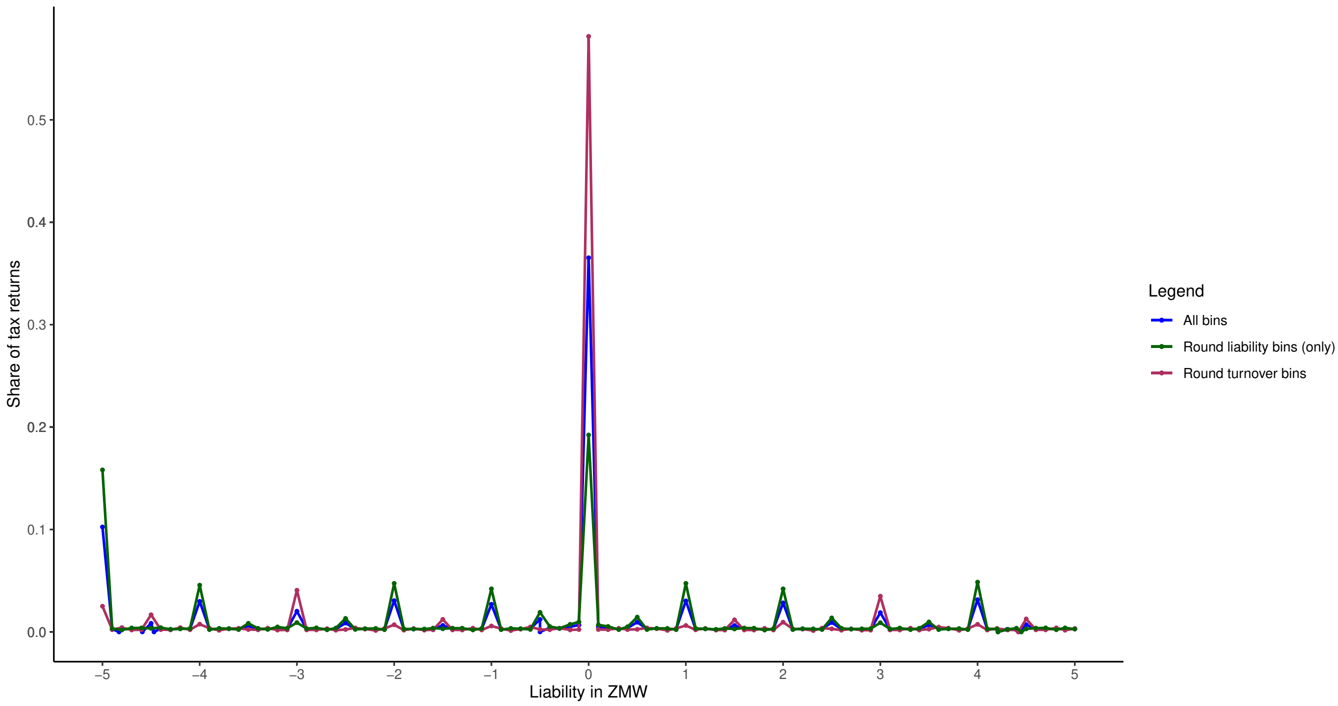

Notes: This figure plots the distribution of how much tax returns deviate from the closest turnover amount, which results in a tax liability that is a multiple of 10. *Round liability bins (only)* implies that we only consider tax returns whose closest liability which is divisible by 10 results from odd turnover (e.e. 1333,33). *Round turnover bins* implies that we only consider tax returns also coincides with round turnover as well (e.g. 2000). Data source: ZRA. Years: 2015, 2016.

1.C Audit Probabilites

This section complements Section 1.6.1. First, we estimate empirical audit probabilities by turnover. Then, we show theoretically that differential audit probabilities on both sides of the threshold are an unplausible explanation for the phenomenon of bunching *above*.

Empirical audit probabilities around the threshold. We leverage the administrative data on tax audits to test whether there were indeed differential audit probabilities on both sides of the threshold. As most audits in the data refer to more than one tax period, we connect the audits to the tax returns by matching the month the tax return refers to to the audit period end date. This assumes that the last return that has been filed in the tax audit period was the one that triggered the audit. We can therefore infer implicit audit rules, i.e. the empirical probability of being audited conditional on reporting a certain amount of turnover. Figure 1.C.1 plots this empirical audit probability according to turnover bins of size 100 ZMK around the 4200 and 8300 thresholds. As can be seen, conditional audit probabilities were very low for all bins but importantly, there was no substantial difference between the probability of being audited when reporting turnover between $(4100, 4200]$ and $(4200, 4300]$ or $(8200, 8300]$ and $(8300, 8400]$, respectively. Note, that for the other thresholds, there had not been any audits on either side of the threshold throughout 2017 and 2018. This evidence alleviates the concern that bunching *above* thresholds is driven by implicit audit rules. We further address concern from a theoretical perspective in the following.

Tax evasion model and simulation. We consider a firm with turnover z which decides on the tax payment T it plans to make. The firm has quasi-linear utility in z according to

$$U(T,z) = z - T + v(T,z) \tag{1.15}$$

where $v(T,z)$ is an increasing and concave function in T and governs the decision of how much taxes to pay.

When thinking about how to choose the tax payment T, the firm considers the probability of being caught and the expected penalty. To arrive at a desired tax payment the firm inverts the tax schedule denoted by $T_{schedule}$ and reports turnover accordingly. Thus, for a firm with turnover z that makes a payment T, the probability of being caught is $p(z - T^{-1}_{schedule}(T))$. The firm also considers the penalty which we model to be proportional

Figure 1.C.1: Empirical audit probabilities

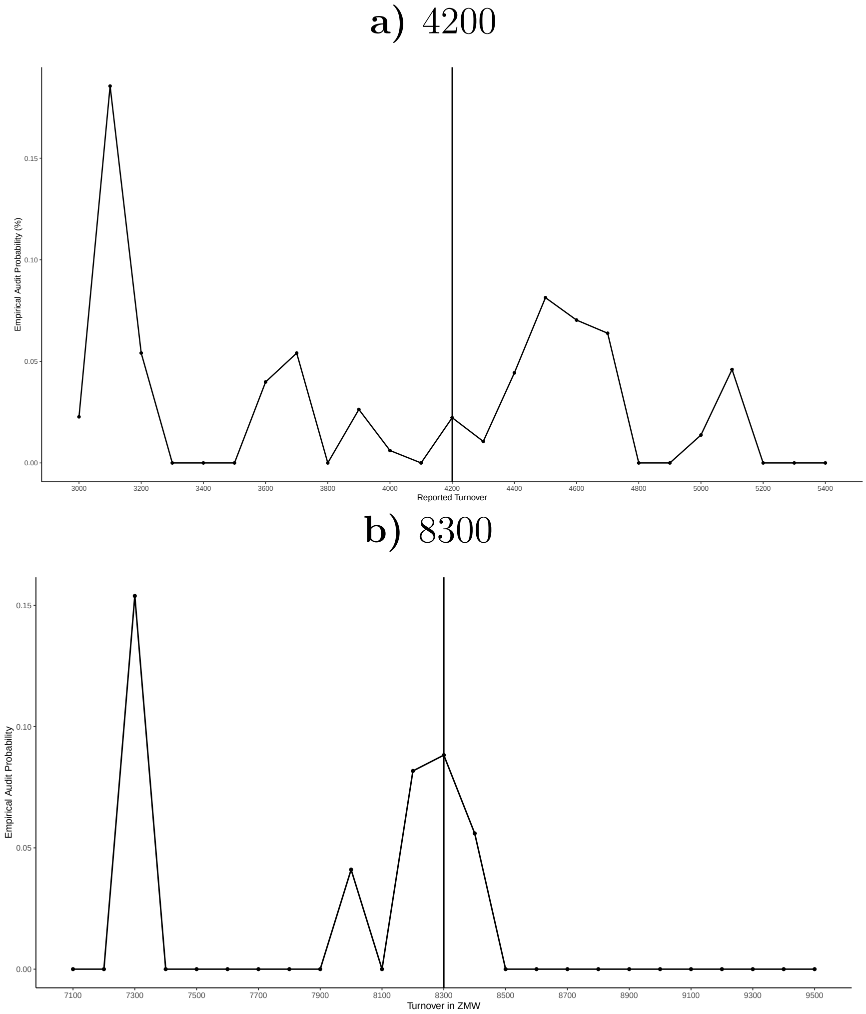

Notes: This figure plots the empirical audit probability (i.e. share of audits per bin) over turnover bins of 100 ZMK. Data source: ZRA. Years: 2017,2018.

to the evaded amount: $\xi(T(z)-T)$. Implementing the notched schedule from Table 1.1, we therefore define

$$v(T,z) = p\left(z - \left(\frac{T-F}{t} + \bar{z}\right)\right)(T - T(z))(1+\xi) \tag{1.16}$$

Figure 1.C.2: Support for bunching *above* in the audit model

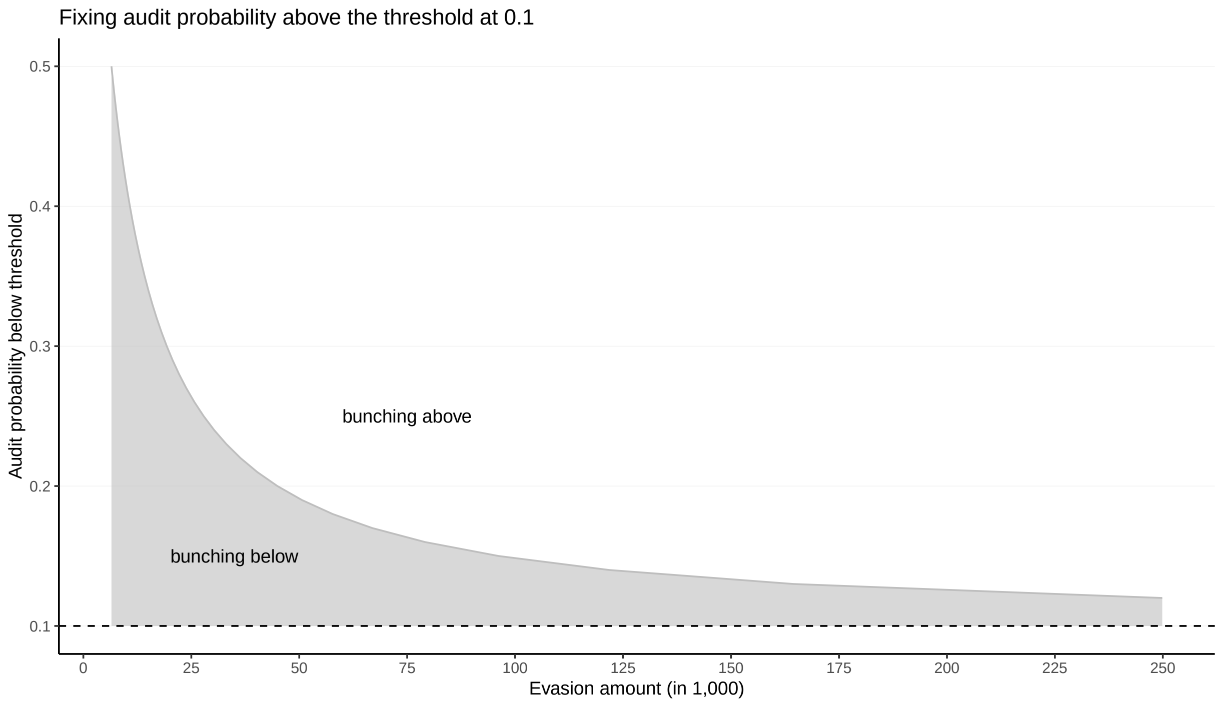

Notes: This figure illustrates under which circumstances firms would choose to bunch *above* – instead of below the threshold due to differential audit probabilities on both sides. The horizontal axis plots the amount of taxes a firm evades while reporting a turnover of 4200 (i.e. bunch below the threshold). When fixing the audit probability above the threshold at 10% for all evaded amounts, the grey area depicts the audit probabilites below the threshold that would support bunching below the threshold in the model sketched above (akin to Kleven, Knudsen, Kreiner, Pedersen, and Saez (2011)), given true turnover i.e. the amount that is evaded. The white area shows the combinations that support bunching *above*.

such that $U(T,z)$ becomes

$$\begin{aligned}
U(z,T) &= z-T+p\left(z-\left(\frac{T-F}{t}+\bar{z}\right)\right)(T-T(z))(1+\xi) \\
&= z-T+p\left(z-\left(\frac{T-F}{t}+\bar{z}\right)\right)(z-z+T-T(z)+(T-T(z))\xi) \\
&= \left(1-p\left(z-\left(\frac{T-F}{t}+\bar{z}\right)\right)\right)(z-T) \\
&+p\left(z-\left(\frac{T-F}{t}+\bar{z}\right)\right)(z-T(z)-(T(z)-T)\xi)
\end{aligned}$$

which is akin to the evasion function from Kleven, Knudsen, Kreiner, Pedersen, and Saez (2011).

Within this framework, we now investigate which assumptions would generate bunching *above* the threshold. In particular, we consider a firm with true turnover way above the threshold $z >> \bar{z}$ which evades taxes up to the threshold and contemplates choosing to report just below or just above the threshold. If the firm chooses to report above the threshold instead of below this would imply:

$$\begin{aligned}
&U(z,225) > U(z,36) \\
&z - 225 + p(z - 4200 + \epsilon)(225 - T(z))(1+\xi) > z - 36 + p(z-4200)(36 - T(z))(1+\xi) \\
&\frac{189}{1+\xi}(1 - p(z-4200)(1+\theta)) < (p(z-4200) - p(z-4200+\epsilon))(z-\bar{z})t
\end{aligned}$$

In the Zambian context, one can assume that $\theta \leq 0.095$ (Zambia Revenue Authority, 2022). To investigate under which circumstances, a firm will choose to bunch above, we simulate turnover and audit probabilities. In Figure 1.C.2 the white area shows the combinations of audit probabilities below the threshold (y-axis) and the amount that is evaded (x-axis) that support bunching *above* when the audit probability above the threshold is fixed at 10%.

Overall, these simulation results do not support the hypothesis that bunching *above* thresholds, is the result of differing audit probabilities on both sides of the threshold. For example, even if the audit probability would jump from 10% to 30% at the threshold, only firms evading at least 25,000 (i.e. 83% of turnover) would choose to bunch *above*.

Simulation with risk aversion. We extend the model to allow for different levels of risk aversion. Instead of quasi-linear utility, we now consider a risk-averse firm with CRRA utility of the following form:

$$U(z,T) = \frac{(z - T + v(z,T))^{1-\sigma} - 1}{1-\sigma}. \tag{1.17}$$

With this utility function at hand, we calculate under which circumstances the firm's expected utility[19] is larger when bunching above than when bunching below, given that the audit probability is 10 % above the threshold. Let the audit probability above the threshold given z be denoted by $p_{a,z}$. The condition for bunching above is then given by:

[19]Note that under the assumption of quasi linear utility, as above, the expectation function is linear, i.e. $\mathbf{E}[z - T + v(z,T)] = z - T + v(z,T)$.

$$\mathbf{E}[U(z,225)] > \mathbf{E}[U(z,36)]$$
$$p_{a,z} > \frac{(z-225)^{1-\sigma} - (z-36)^{1-\sigma} + 0.1((z-225-(z-4200)t(1+\xi))^{1-\sigma} - (z-225)^{1-\sigma}}{(z-36-(225+(z-4200)t-36)(1+\xi))^{1-\sigma} - (z-36)^{1-\sigma}}$$

The resulting combinations of $p_{a,z}$ and the evaded amount $z-4200$ change in comparison to Figure 1.C.2. In particular, the curve becomes flatter. Assuming a risk aversion parameter of $\sigma = 2$, we now find that under the scenario from before, where audit probability triples above the threshold (from 10% to 30%), only firms evading at least 58% of their turnover would choose to bunch above. With higher risk aversion of $\sigma = 3$, this number drops to about 50%.

1.D Revenue Targets

This section relates to arguments made in Section 1.5.3. To test the role of tax officials' incentives, we regress an indicator of whether a tax return classifies as *bunching above* on an indicator of whether a tax office has already reached its revenue target or not. The estimating equation reads:

$$\mathbf{1}(\text{bunched above})_{i,t} = \alpha + \beta \mathbf{1}(\text{target reached})_{o,t-1} + \mathbf{X} + \epsilon_i \tag{1.18}$$

where $\mathbf{X}$ represent fixed effects for tax office (denoted by o), a firm's sector and the threshold at which the tax return bunched (i.e. 4200, 8300,...). Table 1.D.1 presents the results.

Table 1.D.1: Tax office revenue targets and *bunching above*

Panel A: All tax returns	*Dependent variable:*			
	Bunched above			
	Whole year		Second half of year	
	(1)	(2)	(3)	(4)
Revenue target reached (0/1)	**−0.0001**	**−0.0001**	**−0.0004*****	**−0.0004*****
	(0.0001)	(0.0001)	(0.0002)	(0.0002)
Observations	1,403,815	1,403,815	751,349	751,349
R^2	0.0001	0.0002	0.0001	0.0002
Taxoffice FE	✓	✓	✓	✓
Sector FE		✓		✓
Baseline mean	**0.0036**	**0.0036**	**0.0039**	**0.0039**
Panel B: Only bunchers	*Dependent variable:*			
	Bunched above			
	Whole year		Second half of year	
	(1)	(2)	(3)	(4)
Revenue target reached (0/1)	**−0.008**	**−0.008**	**−0.012**	**−0.013**
	(0.007)	(0.007)	(0.008)	(0.008)
Observations	25,598	25,598	14,264	14,264
R^2	0.003	0.006	0.004	0.008
Threshold FE	✓	✓	✓	✓
Taxoffice FE	✓	✓	✓	✓
Sector FE		✓		✓
Baseline mean	**0.204**	**0.204**	**0.206**	**0.206**

Notes: This table shows the estimated correlations between the event of a tax return being *above* a threshold in a given month and whether a tax office's yearly revenue target has been reached in the previous month. The revenue targets are defined as the total turnover tax collections in the previous year. Panel A shows the results when including all tax returns. Panel B shows the results when restricting the sample to tax returns that were bunching either just above (e.g. at 4201) or just below (e.g. at 4200) a threshold. Data source: ZRA. Years: 2016-2018.

1.E Survey experiment

This section provides more detailed information on the survey experiment, outlined in Section 1.6.1.

Sampling and randomization. The 517 survey respondents were sampled from market places and business districts across Lusaka. Interviews were held in person by a total of 5 surveyors, who each handled an approximately equal share of the 517 interviews. Firm-owners or people running shops were randomly approached and interviewed. In some cases, the respondent provided suggestions as to where the enumerator could find other firms nearby which could also be interviewed (*snowball approach*). As on each surveying day, new firms were randomly approached, we view the overall sampling as quasi-random. To account for potential endogeneity induced by enumerators relying on such firm networks, we control for enumerator fixed effects in our estimations.

The experimental component was tied to the end of the survey and consisted of randomly providing 3 different pieces of information. Randomization took place at the individual respondent level. That is, in each interview, a random number generator assigned the respondent to 1 of 3 groups, each associated with a different information treatment. This randomization allows us to estimate the causal effect of the information treatment on the answers recorded after the post-treatment. In a later robustness check, we again randomly assigned the respondents into two further groups.

Treatment Messages. The 3 groups mutually exclusively received the following messages:

1. **Control.** "We are now reaching the end of the survey. At this point, we want to share some information with you that the University of Mannheim has gathered. In the year 2022 over 170,000 firms in Zambia were registered under Turnover Tax. Were you aware of the information we just shared with you?"

2. **Audit Treatment.** "We are now reaching the end of the survey. At this point, we want to share some information that the University of Mannheim has gathered on the audits conducted by ZRA. Over the last two years, the total sales which were audited by ZRA increased by almost 20 million Kwacha. This is an increase of 50%. The penalties that had to be paid amounted to 20 million Kwacha. Were you aware of the information we just shared with you?"

3. **Contract Treatment.** "We are now reaching the end of the survey. At this point,

> we want to share some information with you that the University of Mannheim has gathered. The government of Zambia has comitted to awarding 20% of their business contracts to small firms like yours. This is a large amount of potential business. For securing government contracts, a firm requires clearance from ZRA. Were you aware of this information we just shared with you?"

Outcomes. The aim of the experiment was to test alternative explanations for why firms bunch above a threshold instead of below (besides bargaining). As we, unfortunately, can not link survey respondents to their actual tax declarations, we rely on a stated measure of *bunching above*, as follows. In the beginning of the survey, respondents were asked about the monthly turnover they usually have. Let this stated turnover be denoted by X. Later, but before the treatment occured, the respondent was asked, which tax payment the respondent would find appropriate for a firm with turnover X. Let this preferred tax payment be denoted by Y. After the treatment occured, we confront the respondents with a hypothetical notch, i.e., a region of payments that is not reachable. In this situation, the respondent was asked whether it would rather deviate below or above from its previously stated appropriate payment Y. In particular, the respondent was asked:

> "You have indicated that Y Kwacha would be an appropriate tax payment for a business like yours. Now, please think of a scenario in which the tax schedule does not allow a payment of Y Kwacha. Instead, the tax schedule would only allow payments of either $(1-0.1)\times Y$ Kwacha or $(1+0.1)\times Y$ Kwacha. Which payment would you choose instead?"

If a respondent states to rather pay the larger amount, i.e., $(1+0.1)\times Y$, we interpret this as a propensity to bunch above a threshold. Arguably, this hypothetical situation is similar to the one, a taxpayer would face when filing taxes (without any interaction with tax officials). In the survey, the respondent can simply choose to either pay the larger or the lower amount. In reality, being below or above the threshold is a matter of reporting only a slightly different amount of turnover.

In principle, respondents' answers could be influenced by legal considerations. For example, if firms stated their actual liability as the appropriate payment, i.e., $Y = 0.04\times X$. Then choosing $(1-0.1)\times Y$ could be viewed as illegal by taxpayers. To check whether these concerns matter for the outcomes, we randomly assign respondents into two groups. We reframe the question and exogenously fix the respondents' turnover to $X = 10,000$ and their appropriate payment Y to 300 and 500, respectively. In this case, for the first group

both payments above and below would be illegal (< 400) while for the second group both answers depict legal amounts (> 400).
Estimating the effects of the randomized information treatments can inform us about the channels which could drive bunching above. The results are presented in Section 1.6.1 and show that there is no significant effect of either treatment. This supports the notion that bargaining is the most likely explanation for firms bunching above the threshold. Clearly, as in any survey experiment, the outcomes we measured are only stated and do not necessarily coincide wth actions. However, being unable to match survey respondents to their administrative tax records, we consider our approach second-best.

1.F Proofs

This section provides the proofs for the propositions stated in Section 1.5.

1.F.1 Proof of Proposition 1

We start by characterizing the optimal tax payment T^* in the non-cooperative setting, namely the solution to Eq. (1.6). The first order condition reads:

$$(1-p)v'(z-T) = p\xi v'(z-T(z)(1+\xi)+T\xi)$$
$$\iff$$
$$(v')^{-1}(1-p)(z-T) = (v')^{-1}(p\xi)(z-T(z)(1+\xi)+T\xi)$$

After rearranging, we can write $T^*(p)$ as a linear combination of z and $T(z)$ as follows:

$$T^*(p) = z\underbrace{\left(\frac{(v')^{-1}(1-p)-(v')^{-1}(p\xi)}{(v')^{-1}(1-p)+\xi(v')^{-1}(p\xi)}\right)}_{\equiv \mathbf{K_1}} + T(z)\underbrace{\left(\frac{(v')^{-1}(p\xi)(1+\xi)}{(v')^{-1}(1-p)+\xi(v')^{-1}(p\xi)}\right)}_{\equiv \mathbf{K_2}}. \tag{1.19}$$

Note that $K_1 < 0$, $K_2 > 1$ and $K_1 + K_2 = 1$.[20] Furthermore, if $\frac{1-p}{p} = \xi$, then $K_1 = 0$ and $K_2 = 1$. As $\frac{1-p}{p}$ is decreasing in p, this means that there is full compliance (i.e. $T^* = T(z)$), if either ξ or p is sufficiently large such that

$$1 = (1+\xi)p. \tag{1.20}$$

To prove that $T(z)$ lies outside of the bargaining set (T_G, T_F) it is sufficient to show that $T(z) > T_F$. This is equivalent to:

$$T(z) > z - v^{-1}\left((1-p)v(z-K_1z-K_2T(z)) + pv(z-(1+\xi)T(z)+\xi K_1 z+\xi K_2 T(Z))\right)$$
$$\iff$$
$$v(z-T(z)) < (1-p)v(z-K_1z-K_2T(z)) + pv(z-(1+\xi)T(z)+\xi K_1 z+\xi K_2 T(Z))$$

Now, one can see that as p approaches its maximum value $p = \frac{1}{1+\xi}$, the above inequality will become an equality, because $T^*(\frac{1}{1+\epsilon}) = T(z)$, as can be seen from Eq. (1.19). Thus, to show that the inequality holds strictly for all other audit probabilities, it suffices to show

[20] To reconcile the inequalities note our assumption that $\frac{1-p}{p} \geq \xi$ and that $v()$ is concave.

that the right hand side is strictly decreasing in p for all $p < \frac{1}{1+\epsilon}$. We therefore take the first derivative of the right hand side with respect to p:

$$\begin{aligned}\frac{\partial()}{\partial p} =& -(1-p)T^*(p)v'(z-T^*(p)) - v(z-T^*(p)) \\ &+ \xi p T^{*\prime}(p)v'(z-(1+\xi(T(z)+\xi T^{*\prime}(p)) + v(z-(1+\xi)T(z)+\xi T^*(p)) \\ =& \underbrace{v(z-(1+\xi)T(z)+\xi T^*(p)) - v(z-T^*(p))}_{<0} \\ &+ T^{*\prime}(p)\underbrace{(\xi p v'(z-(1+\xi)T(z)+\xi T^*(p)) - (1-p)v'(z-T^*(p)))}_{=0} \\ &< 0\end{aligned}$$

where the last term is zero by the Envelope Theorem and the optimality condition for $T^*(p)$ from Eq. (1.6). This concludes the proof of Proposition 1. □

1.F.2 Proof of Proposition 2.

We start by revisiting the government's optimal choice of p in Eq. (1.8). The solution reads:

$$c'(p) = T^{*\prime}(p) + (T(z) - T^*(p) - pT^{*\prime}(p))(1+\xi) \tag{1.21}$$

We multiply both sides of the equation by $\frac{p}{c(p)}$ to express the condition in terms of the cost elasticity κ:

$$\kappa = \left(T^{*\prime}(p) + (T(z) - T^*(p) - pT^{*\prime}(p))(1+\xi)\right)\frac{p}{c(p)}. \tag{1.22}$$

Taking the total differential of Eq. (1.22) and rearranging, we get

$$\frac{dp}{d\kappa} = \left[\frac{\Xi}{c(p)^2}\right]^{-1} \tag{1.23}$$

where Ξ is given by

$$\begin{aligned}\Xi =& (T^{*\prime\prime}(p) - (T^{*\prime}(p) + pT^{*\prime\prime}(p) + T^{*\prime}(p))(1+\xi))pc(p) \\ &+ T^{*\prime}(p) + (T(z) - T^*(p) - pT^{*\prime}(p))(1+\xi))c(p) \\ &- (T^{*\prime}(p) + (T(z) - T^*(p) - pT^{*\prime}(p))(1+\xi))pc'(p)\end{aligned}$$

which, again can be simplified to

$$\begin{aligned}\Xi =&(T^{*\prime}(p)-(T(z)-T^*(p)-pT^{*\prime}(p))(1+\xi))(c(p)-pc'(p))\\&+(T^{*\prime\prime}(p)-(T^{*\prime}(p)+pT^{*\prime\prime}(p)+T^{*\prime}(p))(1+\xi))pc(p)\\=&(\underbrace{(T^{*\prime}(p)(1-p(1+\xi))}_{\geq 0}+\underbrace{(T(z)-T^*(p))(1+\xi))}_{\geq 0}\underbrace{(c(p)-pc'(p))}_{\leq 0}\\&+\underbrace{(T^{*\prime\prime}(p)-(T^{*\prime}(p)+pT^{*\prime\prime}(p)+T^{*\prime}(p))(1+\xi))pc(p)}_{\leq 0}\leq 0\end{aligned}$$

Finally, if $\Xi \leq 0$, it follows that $\frac{dp}{d\kappa} \leq 0$.

We have now established that the optimally chosen audit probability p is increasing as the audit cost elasticity κ decreases. It is left to show that in the limit, as $\kappa \to 1$, the audit probability is such that there is no room for bargaining i.e. $T^* = T(z) = T^{nc} = T_F = T_G$. To do so, we recall the optimization of the government:

$$\max_p \mathbf{E}[U_G] = T^{nc}(p) - c(p). \tag{1.24}$$

Clearly, the first part in Eq.(1.24), T^{nc}, will be maximized when there is full compliance i.e. $T^* = T(z)$. From Eq. (1.19), we see that this is the case when $p = \frac{1}{1+\epsilon}$. To prove that $p = \frac{1}{1+\epsilon}$ is a maximum also of the whole equation, it is therefore sufficient to show that the second and negative part of Eq. (1.24) is minimized by $p = \frac{1}{1+\epsilon}$. To do so, we start by considering the optimality condition in the limit i.e. $\kappa \to 1$. Plugging in Eq. (1.22) and rearranging gives

$$c(p) \underset{\kappa\to 1}{=} (T^{*\prime}(p) + (T(Z) - T^*(p) - pT^{*\prime}(p))(1+\xi))p \tag{1.25}$$

Plugging in the conjectured optimum $p = \frac{1}{1+\xi}$ yields that

$$\begin{aligned}&c(\frac{1}{1+\xi}) \underset{\kappa\to 1}{=} (T^{*\prime}(\frac{1}{1+\xi}) + (T(Z) - T^*(\frac{1}{1+\xi}) - \frac{1}{1+\xi}T^{*\prime}(\frac{1}{1+\xi}))(1+\xi))\frac{1}{1+\xi}\\&c(\frac{1}{1+\xi}) \underset{\kappa\to 1}{=} T(Z) - T^*(\frac{1}{1+\xi})\\&c(\frac{1}{1+\xi}) \underset{\kappa\to 1}{=} 0.\end{aligned}$$

where the last equality stems from Eq. (1.19). We have now shown that as $\kappa \to 1$, $p = \frac{1}{1+\xi}$ minimizes the cost function. Hence, $p = \frac{1}{1+\xi}$ will be the optimally chosen audit probability.

It follows immediately that $T^* = T(z) = T^{nc} = T_F = T_G$, such that the bargaining set collapses. $\square$

Chapter 2

Place-based Policies, Structural Change and Female Labor: Evidence from India's Special Economic Zones

Joint with Johannes Gallé, Nadine Riedel and Tobias Seidel.

2.1 Introduction

An increasing number of less-developed countries have implemented Special Economic Zones (SEZs) to foster economic development. According to UNCTAD's World Investment Report (UNCTAD, 2019), the total number of SEZs worldwide increased from 500 in 1995 to about 5,400 in 2018 - with the vast majority of the new zones being located in developing economies. While their specific design can differ, SEZs have in common that they are set up in a clearly defined geographic area where physically present firms have access to lower tax and tariff rates or cost-saving bureaucratic procedures (World Bank, 2008). Their establishment can thus be understood as a place-based policy.

The literature on place-based policies is primarily set in developed economies (e.g. Criscuolo, Martin, Overman, and Reenen, 2019; Grant, 2020; Neumark and Simpson, 2015). Evidence on the effects of SEZs in developing or transitional countries is still scarce (e.g. Duranton and Venables, 2018).[1] This is an important gap in the literature as experiences

[1] An important exception is the growing literature on Chinese SEZs (Alder, Shao, and Zilibotti, 2016; Lu,

with place-based policies in developed countries can hardly be transferred to less-developed economies for various reasons. First, developing countries are characterized by significantly lower institutional quality than their developed country counterparts, which may limit the efficiency of local transfer programs and place-based policies (Becker, Egger, and Ehrlich, 2013; Farole and Moberg, 2014). Second, formal firms operating in developing countries often face substantially higher tax and bureaucratic burdens than firms in developed countries (Gordon and Li, 2009b). Place-based policies that reduce administrative burdens and grant tax exemptions might hence create steeper location incentives. Finally, SEZs in developing countries also differ in purpose and structure from SEZs in the developed world. Among others, they often target exporting firms, for example by offering tariff exemptions for input goods – a feature that is hardly prevalent in developed economies and a clear distinction from other place-based policies that target lagging regions through transfers or tax cuts.

This chapter contributes to the growing literature on place-based policies by evaluating the economic and spatial effects of SEZs that were established after the Special Economic Zones Act in 2005 (SEZ Act, 2005) in India. The policy provided a uniform legal framework for developing and doing business in SEZs and granted firms within SEZs generous tax and tariff exemptions. India was ranked as one of the least business-friendly countries in the *Ease of Doing Business Index* (World Bank, 2005) at the time and the SEZ Act was initiated to improve this situation and create new economic activity. Using a newly compiled data set on the establishment of 147 SEZs between 2005-2013, we show that the SEZ Act led to a substantial increase in non-agricultural employment in SEZ-hosting municipalities.[2] The policy also induced positive employment effects in neighboring locations up to 10km. The rise in local manufacturing and service employment was mirrored by a decline in agricultural work, especially by women. We interpret this pattern as an indication for local structural change from the primary sector towards better-paying jobs in non-agricultural industries. Additional analyses suggest that the establishment of SEZs led to a genuine increase in non-agricultural employment rather than a relocation of jobs.

Methodologically, we identify the effects of SEZs on local employment based on census data in a spatial difference-in-differences (DiD) framework. The estimator compares changes in the economic outcomes in municipalities where SEZs were established with municipalities

Sun, and Wu, 2022; Lu, Wang, and Zhu, 2019; Wang, 2013). Other papers on place-based policies in developed countries include Gobillon, Magnac, and Selod (2012), Busso, Gregory, and Kline (2013), Kline and Moretti (2014) and Ehrlich and Seidel (2018).

[2] We use *municipality* as a collective term for villages and towns in India.

in the same region without SEZs. To this end, we define 5km-distance bins around each SEZ up to a radius of 50km and determine the spatial gradient of the SEZ-effect without parametric restrictions. The main empirical identification concern is that SEZs are not randomly allocated in space, but that their location systematically correlates with the economic trajectories before SEZ-establishment.[3] In the parlance of DiD design, there might be a violation of the common trend assumption.

We address this concern in two ways: First, we examine baseline differences between treated and other municipalities. While such differences are absorbed in DiD designs if time-constant, we also allow outcome trajectories to differ in municipalities' pre-treatment characteristics. In a complementary analysis, we use matching techniques to reduce the imbalance in pre-treatment characteristics of treated and other municipalities in the estimation sample. Second, we run placebo tests, where we draw on census data prior to SEZ establishment to show that employment outcomes did not develop systematically differently between SEZ-hosting municipalities, their neighbors and municipalities in further distance prior to SEZ establishment. This further corroborates the plausibility of the common trend assumption. As census rounds are infrequent, we also show that parallel pre-trends between treated and reference municipalities hold when we proxy economic activity by annual nightlight data, which allows for a more detailed picture in the immediate lead-up to SEZ establishment.

Our empirical analysis builds on a novel data set that combines census data with georeferenced data on SEZs for the period 1998-2013. In the main model, we draw on employment information from the 2005 and 2013 waves of the Economic Census and on population information from the 2001 and 2011 waves of the Population Census, which cover the universe of firms and households in India, respectively. We, moreover, identify the location of all SEZs and the date when they went into operation from newspaper articles, official statistics by the Ministry of Commerce and Industry as well as from minutes of the Central Board of Approval and match them with their hosting municipality using the India Village-Level Geospatial Socio-Economic Data Set (Meiyappan, Roy, Soliman, Li, Mondal, Wang, and Jain, 2018). Having identified the SEZ-hosting municipality allows us to add rich granular census information like municipal employment by sector and gender and the number of firms (Asher, Lunt, Matsuura, and Novosad, 2021). Our final sample includes

[3] Note that we restrict our data to municipalities within a 50km radius of an SEZ. Even if SEZ developers condition the location of SEZs on regional economic outcomes, it is – e.g. due to land restrictions (Levien, 2012; Parwez and Sen, 2016) – hardly feasible to precisely target SEZs to specific subareas. This dampens concerns that the location of SEZs across sample municipalities correlates with outcome trends in our data.

almost 50K Indian municipalities with a total population of 146M people in 2011.

Our baseline results uncover a sizable effect of SEZ-establishments on local employment in manufacturing and services. In SEZ-hosting municipalities, employment growth over the 8-year time frame between 2005 and 2013 is estimated to exceed employment growth in reference locations – defined as municipalities in the 20-25km distance bin – by 52 percentage points (*pp*). To put this sizable effect into perspective, note that India experienced high overall employment growth in this time period and municipalities are mostly small entities with average non-farm employment of 290 workers (median of 41). Our findings indicate that the policy also contributed to local economic development beyond the boundaries of SEZ-hosting municipalities up to a distance of 10km. In the first distance bin around SEZs ($<$ 5km), non-agricultural employment growth is 22*pp* higher than in the reference location after SEZ-establishment; in the second distance bin (5-10km), it is 16*pp* higher. For municipalities in a distance of 10-50km from an SEZ, we find no significant difference in employment trajectories relative to the reference location.

In additional analyses, we show that the SEZ policy reduced the number of workers in the agricultural sector: SEZ municipalities experienced a 17*pp* lower agricultural employment growth than reference locations. This pattern suggests that the SEZ Act contributed to a local transition from an agrarian-based towards an industrial and service economy. This transition is widely considered to be one of India's main development challenges (Sud, 2014) as productivity in the agricultural sector is low: the 50-60% of Indian workers employed in agriculture contribute only 18% to GDP (World Bank, 2023a,b). The drop in agricultural work is centered around marginal employment (i.e. employment of 183 days per year or less) and hence around the least-paying jobs in the agricultural sector. This further supports the interpretation that SEZs created better employment opportunities for local workers. This finding connects well with previous research that has emphasized the importance of sectoral shifts from agriculture to more productive industries as a key driver of economic development (Eichengreen and Gupta, 2011; Gollin, Lagakos, and Waugh, 2014; McMillan, Rodrik, Dani, and Verduzco-Gallo, 2014). A back-of-the-envelope calculation suggests that the decrease in agricultural employment amounts to around one third of the increase in non-agricultural employment.

Moreover, we find that the decline in agricultural employment was in particular driven by female workers. Men experienced a weaker and statistically insignificant drop in agricultural work. One possible explanation for the latter finding is that men own most agricultural land in India (Agarwal, Anthwal, and Mahesh. Malvika, 2021), which might limit their responsiveness to alternative job opportunities. Female non-agricultural employment in SEZ municipalities went up markedly after SEZ establishment: the growth rate of female manufacturing workers in SEZ-hosting municipalities exceeded that in reference municipalities by 55*pp*. The policy hence contributed to better employment opportunities for women in the secondary sector. Female employment in services increased only marginally (and insignificantly) in turn, while male employment went up to a similar extent in manufacturing and services.

The finding on gender effects resonates well with observers' expectation that SEZ policies would generate new and better jobs for women (Bacchetta, Ernst, and Bustamante, 2009; Rama, 2003; World Bank, 2011a). And it connects to recent literature that has documented rising shares of female employment caused by free-trade policies in many countries (Bussmann, 2009; Ozler, 2000). We also consider SEZs' effect on female employment to be of particular relevance as women in India – similar to other less developed countries – are a vulnerable group in the labor market: Gender discrimination is a prevalent and long-standing phenomenon, and unemployment rates among women are significantly higher than among men (Klasen and Pieters, 2015; Srivastava and Srivastava, 2015).

Finally, we offer two further insights. The first concerns the impact of SEZ policies on the formal and informal sector. Our empirical analysis draws on census data that allows us to observe the universe of Indian manufacturing and service firms and to proxy for formal and informal firms. We show that smaller, informal entities – which, despite accommodating over 90% of the Indian workforce, are ignored in many previous studies – also respond strongly to SEZ establishment and contribute significantly to the aggregate creation of non-agricultural jobs by SEZs. Ignoring these firms hence underestimates the local employment impact of SEZs and other place-based policies. In further analyses, we exploit that India hosts a variety of SEZs, which differ in two key dimensions: there are zones that are developed by private and public developers respectively and zones with different industry denominations. Our analysis shows that SEZs of different types exert broadly comparable effects on overall local employment (while the industry composition of the new employment can naturally differ). Finally, we combine our estimates with official statistics on foregone tax revenues. This tentatively suggests that the SEZ scheme supported job creation at relatively low fiscal costs.

Beyond the referenced literature so far, our study relates closely to research on the spatial economic effects of place-based policies. Most existing work is set in developed countries (Busso, Gregory, and Kline, 2013; Ehrlich and Seidel, 2018; Gobillon, Magnac, and Selod, 2012; Neumark and Kolko, 2010; Neumark and Simpson, 2015). Evidence on the effects of SEZs in less-developed countries is scarce – with studies on SEZs in China being the notable exception (e.g. Lu, Sun, and Wu, 2022; Lu, Wang, and Zhu, 2019; Wang, 2013). While offering valuable insights, it is unclear whether these Chinese experiences translate to other countries. Chinese SEZs were established during a time when China transitioned from a central planning to a market economy. The country's institutional context at that time differed from many other less developed economies. In particular, the high degree of government intervention in the economy stands out (Bosworth and Collins, 2008; Farole and Moberg, 2014). As SEZs were granted more free market-oriented economic policies and flexible governmental measures compared to the planned economy elsewhere, incentives for firms to locate inside SEZs might have been steeper than in other countries.[4] Observers have thus raised concerns that results from existing studies may not be externally valid for other countries (The Economist, 2015). The World Bank Group writes: "Extracting wide-ranging policy implications from [...] [such] analysis remains risky" (World Bank, 2017). This calls for evidence for other countries. We fill this gap and augment the evidence on China by documenting a sizable SEZ-effect on local employment in India, another leading emerging economy, which has firmly embraced SEZ policy.[5]

A comprehensive overview of the history and development of the Indian SEZ experience is offered by Mukherjee, Pal, Deb, Ray, and Goyal (2016). An earlier working paper by Hyun and Ravi (2018) mostly relies on nightlight intensity and sample survey data to assess the effect of SEZ establishment on economic development within broad SEZ-hosting districts in India.[6] We offer the first granular analysis based on census data and deliver novel insights on the channels through which SEZs shape local economic outcomes as well

[4] Other SEZ-related benefits, in part, overlap: Investors in China and India are e.g. granted special tax cuts when locating in SEZ areas. One particularity, in turn, of the Indian scheme is that SEZs can be developed bottom-up by private investors. Our evidence suggests that these private SEZs induce significant local employment effects.

[5] Our estimated employment effects for directly affected municipalities are quantitatively broadly comparable to those measured in China (e.g., Lu, Wang, and Zhu, 2019; Wang, 2013). Other studies on place-based policies in China include Koster, Cheng, Gerritse, and Oort (2019) who find 10-15% higher firm productivity following the opening of science parks in Shenzhen. Chen, Lu, Timmins, and Xiang (2019) document a decline in TFP by 6.5% due to closures of development zones. Jia, Ma, Qin, and Wang (2020) explore China's Great Western Development Programme finding no evidence for employment or wage effects, but higher local GDP through physical investment. Also note that prior research for China finds little evidence for SEZ spillovers to neighboring jurisdictions (Lu, Wang, and Zhu, 2019), while we find statistically significant and quantitatively relevant spillover effects in a more granular spatial setting in India.

[6] Alkon (2018) focuses mainly on infrastructure effects of SEZs at the sub-district level. Another paper by Görg and Mulyukova (2024) uses a sample of large Indian firms to study the effect of SEZs on exporting behavior and factor productivity.

as the geographical and social dispersion of economic growth. Previous empirical work on the economic consequences of other regional and local public policies in India has studied distinctively different programs, namely preferential tax policies for industrially backward districts (Abeberese, Chaurey, and Menon, 2024; Hasan, Jiang, and Rafols, 2021), state-level tax incentives (Chaurey, 2017; Shenoy, 2018), rural road construction programs (Asher and Novosad, 2020a). The latter policies aimed at reducing spatial disparities by targeting lagging regions or improving the infrastructure network. SEZ policies differ from these interventions in terms of objectives and instruments. They aim to promote exports and investments, mostly in non-lagging regions (SEZs are often located close to urban areas or transport infrastructure like ports), target internationally active firms (by granting tax and tariff benefits to exporters and importers) and contrary to other place-based policies allow for a bottom-up approach, where private developers can develop and run SEZs. The effectiveness of these policies and agglomeration dynamics may differ, rendering it difficult to apply lessons from prior research to the SEZ policy.

To the best of our knowledge, we are also the first to empirically link SEZ establishment to sectoral shifts from agriculture to manufacturing and services. This adds to the literature on structural change and economic growth (Gollin, Lagakos, and Waugh, 2014; Kline and Moretti, 2014; Laitner, 2000; McMillan, Rodrik, Dani, and Verduzco-Gallo, 2014). For India, Eichengreen and Gupta (2011) identify the sectoral shift from agriculture to services as a key driver of economic growth; Blakeslee, Chaurey, Fishman, and Malik (2022) study the effects of a land-rezoning program in Karnataka on local sectoral shifts. Previous work in other countries has mostly focused on the role of trade liberalization and international integration for structural change, see e.g. Uy, Yi, and Zhang (2013) for Korea and McCaig and Pavcnik (2013) for Vietnam.

The results on changes in female employment in agriculture, manufacturing and services further inform the extensive literature that has documented the positive effects of female labor force participation and empowerment for economic development as summarized, for example, by Duflo (2012) and World Bank (2012) in general and by Das, Jain-Chandra, Kochhar, and Kumar (2015) for India. According to statistics by the International Labour Organization, India features a comparably low female labor force participation rate of around 25%. Policies, which create labor market opportunities for women, may hence come with high socio-economic returns. Our paper contributes to this line of research by connecting novel gender-specific labor market effects with the place-based policy literature.

The remainder of the paper is organized as follows. Section 2.2 describes the institutional background. Section 2.3 presents the empirical methodology. Section 2.4 introduces the construction of our data set and descriptive statistics. We discuss our findings in Sections 2.5 and 2.6. Section 2.7 offers a back-of-the-envelope calculation on the cost effectiveness of the policy before we conclude.

2.2 Institutional background

In the 1960s, India became one of the first countries to establish export-processing zones which were later relabeled as SEZs in the early 2000s. But for long, SEZs were rare in the country. Between the 1960s and the 1990s, only seven SEZs were established by the central government. This changed drastically when the Indian government implemented the Special Economic Zones Act in 2005, allowing for private investments in SEZs and a much more flexible environment than the precedent framework in which all zones were owned and managed exclusively by the central government. Until 2020, the number of operational SEZs, i.e. zones with at least one active company, increased markedly to 240 of which more than 90% were established under the SEZ Act (see Figure 2.2.1).

Figure 2.2.1: Operational SEZs in India

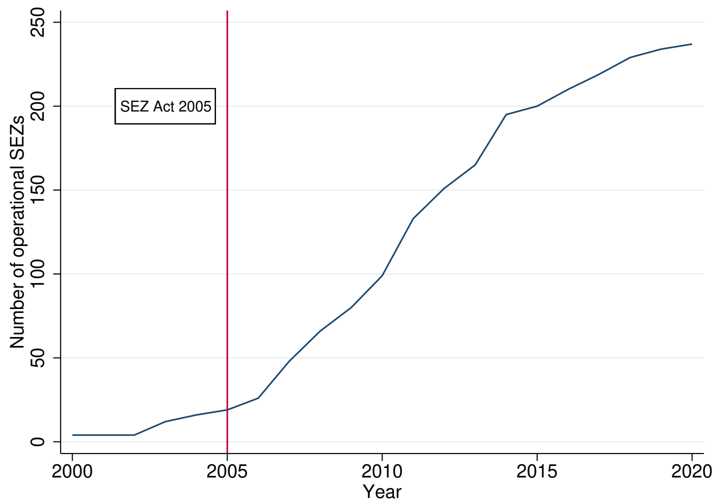

Notes: This figure plots the cumulative sum of operational SEZs in India by year. SEZs are defined as being operational as soon as one firm commenced with its production. The individual SEZ data are obtained from the Indian Ministry of Commerce and Industry. The date of operation is sourced from newspaper articles and administrative records.

Against the background of India's economy being highly regulated and poorly integrated into the global economy (World Bank, 2005), the main goals of the SEZ Act were to (i) generate additional economic activity, (ii) promote exports of goods and services, (iii) promote investment from domestic and foreign sources, (iv) create employment opportunities, and (v) develop local infrastructure facilities (SEZ Act, 2005).[7] To achieve these goals, the SEZ Act provided a uniform legal framework for developing and doing business in these specially designated areas. Firms in SEZs, moreover, enjoyed various administrative and fiscal benefits. On the administrative side, there was so-called "single-window clearance", that is, all approvals were issued by a single authority.[8] Businesses in SEZs, moreover, received a 100% income-tax exemption on export income for the first five years of operation, which reduced to a 50% exemption for the following five years. Thereafter, SEZ firms received a tax benefit of 50% on reinvested profits for a final period of five years. SEZ business units were, furthermore, exempted from sales and service taxes and, until 2012, from the Minimum Alternate Tax (MAT), a minimum tax on profits of 18.5%. SEZ firms also benefited from duty-free imports and domestic procurement of goods and services. Note that SEZs were treated as being outside of the domestic tariff area (DTA), so that goods that were produced in the SEZ and sold into the DTA were considered as imports to the Indian market. In consequence, companies in the DTA had to pay import tariffs if they purchased goods from a SEZ company. In turn, goods and services supplied by DTA companies to SEZ units were considered as exports from the DTA and exempted from any taxes and tariffs. Hence, the flow of goods from DTA into SEZs was subject to no taxes or tariffs, but not vice versa.

Applications for establishing a SEZ were assessed by the Central Board of Approval. One of the main criteria for an approval by the board was that SEZ developers were in the rightful possession of sufficiently large parcels of land depending on the industry denomination. For example, multi-product zones required a minimum contiguous area of 10 square kilometers while sector-specific zones such as IT zones required only 0.1 square kilometers. After the formal approval by the board, the proposal to develop the SEZ was recommended for notification to the Ministry of Industry and Commerce, which officially declared the designated area as an SEZ area.

[7] Mukherjee, Pal, Deb, Ray, and Goyal (2016) conducted a survey of 145 businesses from 32 SEZs and asked for their motivation to locate in SEZs. Duty-free imports, ease of exports, tax holidays, single window clearance and ease of business were mentioned as the most important reasons.

[8] According to the World Bank's Enterprise survey in 2022, 12.1% of IT-firms mentioned "business licensing and permits" to be their biggest obstacle to doing business.

2.3 Empirical approach

To identify the causal economic impact of SEZs across space, we draw on two economic census waves (2005 and 2013) and implement a difference-in-differences-style analysis comparing changes in outcome variables between municipalities that host an SEZ and municipalities in the same region without an SEZ before and after the treatment, i.e. the start of the SEZ Act in 2005.[9] To this end, we group municipalities in 5km-distance rings around their closest SEZ up to 50km.[10] This allows us to non-parametrically study the spatial effects of the policy. Municipalities outside of the 50km radius around an SEZ are dropped from the analysis. The main analysis relies on a spatial difference-in-differences model of the following form:

$$\ln(y_{it}) = \sum_{d=0,d\neq 5}^{10} \beta_d (D_{[d_i=d]} \times POST_t) + \boldsymbol{\eta}' (\mathbf{X}_i \times POST_t) + POST_t + \alpha_i + \varepsilon_{it}, \tag{2.1}$$

where y_{it} represents outcomes like employment or the number of firms in municipality i in year t. $D_{[d_i=d]}$ indicates whether a municipality i is in distance bin d to an operational SEZ in the post-treatment year. $d_i = 0$ indicates SEZ-hosting municipalities, $d_i = 1$ SEZ-neighboring municipalities within a 5km-distance to the SEZ, $d_i = 2$ municipalities in a 5-10km distance etc. up to 50km. Distance bin $d = 5$ (distance of 20-25km) is omitted and serves as the reference category. We interact the distance dummy with a post-reform dummy $POST_t$. The model further includes municipality fixed-effects, α_i, and additional control factors, $\mathbf{X}_i \times POST_t$, which are specified in further detail below. ϵ_{it} is the error term. The β_ds are the parameters of interest capturing differences in outcome trends in municipalities in distance bin d relative to municipalities in the reference category. In the baseline specification, we cluster standard errors at the district level to account for spatial correlation. In additional specifications, we cluster at the level of the "closest SEZ groups"

[9] The prior literature on place-based policies has also pursued identification strategies, where locations that are targeted by a given policy are compared with locations that were considered but not finally picked for treatment (see e.g. Greenstone, Hornbeck, and Moretti (2010)). Approaches along these lines are, unfortunately, not feasible in our setting, as documentations from meetings of the SEZ Board of Approval (http://sezindia.nic.in/cms/boa-minutes.php) show that the vast majority of SEZ applications are approved. Another strategy that has been pursued in prior literature is to compare treated municipalities with municipalities that are selected for treatment in the future. Again, this is not viable in our setting as a substantial part of these potential control SEZs are located within close geographic proximity of SEZs that became operational by 2013.

[10] There are some municipalities within these 50km radii of our SEZs that are also within a 50km radius to an SEZ established before the SEZ Act in 2005. Excluding them from the sample does not change our estimation results.

comprising all municipalities whose d_i is determined by the same SEZ and apply Conley (1999) standard errors.

Note that the concentric ring analysis allows us to capture the spatial effect of the policy. The choice of the reference category is arbitrary and anchors the interpretation of the coefficient estimates for β_d as the effect of the SEZ on the *relative* economic development of municipalities in radius d to the reference municipalities. Prior research has shown that the spillover effects of place-based policies tend to be very local (Ehrlich and Seidel, 2018; Einiö and Overman, 2020). Our results suggest that the same holds true for SEZs in India. If we were willing to assume that the reference municipalities in 20-25km distance are unaffected by the policy, the β_ds can be interpreted as the effect of the SEZ policy on the treated municipalities.

The main threat to our empirical identification strategy and to obtaining unbiased estimates for β_d is the violation of the conditional mean independence assumption. If SEZ developers systematically place SEZs in areas whose outcome trends differ from other municipalities, conditional mean independence is violated – or in the parlance of DiD design – there is a violation of the common trend assumption.[11]

We address this concern in two ways. First, we explore differences in the baseline characteristics of treated and other municipalities. While differences in municipal baseline characteristics are absorbed by α_i if time-constant, these characteristics might also correlate with changes in economic outcomes.[12] We thus control for municipal baseline characteristics interacted with the post-treatment dummy, $\mathbf{X}_i \times POST_t$. The vector $\mathbf{X}_i$ models differences in municipality size (dummy variables for the quartiles of the population distribution), employment structure (a dummy variable indicating that there are formal firms in the locality), industry composition (dummy variables for the dominant industry measured by employment share), distance to key infrastructure (airports, ports, highways, railroad, power plants) and to the next urban center. Under the assumption, that $\mathbf{X}_i$ is randomly drawn from all factors influencing SEZ locations, obtaining similar estimates for β_d with and without the control variables $\mathbf{X}_i \times POST_t$ mitigates the concern that locational differences cause a bias (Altonji, Elder, and Taber, 2005).

[11] In particular, unbiasedness relies on the following assumption:

$$\mathbf{E}[\varepsilon_{it}|D_{[d_i=d]} \times POST_t, \mathbf{X}_i \times POST_t, POST_t, \alpha_i] = \mathbf{E}[\varepsilon_{it}|\mathbf{X}_i \times POST_t, POST_t, \alpha_i] \tag{2.2}$$

That is, conditional on α_i, $POST_t$ and $\mathbf{X}_i \times POST_t$, the regressor of interest $D_{[d_i=d]} \times POST_t$ and the error term ε_{it} are mean-independent.

[12] One potential example are differences in the proximity of municipalities to the Golden Quadrilateral National Highway in India whose construction was completed in 2013.

In a complementary line of inquiry, we turn to matching techniques to reduce imbalances in the characteristics of treated and other municipalities. We employ coarsened exact matching (CEM), that is we temporarily coarsen the data based on the observed $\mathbf{X}_i$ using automated binning strategies and define unique observations of the coarsened data, each of which is a stratum. Treated and control municipalities are then exactly matched on these strata. Observations whose strata do not contain at least one treated and one control observation are dropped and weights are used to compensate for the different strata sizes (Iacus, King, and Porro, 2012). Importantly, and contrary to many other matching strategies, coarsened exact matching does not only account for imbalances in means, but also for imbalances in higher moments and interactions (Blackwell, Iacus, King, and Porro, 2009; Iacus, King, and Porro, 2012).

The second strategy to further corroborate the common-trend assumption is twofold. First, we use the Census waves in 1998 and 2005 to run placebo regressions for the pre-treatment period. If running the spatial difference-in-differences model in Eq. (2.1) on data prior to SEZ introduction reveals no differential outcome trends between treated and control units, this supports the common-trend assumption. Second, as census data are available only infrequently, we augment our analysis by annual nightlight data that have been shown to serve as a good proxy for economic activity and are widely used in the literature (Henderson, Storeygard, and Weil, 2012). This allows us to test for differential outcome pre-trends between SEZ-treated municipalities and reference municipalities in the years directly leading up to treatment.

While our main analysis follows a classic two-by-two difference-in-differences approach (tracking economic outcomes between two censuses), the nightlight data allow us to model the staggered implementation of SEZs in an event study analysis. To obtain unbiased estimates in this setting in the presence of heterogeneous and dynamic treatment effects, we rely on the estimator proposed by Callaway and Sant'Anna (2021), which – similar to other estimators in the literature – ensures that already-treated units are not used as a control group for later-treated units.[13] The model compares the evolution of employment outcomes of municipalities treated by SEZs with reference municipalities (in a distance of 20-25km). It reads:

$$\ln(nl_{it}) = \sum_{k=-5, k\neq -1}^{5} \theta_k \mathbf{1}[t - T_i = k] + \gamma_t + \alpha_i + \epsilon_{it}, \tag{2.3}$$

[13]Other estimators, e.g. Sun and Abraham (2021), yield similar results to the ones presented below.

where nl_{it} denotes the average nightlight intensity in municipality i in year t and T_i denotes the year in which the SEZ related to municipality i became notified. α_i and γ_t denote municipality and year fixed effects, respectively. The θ_ks can therefore be interpreted as the dynamic treatment effects (in relative time k) of SEZs on municipal nightlight intensity.

2.4 Data

Data on SEZs. We compiled information on all 147 Indian SEZs that were established under the SEZ Act and became operational until 2013 from various sources. Data on the name of the SEZ, whether the SEZ was privately or publicly developed, its location, size, industry type and date of notification are readily available from the Ministry of Commerce and Industry.[14] We georeference each SEZ at the municipality-level or, if available, even at its exact location. We verify our strategy by comparing our SEZ coordinates with a subsample of officially georeferenced SEZs that is accessible at the development commissioner's website of the Visakhapatnam SEZ.

A key variable for our empirical analysis, the start of operation of a zone, was not directly accessible and had to be hand-collected from newspaper articles, official statistics by the Ministry of Commerce and Industry as well as from minutes of the Central Board of Approval. We define the date of operation as the earliest date available, where we find at least one firm in the SEZ that went into operation. Figure 2.4.1 illustrates the geographical location of SEZs.

Link to municipal data. Using GIS techniques, we spatially join the georeferenced SEZ data with the India Village-Level Geospatial Socio-Economic Data Set (Meiyappan, Roy, Soliman, Li, Mondal, Wang, and Jain, 2018), which provides the administrative boundaries of every municipality in India based on the Population Census of 2001. To identify SEZ-hosting municipalities and municipalities in close proximity to SEZs, we approximate the area of the SEZ based on the geo-coordinates and information on the SEZ's area which by the SEZ Act is required to be contiguous (SEZ Act, 2005). As information on precise SEZ boundaries is unavailable, we assume SEZs to be circular. Based on the total area, we then calculate the radius of the zone and consider all municipalities that fall within this radius as SEZ-hosting municipalities (see Appendix 2.A for details). The geo-referencing further allows us to compute distances from sea ports, airports, railway networks, highways, cities

[14] http://sezindia.nic.in/index.php.

Figure 2.4.1: Geographical distribution of operational SEZs

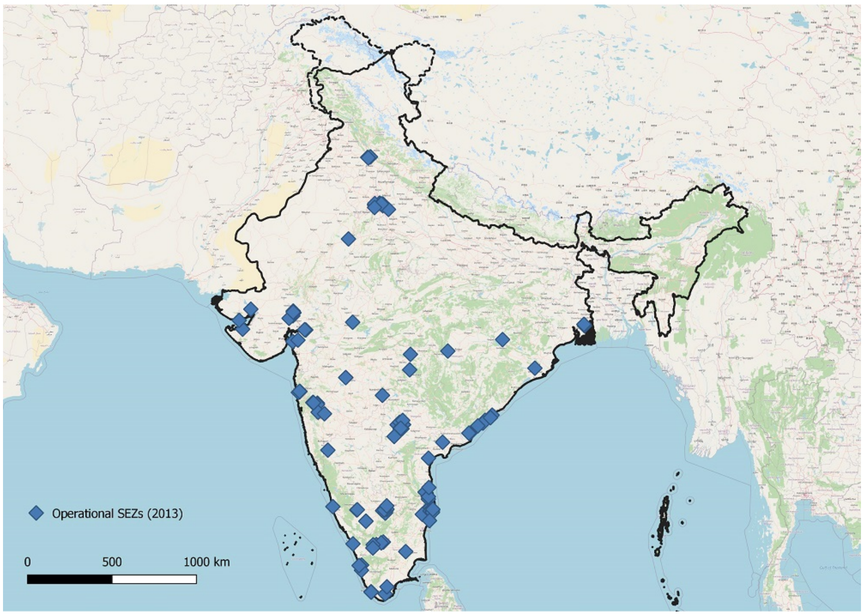

Notes: This figure plots the location of all SEZs in India that were established under the SEZ Act 2005 and became operational until 2013.

or power plants that we will use as control variables in the empirical analysis.[15]

Data on outcome variables. Having information on the start of operation of each SEZ and knowing their hosting municipalities, we finally use both the *Economic Census* and the *Population Census* to add economic variables like employment, population and the number of firms. The *Economic Census* contains the population of all non-agricultural (i.e. manufacturing and service) firms in India including the informal sector. We can draw on three repeated cross-sections of data for the years 1998, 2005 and 2013. We link municipalities across the three Economic Census waves by using the time-consistent municipality identifiers provided by the Socioeconomic High-resolution Rural-Urban Geographic Platform for India (Asher, Lunt, Matsuura, and Novosad, 2021, SHRUG). For every non-agricultural firm in India, the Economic Census contains information on employment (total and separate by gender), a firm's industry code and its host municipality. We disregard public administration employment and employment in international organizations. The Economic

[15]We retrieved data on the geo-coordinates of these infrastructure facilities as follows: Airports from the WFP SDI-T Logistics Database (https://data.humdata.org/dataset/global-logistics), Ports from the World Port Index (https://msi.nga.mil/Publications/WPI), Power Plants from the Global Power Plant database (https://datasets.wri.org/dataset/globalpowerplantdatabase), railways and roads from the Digital Chart of the World (https://www.soest.hawaii.edu/pwessel/dcw). Last accessed: June 2nd, 2024.

Census for 2013 lists 58.5 million firms employing 131.3 million workers. We collapse each Economic Census round to the municipality level and calculate the municipalities' number of firms, total employment, employment by gender and by industry as well as employment for small and large firms, defined as firms with less than 10 employees and firms with 10 employees or more, respectively.[16] The latter distinction is of particular importance as firm size in India discontinuously impacts firm formality. While firms of all sizes may decide to operate outside the formal sector, all firms with less than 10 employees are by official statistics classified as informal, reflecting that they are subject to a light regulatory burden under Indian law (NCEUS, 2009). For example, they do not need to register with official statistics, are exempted from social security taxes and subject to light bureaucratic procedures (Amirapu and Gechter, 2020; Mehrotra, 2019). We will show below that these small, informal firms employ the majority of Indian workers - ignoring them in empirical analyses hence implies that aggregate employment effects of place-based policies can be severely underestimated. We thus consider it to be a decisive advantage of our census data that it provides a complete picture of economic activity, accounting for formal and informal firms as well as for manufacturing and service entities.

We further complement the data with three waves of the *Population Census* containing a repeated cross-section of data for the years 1991, 2001 and 2011. The data contain information on the total population, literacy and infrastructure facilities such as number of schools, road access or electricity for every municipality in India. Most importantly, the Population Census contains information on persons working as cultivators or agricultural laborers, which are not covered by the Economic Census. As the last wave of the Population Census was 2011, we restrict the sample to municipalities in 50km radii of SEZs which became operational up to 2011 for analyses based on Population Census variables. Finally, we use annual information on average nightlight intensity matched to the municipality level (Asher, Lunt, Matsuura, and Novosad, 2021; NOAA, 2013).

Descriptive statistics. Figure 2.4.2 illustrates that the majority of SEZ-hosting municipalities are relatively small as measured by their inhabitants in 2001. There are a few SEZs in India's leading cities – defined as cities with more than 500K inhabitants in 2001

[16] We use the concordance tables provided by the Ministry of Statistics and Programme Implementation to harmonize industry codes across time. While the Economic Census of 2013 uses the National Industry Classification (NIC) of 2008, the Economic Censuses of 2005 and 1998 use the NIC codes of 2004 and 1987, respectively. We match the three-digit NIC-04 Codes to three-digit NIC-08 codes and aggregate them to one digit NIC-08 codes for our analysis. In cases of industry splits across industries, we assign the industry code, that has a higher employment share according to the Economic Census of 2013. Hence, while the harmonization of industry codes is not entirely time consistent, note that most of the industry splits are between NIC-04 and NIC-08 are within the same one-digit industry.

Figure 2.4.2: Size distribution of SEZ-municipalities

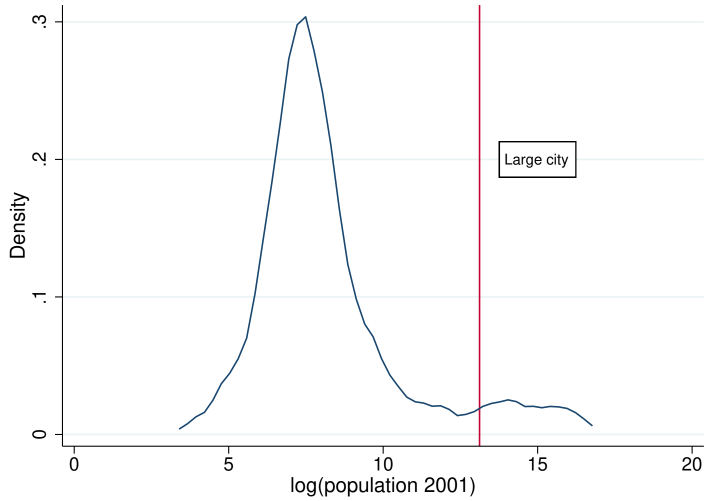

Notes: Large cities are defined as $> 500K$ population.

– which we take out of our base analysis as effects related to SEZ establishment in these metropolitan areas are difficult to detect in the data. Since it concerns few observations, this sample restriction is not decisive for any of the results presented in this paper.[17]

The final sample comprises 49,669 municipalities with a total population of 146M people according to the latest Population Census in 2011. As shown in Appendix 2.A.2, the average municipality employs 290 non-agricultural employees with a median of 41 workers and accommodates 3,061 residents. On average, there are 70 (220) female (male) non-agricultural workers per municipality and 189 (330) female (male) agricultural workers. Small informal firms with less than 10 workers account for about two thirds of average municipal employment.

With respect to ownership, 77% of the SEZs in our 2005-2013 sample were developed by private companies versus 23% by public bodies. In terms of industry denomination, 57% are IT zones, followed by engineering (12%), pharmaceutical (9%) and multi-product zones (9%). The average SEZ covers 1.76 square kilometers, but the size varies systematically by industry denomination. IT-zones, on average, cover 0.25 square kilometers, multi-product SEZs 14.02 square kilometers.[18]

[17] We show in Appendix 2.B that the estimated SEZ effects do not change when large cities are included.

[18] Based on a self-compiled firm-level dataset that we describe in Appendix 2.B.3, we further observe that about 65% of SEZ firms are privately owned by Indians and about 25% belong to a company group. Private firms owned by foreigners only account for about 10% of SEZ firms.

Table 2.4.1: Pre-treatment location characteristics

	Mean values and standard deviations (in brackets)											
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(6)-(1)
	0km	0-5km	5-10km	10-15km	15-20km	20-25km	25-30km	30-35km	35-40km	40-45km	45-50km	Difference
log distance to city (km)	3.808	3.645	3.712	3.658	3.749	3.806	3.861	3.952	4.035	4.108	4.159	-0.002
	(1.011)	(0.909)	(0.874)	(0.808)	(0.838)	(0.763)	(0.699)	(0.651)	(0.594)	(0.549)	(0.536)	(0.066)
log distance to power plant (km)	3.718	3.652	3.699	3.531	3.733	3.753	3.816	3.896	3.925	3.929	3.939	0.035
	(0.750)	(0.774)	(0.784)	(0.851)	(0.769)	(0.798)	(0.767)	(0.752)	(0.753)	(0.776)	(0.765)	(0.068)
log distance to airport (km)	4.519	4.395	4.431	4.370	4.607	4.618	4.701	4.783	4.796	4.825	4.855	0.099
	(1.350)	(1.264)	(1.193)	(1.052)	(1.053)	(1.016)	(0.970)	(0.917)	(0.881)	(0.838)	(0.810)	(0.088)
log distance to port (km)	4.459	4.756	4.842	4.753	5.046	4.960	5.025	5.102	5.087	5.143	5.126	0.501
	(1.327)	(1.345)	(1.273)	(1.277)	(1.175)	(1.132)	(1.138)	(1.117)	(1.080)	(1.094)	(1.058)	(0.097)
log distance to railway (km)	1.735	2.002	2.073	2.020	2.112	2.156	2.220	2.316	2.382	2.461	2.467	0.421
	(1.153)	(1.153)	(1.085)	(1.054)	(1.068)	(1.084)	(1.116)	(1.116)	(1.092)	(1.114)	(1.148)	(0.093)
log distance to highway (km)	1.941	2.151	2.312	2.419	2.553	2.618	2.733	2.863	2.954	3.014	3.063	0.677
	(1.290)	(1.228)	(1.150)	(1.062)	(1.118)	(1.112)	(1.116)	(1.062)	(1.047)	(1.040)	(1.055)	(0.096)
log population in 2001	7.643	7.428	7.204	7.257	7.115	7.092	7.088	7.044	6.997	6.980	6.949	-0.551
	(1.432)	(1.220)	(1.098)	(1.063)	(1.098)	(1.071)	(1.034)	(1.049)	(1.079)	(1.080)	(1.079)	(0.093)
Formal employment share in 2005	0.222	0.135	0.110	0.102	0.105	0.0995	0.0873	0.0745	0.0735	0.0715	0.0687	-0.122
	(0.310)	(0.235)	(0.214)	(0.204)	(0.215)	(0.211)	(0.192)	(0.178)	(0.176)	(0.174)	(0.167)	(0.018)

Notes: This table reports the mean values and their standard deviations for municipalities in the respective distance bins relative to SEZs. The last column shows the differences between the control group (column 6) and SEZ-hosting municipalities (column 1). *Distance* measures the distance in kilometers to the closest respective amenity. *City* denotes municipalities with a population of more than 500K. *Formal employment share in 2005* denotes the share of of formal employment (i.e. in firms with more than 10 employees) in total municipal employment. Standard deviations in brackets.

As we compare municipalities across space, Table 2.4.1 provides an overview of locational characteristics by distance bin, mostly in logs as they enter our estimation. The last column shows differences between the reference locations and SEZ-hosting municipalities. There are no significant differences between the two groups with respect to proximity to large cities, airports and power plants, but SEZ-hosting municipalities tend to be closer to other infrastructure facilities such as railways or highways compared to reference locations. Further, municipalities with an SEZ tend to be larger in terms of population and are characterized by a higher formal employment share. We ensure that these differences in locational characteristics are not driving our estimation results by interacting all depicted covariates with the post-treatment dummy according to Eq. (2.1).

2.5 Baseline results

In this section, we will present estimation results for the models specified in Section 2.3. In the following, we will show that SEZ establishment increased local manufacturing and service employment in India (Section 2.5.1), illustrate that SEZs generated genuinely new jobs, rather than inducing relocation of jobs in space (Section 2.5.2) and present evidence that relates SEZs to structural change (Section 2.5.3). In the appendix, we furthermore document that SEZ establishment did not improve local infrastructure provision (which was another goal of the SEZ Act as outlined in Section 2.2).

2.5.1 Employment effects

Using the log of municipalities' manufacturing and service employment as the dependent variable, Figure 2.5.1 (a) presents estimation results of the spatial model in Eq. (2.1) by plotting the coefficients $\hat{\beta}_d$ with the corresponding 95%-confidence intervals for all distance bins. We find a sharp difference in the employment growth of SEZ-hosting municipalities and reference locations between 2005 (the year of the SEZ Act) and 2013. SEZ-hosting municipalities and direct neighbors significantly gained employment relative to municipalities in further distance to the SEZ, suggesting that SEZs had a strong impact on local economic activity. Quantitatively, the point estimate suggests that SEZ establishment increased employment growth in SEZ-hosting municipalities by $52pp$ $(= (e^{0.418} - 1) \times 100)$ relative to the reference municipalities.[19] Employment growth in municipalities in the <5km distance bin and the 5-10km distance bin increased by $22pp$ and $16pp$, respectively, indicating substantial positive spillovers to adjacent regions. For more distant municipalities, the estimates for β_d turn out to be small and statistically insignificant, suggesting that employment changes between municipalities in further distance to the SEZ did not differ systematically. The magnitude of the estimated employment response is fairly large, but not implausible given the high general employment growth in India between the 2005 and 2013 Census waves and the relatively small sizes of our sample jurisdictions. The average SEZ municipality in the sample hosts only 3,139 non-agricultural employees prior to treatment, so the estimated relative effect translates into moderate absolute values. We show in Appendix 2.B.1 that these results are robust to using alternative distance bin classifications, alternative standard error clustering, including municipalities up to a distance of 200km, including large cities in the sample, re-estimating Eq. (2.1) without the control vector $\mathbf{X}_i \times POST_t$ or applying coarsened exact matching to reduce the imbalance between treated and other municipalities. We further show that the Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA) – a public employment program that was enacted in 2005 – does not act as a confounder in the analysis.

Figure 2.5.1 (b), moreover, presents estimates of a placebo test that reruns the spatial difference-in-differences analysis for the pre-treatment period 1998-2005. Evidently, all estimated coefficients are close to zero and statistically insignificant which supports the common trend assumption of our spatial difference-in-differences design.

[19] As Eq.(2.1) includes municipality fixed effects, β_d is the difference between changes over time in $\ln(y)$ for municipalities in distance band d relative to changes over time in $\ln(y)$ for municipalities in the reference distance band. Thus, it captures percentage point differences in growth rates of y.

Figure 2.5.1: Spatial difference-in-differences model

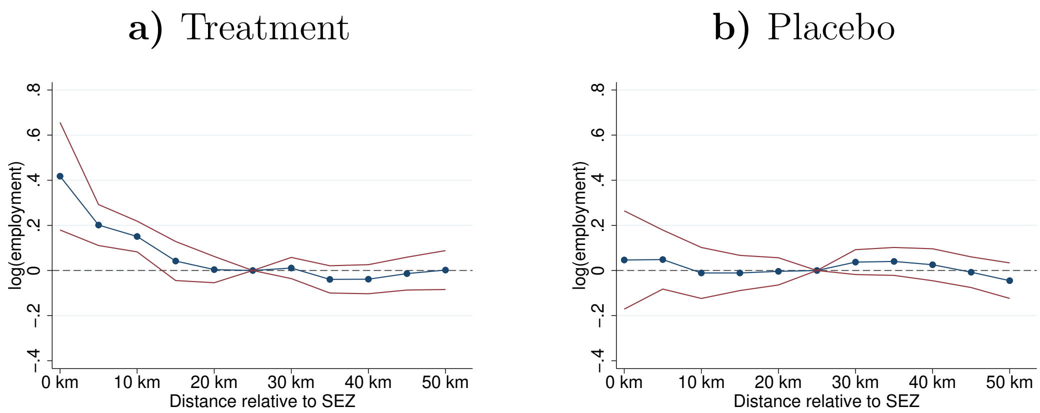

Notes: The dots indicate the estimated parameters $\hat{\beta}_d$. Each subscript d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Red lines indicate 95%-confidence intervals. Panel (a) refers to specification Eq. (2.1), panel (b) depicts the placebo test, where we rerun the analysis for the pre-treatment period 1998-2005. Standard errors are clustered at the district level. Employment data are based on the Economic Census for 1998, 2005 and 2013.

Given the 7-year gap between the censuses prior to treatment, we augment our data by annual nightlight information to more granularly assess outcome trends in treated and control municipalities in the years leading up to treatment. Estimates based on Eq. (3.2) are presented in Figure 2.5.2. We define municipalities up to 10km as treated (as they feature positive employment effects in the base analysis) and compare their nightlight outcomes to those in reference municipalities in a 20-25km distance radius. Treatment time is set to the year of SEZ notification by the board of approval, reflecting that SEZ

Figure 2.5.2: Nightlights in event study

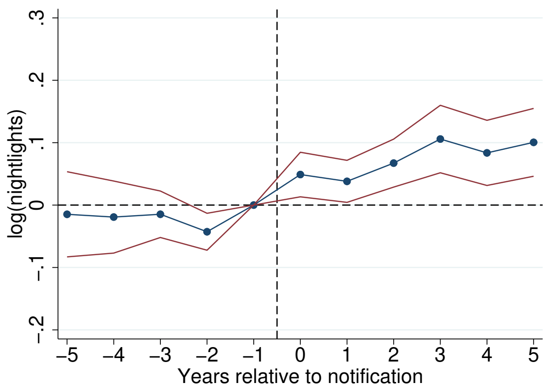

Notes: Event study estimates for 10km-regions around SEZs, municipalities in 20-25km distance serve as controls. The figure plots the $\hat{\theta}_k$ as estimated from Eq. (3.2) following Callaway and Sant'Anna, 2021. Endpoints are binned. Red lines indicate 95%-confidence intervals. Standard errors are clustered at the district level.

construction – and hence nightlight intensity – plausibly emerges from SEZ notification onward. Figure 2.5.2 shows that nightlights developed in parallel between treated and reference municipalities in the years prior to treatment, which corroborates the common trend assumption and the causal interpretation of our baseline estimates. Intuitively, the effect of interest is largest for manufacturing SEZs, whose production sites emit relatively much nightlight and tend to be located in rural areas with low underlying nightlight levels (making it easier to detect changes in nightlight intensity), see Appendix 2.B.2.

2.5.2 Job relocation or genuinely new employment?

An important aspect to understand is the extent to which the policy has generated *new* economic activity, relative to a mere relocation of manufacturing and service employment in space (Criscuolo, Martin, Overman, and Reenen, 2019; Ehrlich and Seidel, 2018; Kline and Moretti, 2014). Relocation can, in principle, be the sole driver behind the estimated employment effects. To rebut this concern, we suggest three pieces of evidence.

First, our baseline estimates show a stark picture in the sense that employment growth differs strongly between SEZ-hosting municipalities and their neighbors in distance circles up to 10km, while there is no significant difference between the employment growth of municipalities in further distance from the SEZ (10-50km). For this pattern to be consistent with relocation of economic activity, relocation costs must be invariant in space, i.e. additional employment must have been sourced from municipalities in distance radii of 10-50km at about equal rates, irrespective of their precise distance to the SEZ. This is at odds with existing empirical evidence, which shows a rather stable inverse relation between geographic distance and relocation costs (Bodemann and Axhausen, 2012; Rossi and Dej, 2020). Note that extending the distance radius to 200km from SEZs does not change this pattern (see Appendix 2.B.1).

Second, we explore whether the additional employment or the number of firms in SEZ municipalities and their direct neighboring jurisdictions in distance bands of up to 10km systematically correlate with changes in employment or the number of firms in municipalities in further distance. If the strong relative employment increase in SEZ-hosting municipalities and jurisdictions in close proximity to an SEZ (less than 10km distance) reflects relocation, we expect that larger employment increases in SEZ municipalities and surroundings are associated with stronger employment declines in jurisdictions in further

Table 2.5.1: Outcome changes in SEZs vs distant municipalities

	Distance to SEZ							
	10-15km	15-20km	20-25km	25-30km	30-35km	35-40km	40-45km	45-50km
Employment ($\leq$ 10km)	-0.021	-0.031	-0.023	-0.027	-0.000	0.057	0.009	0.019
	(0.047)	(0.039)	(0.048)	(0.029)	(0.046)	(0.038)	(0.043)	(0.046)
Firms ($\leq$ 10km)	0.008	-0.039	-0.039	-0.030	-0.009	0.034	-0.015	-0.011
	(0.055)	(0.039)	(0.038)	(0.030)	(0.044)	(0.036)	(0.042)	(0.041)
Observations	6,940	7,864	9,070	10,556	11,656	12,334	13,054	13,534
Municipality fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓

Notes: Regression results from Eq. (2.4). The upper panel depicts the effects of employment within a 10km radius around a SEZ on employment in municipalities in further distance bins. The lower panel reruns this specification using the number of firms as the dependent variable. Standard errors are clustered at the district level. Years included: 2005 and 2013. *** p<0.01, ** p<0.05, * p<0.1.

distance ($>$ 10km). We run a regression model of the following form:

$$ln(y_{i,t}) = \beta_0 + \beta_1 ln(y_{i,t}^{0-10}) + POST_t + \alpha_i + \epsilon_{it}, \tag{2.4}$$

where $y_{i,t}$ measures non-agricultural employment or the number of firms in municipalities in a distance of more than 10km to their closest SEZ while $y_{i,t}^{0-10}$ depicts either variable in SEZ-municipalities and its neighbors up to 10km. We run this regression separately for each distance bin $> 10km$.

The estimates for β_1 on employment are reported in the upper panel of Table 2.5.1. The columns reflect specifications for neighboring municipalities in different distance bins (Specification (1) comprises municipalities in a distance between 10-15km from an SEZ; Specification (2) municipalities in a distance between 15-20km etc.). Throughout all specifications the β_1-estimate turns out small and statistically insignificant, corroborating the notion that the observed baseline findings reflect a genuine increase in local non-agricultural economic activity rather than relocation of economic activity in space. Similar results emerge if we use the number of firms as the measure of economic activity (see the lower panel of Table 2.5.1).[20]

In Appendix 2.B.3, we also show in a back-of-the-envelope calculation that, even if we take the negative (and statistically insignificant) coefficient estimates for some distance bins as depicted in Table 2.5.1 at face value, the estimates suggest that only around 1% of

[20]Note that the number of municipalities per bin increases mechanically with distance to SEZs. Thus, the number of sourcing municipalities becomes larger relative to the number of potentially receiving municipalities (municipalities in <10km from an SEZ). Nevertheless, relocation would still imply that the estimated coefficients β_d decline in distance d.

the observed employment gain in SEZs and neighboring jurisdictions up to 10km relates to relocation from municipalities in further distance. This points to genuine increases in aggregate economic activity through SEZ establishment rather than relocation across space, which is consistent with prior findings in the literature (Criscuolo, Martin, Overman, and Reenen, 2019; Ehrlich and Seidel, 2018).

One remaining concern may be that some of the firms located in SEZs belong to larger firm groups, which may be active throughout India, including locations outside of the 50km radius accounted for in the above analyses. If firms relocate economic activity from other group locations to SEZ areas, it may hence not be captured in our prior analyses. To address this point, we undertake a third analysis, where we compile a firm dataset for India (see Appendix 2.B.3 for details), which allows us to identify SEZ firms and other corporations within the same firm group. We use this data to first show that most firms, which are active in SEZs (> 60%) are standalone Indian companies, limiting the scope for relocation. We, furthermore, use this data to construct firm groups and determine the impact of entering an SEZ on firm activity of treated company groups *outside of* SEZ areas – specifically, outside of the 50km radius reflected in our baseline analysis. The effect is identified by comparing the development of firm activity of companies outside of SEZ areas that belong to groups, which enter an SEZs, to the development of firms in company groups without SEZ connection. Event study estimates suggest that the economic activity of these treated and non-treated firms emerged in parallel prior to the implementation of the SEZ Act in 2005. After 2005, we find no indication for a reduction of sales, profits, assets, or the total wage bill of treated firms versus non-treated firms – which rejects that economic activity is relocated within company groups towards SEZ areas. If anything, we observe an *increase* in economic activity of treated firms outside of SEZ areas after the group entered an SEZ. This – as discussed in more detail in the appendix – is consistent with existing evidence in the literature, which suggest that economic activities at different firm group locations are complements rather than substitutes (see e.g. Desai, Foley, and Jr., 2009, Becker and Riedel, 2012, Chodorow-Reich, Smith, Zidar, and Zwick, 2024). See Appendix 2.B.3 for details.

2.5.3 Structural change and migration

If genuinely new jobs were created, then a natural follow-up question is who took up these jobs. We explore two channels: structural change and regional migration.

India is characterized by a large agricultural sector that accommodates about half of the

Figure 2.5.3: Sources for local non-agricultural employment growth

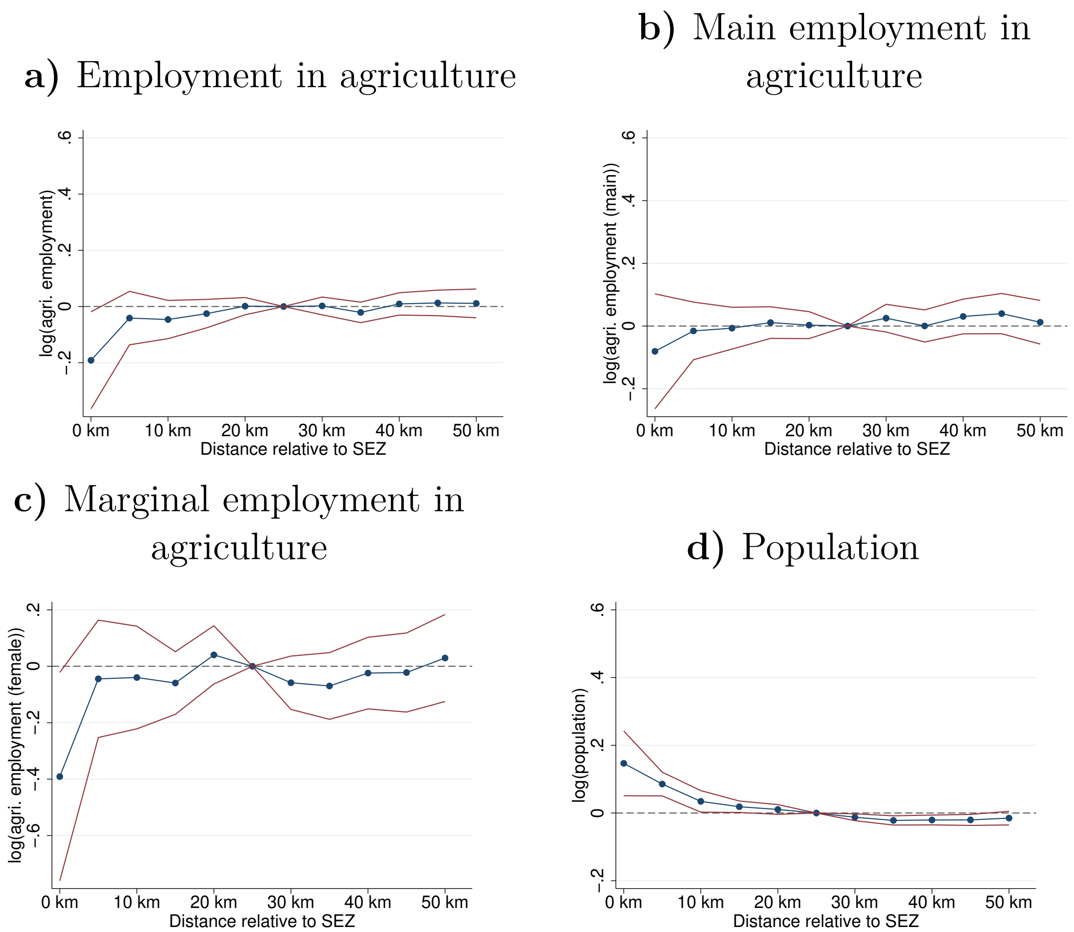

Notes: The dots indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.1). Each d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Panel (a) depicts results for agricultural employment. Panels (b) and (c) show results for main and marginal agricultural employment, respectively. Panel (d) depicts results for total municipal population. Red lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. All panels are based on the Population Census for the years 2001 and 2011.

working population, mostly in low-productivity jobs and in marginal employment relationships (International Labour Organization, 2013). Managing the transition from an agricultural to a manufacturing and service economy is widely believed to be one of the country's top challenges (Binswanger-Mkhize, 2013) and a promising avenue to higher-paid jobs and economic growth (Eichengreen and Gupta, 2011; Gollin, Lagakos, and Waugh, 2014; McMillan, Rodrik, Dani, and Verduzco-Gallo, 2014). We test whether SEZs contributed to this transition.

Specifically, we ask whether the documented increase in local non-agricultural employment in SEZ-areas is paralleled by a decline in agricultural employment. Based on the population census, we assign agricultural employment to municipalities following the procedure outlined in Section 2.4 and then rerun our baseline model in Eq. (2.1) using the log of the number of agricultural workers as the dependent variable. Panel (a) of Figure 2.5.3 indicates that the number of workers in the primary sector declined in SEZ municipalities after SEZ establishment. Quantitatively, the drop amounts to 17*pp* (p-value: 0.03). Neighboring municipalities up to 10km also experience a negative, but smaller effect.[21]

We can go one step further and split up the overall reduction of agricultural jobs into main and marginal employment. As shown in panels (b) and (c) of Figure 2.5.3, SEZs in particular led to a reduction in marginal agricultural employment – that is, in the number of agricultural workers that are employed for less than 183 days per year. Quantitatively, their number declined by 33*pp* (p-value: 0.04) in SEZ-municipalities relative to municipalities in the reference category. The point estimate for the response of the number of main agricultural workers is close to zero, in turn. It is thus the least attractive jobs in the agriculture sector, which drop off the market. In Appendix 2.B.4, we further show in a back-of-the-envelope calculation that the decline in agricultural employment explains about one third of the increase in non-farm jobs. We also provide additional evidence for employment increases in low-skilled services and manufacturing. Although we cannot follow individual workers across space and jobs, these results provide novel evidence in line with a transition from agricultural to low-skilled non-farm employment.

The result adds to a growing literature on structural change in less developed countries, which has documented large differences in productivity between agriculture and non-agriculture sectors and across regions (e.g. Gollin, Lagakos, and Waugh, 2009), stemming from various frictions (e.g. imperfect transport infrastructure (Asher and Novosad, 2020b), rural insurance networks (Munshi and Rosenzweig, 2016) or social norms that stigmatize migration (Beegle, Weerdt, and Dercon, 2011)). While the SEZ policy did not directly eliminate or lower these frictions, local productivity gains in non-agricultural employment in SEZs areas, related to agglomeration advantages (e.g. Combes and Gobillon, 2015) or the location of high-productivity (multinational) firms (see e.g. Alfaro-Urena, Manelici, and Vasquez, 2022), may have induced worker movements across sectors and regions, consistent with our findings.

Turning to the second channel, workers may be sourced from outside of the SEZ-municipality. While we have shown above that there is little evidence for net job relocation in space, workers might migrate towards SEZ-municipalities, resulting in higher local population

[21] Data on agricultural employment come from the Population Census that is available for 1991, 2001 and 2011.

growth. The pronounced population growth in India provided an ideal environment for such an effect. In our sample frame, the population increased from 127M to 146M between 2001 and 2011. Panel (d) of Figure 2.5.3 shows that population growth in SEZ areas was systematically higher than in control jurisdictions and there is indication of SEZ-induced population gains in neighboring areas. In principle, the difference in population growth might also reflect differences in fertility rates (e.g. triggered by higher income opportunities in SEZ areas). Given that we study a rather short time frame, we consider this explanation to be of second-order importance at best.[22]

2.6 Heterogeneous effects

In this section, we shed light on heterogeneous treatment effects by gender (2.6.1), firm size (2.6.2) and zone characteristics (2.6.3) to explore the anatomy of the employment response.

2.6.1 Employment effects by gender

Female workers are a particularly vulnerable group in the Indian labor market as unemployment rates among women tend to be high and discrimination is a long-standing phenomenon (Klasen and Pieters, 2015; Srivastava and Srivastava, 2015). Against this background, providing better income opportunities to women by integrating them into the formal labor market would be an important effect of the policy. One presumption proponents of the SEZ-policy have expressed is that additional jobs in manufacturing or services would be sourced from the unused female workforce or from women being employed marginally in the agricultural sector and that women might be the main beneficiaries of such policies (e.g. Bacchetta, Ernst, and Bustamante, 2009; Rama, 2003; Brussevich and Dabla-Norris, 2020). These hopes were further spurred by rising female employment shares in export-oriented industries in many less-developed countries (Bussmann, 2009; Ozler, 2000).

Our data allow us to split up employment effects by sector and gender. Panels (a) and (b) in Figure 2.6.1 reveal that in particular female employment declined in the agricultural

[22] Employment growth in manufacturing and services could also be associated with higher labor-force participation or lower unemployment. As related data are unavailable at the municipality level, we can explore neither of these underlying sources empirically. Workers could, on top of that, also commute from neighboring locations to SEZ areas. Commuting is rather uncommon in India, however, as public transport networks are not well developed and services tend to be infrequent. Census data for 2011 suggests that only around 18% of the Indian workforce travels more than 10km to work (own calculation based on the Population Census 2011).

sector ($-29pp$, p-value: 0.001) relative to reference municipalities, while the effect on men was closer to zero and statistically insignificant. An explanation for this gender effect

Figure 2.6.1: Employment effects by gender

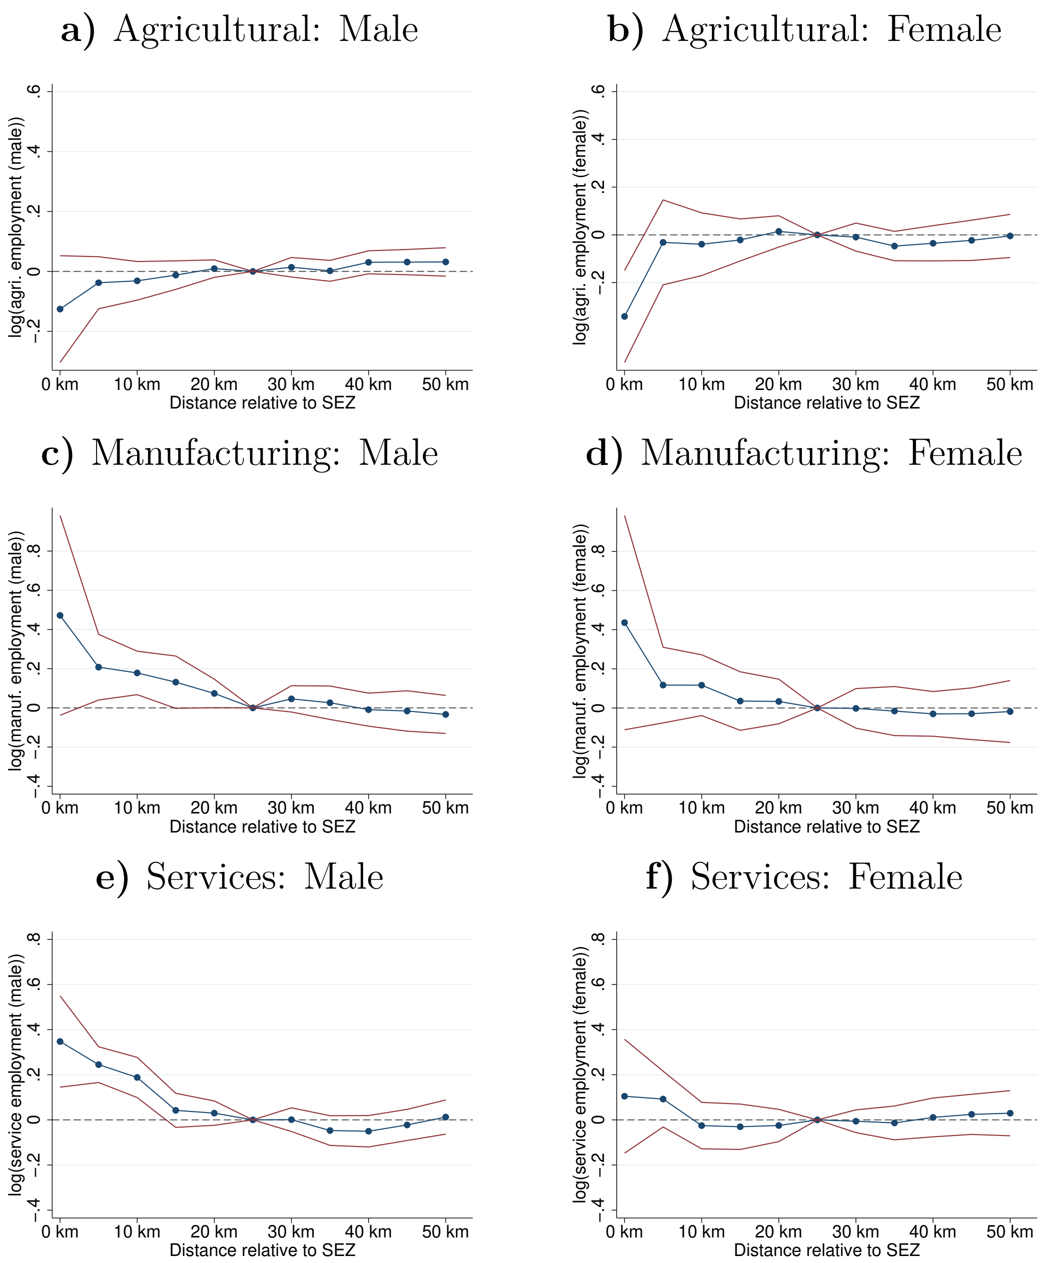

Notes: The dots indicate the estimated parameter $\hat{\beta}_d$ according to Eq. (2.1). Each d refers to a distance on the horizontal axis e.g. the coefficient at 0km refers to $d = 0$. Red lines indicate 95% confidence intervals. Standard errors are clustered at the district level. Panels (a)-(d) are based on the Economic Census for the years 2005 and 2013. Panels (e)-(f) are based on the Population Census for the years 2001 and 2011.

might be that only about 15% of agricultural businesses are owned by women rendering them more responsive to new job opportunities in manufacturing (Agarwal, Anthwal, and Mahesh. Malvika, 2021). Moreover, our data reveal that female workers account for 59% of marginally employed agricultural workers such that non-agricultural jobs might offer an appealing alternative for many. The decline of female employment in agriculture is paralleled by a pronounced increase of female workers in manufacturing by 55*pp* (which just fails to gain statistical significance at conventional levels, p-value: 0.118, panel (d)), but a much smaller and insignificant effect in services. Male employment, in contrast, rises in both manufacturing (60*pp*, p-value: 0.069) and services (41*pp*, p-value: 0.001) as can be seen from panels (c) and (e). As high-skill-intensive IT-zones play a quantitatively important role within the service industry in our sample, it is plausible that additional employment is taken up by skilled workers rather than being drawn from the predominantly low-skilled agricultural sector. A potential source of skilled-workers could be regional migration (see panel (d) of Figure 2.5.3).

In sum, we conclude that employment changes in agriculture, manufacturing and services point in the same direction for men and women. The decline in agriculture appears more pronounced for women, the increase in services is higher for men.[23] The sectoral shifts are hence centered around female employment (see Section 2.5.3).

2.6.2 Employment effects by firm size

Our data further allow us to decompose the overall employment effect by firm size. While some elements of the SEZ-policy (e.g. tariff-related benefits or lighter regulations) mainly target large firms, others – e.g. the corporate tax holidays provided – are equally attractive for smaller entities. Smaller and informal firms may also find it attractive to co-locate in or close to SEZs if they are connected to other (exporting) firms through input-output links. In the following, we assess the impact of SEZ establishment on employment in firms with more or less than 10 workers, where the latter are classified as informal by official statistics in India (NCEUS, 2009). Our analysis hence provides indication as to what extent studies underestimate aggregate employment responses to local policies if the focus is on the formal sector only. Second, distinguishing between small and large firms is also of interest as firm size correlates with economic outcomes like worker productivity and workers' wages (Idson and Oi, 1999; Oi and Idson, 1999) and with firms' fiscal contributions

[23] Only the estimates for services are statistically significantly different from each other.

Figure 2.6.2: Employment effects by firm size

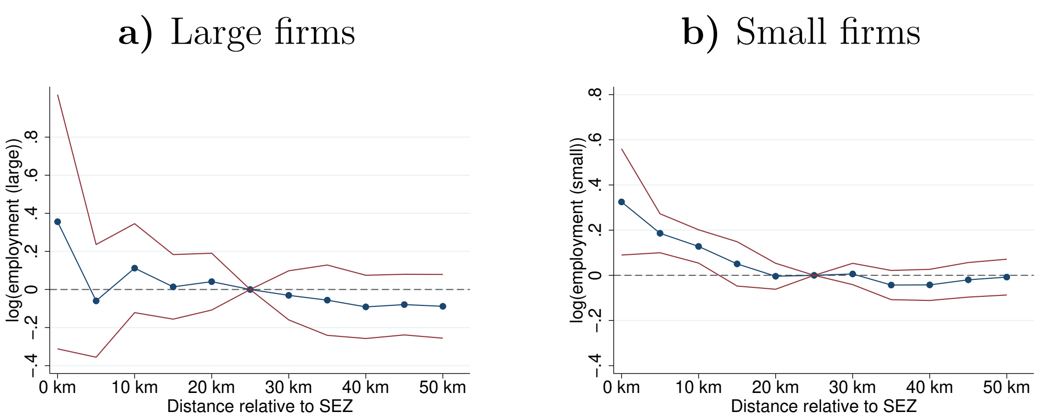

Notes: A firm classifies as small if it employs not more than 10 workers. The dots indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.1). Each d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Red lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. All panels are based on the Economic Census for the years 2005 and 2013.

(LaPorta and Shleifer, 2014; McCaig and Pavcnik, 2021).[24]

Panels (a) and (b) of Figure 2.6.2 report employment responses separately for large (more than 10 workers) and small firms (up to 10 workers), respectively. We find a strong, but insignificant effect for large firms of 53*pp* and a somewhat smaller, but significant employment gain of 38*pp* for small firms. The insignificant estimate for large firms likely relates to the relatively small number of large firms per municipality (which lowers the statistical power of the analysis). Complementary, we show in Appendix 2.B.5 that the SEZ-policy has stimulated entry of small informal firms, especially in areas outside SEZs. Small informal firms are hence found to add significantly to the observed positive local economic effect induced by SEZ establishment.

2.6.3 Zone characteristics

One feature of the small existing literature on the spatial effects of SEZs is that studies largely assume SEZs to be homogeneous entities (e.g. Lu, Wang, and Zhu, 2019; Wang, 2013). This is at odds with real-world settings (World Bank, 2008). Zones in India differ in two key dimensions: First, there is heterogeneity in zones' main industry denomination.

[24]Note that productivity and wages of small firms in the manufacturing and service sector are arguably still higher than wages in agriculture, especially in comparison with marginal agricultural work (workers would otherwise not switch jobs). Fiscal contributions also correlate with firm size as small firms are exempt from certain insurance and social security tax payments and, in general, show weaker tax compliance behavior than larger entities (LaPorta and Shleifer, 2014; McCaig and Pavcnik, 2021).

There are IT, pharma, engineering, apparel or manufacturing zones (the latter are tabbed 'multiproduct zones'). Zones further differ in whether they are developed and run by a private or a public body. In this section, we assess how these characteristics shape the impact of SEZs on local economic activity.

Public vs. private SEZs. As depicted in panel (a) of Figure 2.6.3, more than two thirds of the zones that went into operation during our sample period were developed and run by a private entity. While privately developed zones do not systematically differ from their publicly developed counterparts in terms of area size (see panel (b)), they tend to be located in larger and more prosperous areas (as determined by host municipalities' employment and nightlight intensity, see panels (c) and (d)).[25] This is consistent with public developers putting a stronger emphasis on creating new employment in less prosperous regions compared to private developers, who primarily seek to maximize profits.

There are also reasons to believe that the local employment impact of public and private SEZs may differ. On the one hand, public bodies have less incentives to run projects efficiently (see e.g. Megginson and Netter, 2001) and the optimal size of publicly developed zones may therefore, ceteris paribus, be smaller than the optimal size of private zones. On the other hand, public zones may exert stronger local employment effects as public developers often pursue employment goals when designing SEZs, while private developers first and foremost aim for profit maximization. To test for effect heterogeneity along these lines, we estimate a model of the following form:

$$\begin{aligned} ln(y_{it}) = &\sum_{d=0,d\neq 5}^{10} \beta_d D_{[d_i=d]} \times POST_t + \sum_{d=0,d\neq 5}^{10} \theta_d D_{[d_i=d]} \times POST_t \times priv.developer_i \\ &+ POST_t \times priv.developer_i + \boldsymbol{\eta}' \left(\mathbf{X}_i \times POST_t\right) + POST_t + \alpha_i + \epsilon_{it}, \end{aligned} \tag{2.5}$$

where the variable definitions correspond to Eq. (2.1) and $priv.developer_i$ is a dummy variable indicating that the closest SEZ to municipality i is developed by a private developer. One challenge when estimating Eq. (2.5) is that SEZs do not only differ in their status of being developed by a private or public body, but also in their industry denomination. If the industry denomination correlates systematically with private and public development status and with SEZs' local employment impact, estimates of θ_d may be confounded. Descriptive statistics indeed suggest that the fraction of IT zones is, for example, larger among

[25] Consistent GDP data are, unfortunately, not available at the level of Indian municipalities. Henderson, Storeygard, and Weil (2012) show that nightlights are a reasonable proxy for economic development and income growth at subnational levels.

Figure 2.6.3: SEZ characteristics by industry and ownership

Notes: SEZ-municipality characteristics are based on the year 2005. Authors' own calculations based on SEZ information from the Ministry of Industry and Commerce, the Economic Census and DMSP-OLS Nighttime Lights Time Series provided by the National Oceanic and Atmosphere Administration (NOAA).

private than among public SEZs. We draw on exact matching to address this concern. In the base analysis, we match observations according to the industry class of the closest SEZ located in distance d_i from municipality i to balance differences in industry denomination across SEZs developed and run by private and public entities.

Panel (a) of Figure 2.6.4 plots the effects of SEZs on local employment conditional on industry denomination and separately for public and private SEZs (β_d and $\beta_d + \theta_d$ in Eq. (2.5)). It is evident that the effects do not differ systematically between publicly and privately developed SEZs. If anything, employment effects are larger in publicly developed zones, but the effects are not statistically different from each other. In Appendix 2.B.1, we report additional results where we re-estimate Eq. (2.5), first, without matching and, second, applying coarsened exact matching and accounting for SEZ's industry denomination *and* the area size of the SEZ relative to the area of its hosting municipality (Blackwell, Iacus, King, and Porro, 2009; Iacus, King, and Porro, 2012). The latter variable is coarsened

Figure 2.6.4: Employment effects by zone type (CEM)

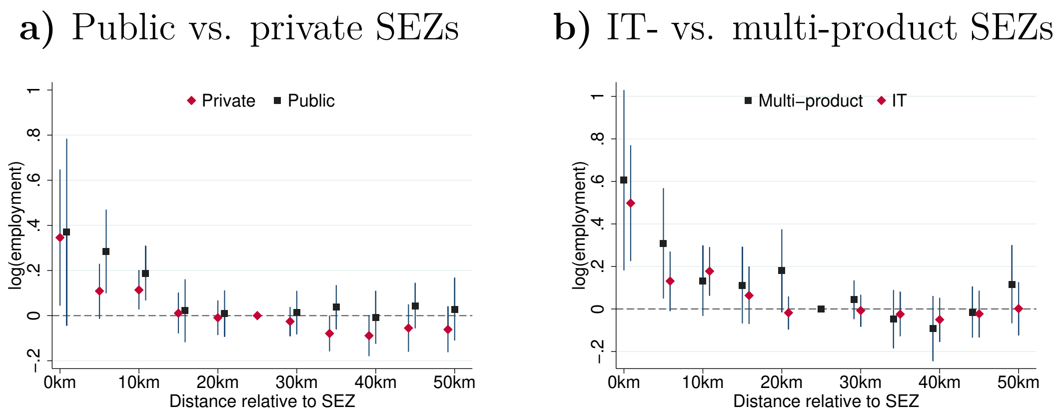

Notes: The plotted coefficients are estimated according to Eq. (2.5). In panel (a) (panel (b)), black squares depict the effects of public (multi-product) SEZs on employment in the respective distance bins ($\hat{\beta}_d$). Red diamonds show the effects for private (IT) SEZs ($\hat{\beta}_d + \hat{\theta}_d$). Each d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Black lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. Regressions include municipality and year fixed effects. Observations are re-weighted using coarsened exact matching over designated industry (ownership-type) and with *private* (*IT*) as the treatment category. For the purpose of giving a comprehensive picture of the full set of SEZ location choices the IT-sample includes also large municipalities. Employment data are based on the Economic Census for the years 2005 and 2013.

based on the default autocut algorithm as in Blackwell, Iacus, King, and Porro (2009). All specifications yield similar results.

Sector-specific effects. The impact of SEZs on local economic activity may also hinge on SEZs' industry denomination. In the following, we will in a first step compare IT and multi-product (i.e. manufacturing) zones. Testing for effect heterogeneity in this dimension again comes with the challenge that industry denomination might correlate with other zone characteristics like the type of developer and zones' size relative to the size of the host municipality. Our data indeed suggest that IT-zones tend to be hosted by systematically larger jurisdictions than multi-product zones. This is intuitive since IT-firms demand high-skilled labor, which can be found predominantly in big cities.[26] Furthermore, the minimum area size requirement for IT-zones is substantially smaller than for other zone types, facilitating the establishment of IT-SEZs in areas where land is scarce and costly. Multi-product SEZs are, in turn, observed to be located in smaller municipalities at the coast, reflecting their need for proximity to physical infrastructure such as ports for exporting manufactured goods.

We apply coarsened exact matching to account for these features by estimating a model sim-

[26] Note that we include municipalities with more than 500K inhabitants when studying heterogeneous effects across industries since a significant share of IT-SEZs is located in large cities.

ilar to Eq. (2.5) where we replace $priv.developer_i$ by an industry identifier $multiproduct_i$. In panel (b) of Figure 2.6.4, we match zones by developer type (private vs. public body). In Appendix 2.B.1, we present additional results, where we match on zones' size relative to the host municipality. Across both specifications, point estimates are somewhat higher for multi-product zones in some distance bins, but are never statistically different from IT-zones. Similar conclusions emerge for other industries (pharma, engineering, apparel), see Appendix 2.B.1. This suggests that the aggregate local employment effects are comparable across SEZs of different type.

Note that the existence of (privately-developed) SEZs with different industry denominations, which also exert comparable local employment effects, suggests that different SEZ features – tax cuts for export income, tariff reductions and ease of regulatory burden – have pull and attract firms to SEZ areas. In Appendix B.1, we substantiate this point by showing that IT and manufacturing firms substantially differ in relevant underlying characteristics – the regulatory burden (which is particularly high for firms in the IT sector), import intensity (which is particularly high for firms in manufacturing) and export intensity (which is broadly comparable for firms in manufacturing and IT) – and that there is within-industry selection of firms with high export and import-intensity into SEZ areas.

2.7 Has the SEZ policy been cost-effective?

We finally draw on a simple back-of-the-envelope calculation to obtain an understanding whether the Indian SEZ-policy has been cost-effective. The exercise relates net employment changes to the fiscal expenditures of the program (Criscuolo, Martin, Overman, and Reenen, 2019; Lu, Wang, and Zhu, 2019). Information on foregone revenues is taken from the Indian Ministry of Finance, which monitors the SEZ-policy and publishes foregone revenues as the total amount of income tax concessions claimed by SEZ-firms and SEZ-developers (Ministry of Finance, 2015). For the years 2006-2013, these concessions amounted to INR 596.2 billion, equivalent to USD 9.85 billion based on 2013 purchasing power parity (PPP) exchange rates.[27] Our baseline estimates suggest that the policy, in the aggregate, created 1.25 million new jobs (see Appendix 2.B.3 for details).[28] This

[27] For 2006, official statistics only include the aggregate income tax concessions for all incentive programs in India. We approximate the SEZ-related foregone revenues in 2006 by extrapolating the share of SEZs in total revenue loss for 2007 (where the revenue losses were split up by incentive programs) to 2006.

[28] One concern might be that SEZ-induced employment effects may systematically differ across municipalities of different size, which may potentially bias the estimate for the specified aggregate employment response. Note that, if we split the set of municipalities with less than 500K inhabitants into four equally-sized bins

translates into revenue costs of INR 475,158 (USD 7,853, PPP) per newly created job. The ratio between workers' wages and fiscal costs per job is 0.72 if jobs are created for the eight years we study and workers earn the Indian minimum wage ($3,562$ INR per month). If we additionally account for the decrease in agricultural employment and assign a value of INR 50 per day as agricultural income (Saini, Gulati, Braun, and Kornher, 2020), we retrieve a ratio of 0.62. These estimates are within the broad range of prior studies on place-based policies (Chodorow-Reich, 2019; Criscuolo, Martin, Overman, and Reenen, 2019). Note, however, that the latter are largely set in the developed world, limiting comparability with our findings.[29] For India, most existing studies fail to report program costs. An exception is Chaurey (2017), who shows that tax concessions granted to firms in two Indian states created employment at much higher fiscal costs than the SEZ program studied in our paper.[30]

While these back-of-the-envelope estimates offer a valuable benchmark, caveats need to be kept in mind. They include that the foregone revenues calculated by the Indian Finance Ministry abstract from firms' behavioral response to the SEZ policy by assuming that all foregone taxes would have been paid under the counterfactual and by abstracting from spillovers to other tax bases.[31] More broadly, note that our bang-for-the-buck estimates follow the spirit of several prior papers on place-based policy interventions (e.g. Lu, Wang, and Zhu, 2019 and Criscuolo, Martin, Overman, and Reenen, 2019) but do not directly speak to the welfare implications of the SEZ policy (e.g. accounting for sectoral or regional migration costs and productivity shifts induced by the policy (Combes and Gobillon, 2015)). We consider analyses in this direction to be a fruitful avenue for future research.

and run our regression model separately in each subset of municipalities before aggregating, this yields very similar estimates for the aggregate employment effect (1.24 million new jobs).

[29] Information on the minimum wage is taken from: https://countryeconomy.com/national-minimum-wage/india, Last retrieved: June 21, 2023. Also note that the minimum wage only binds in the formal sector. But prior evidence for India shows that it also shapes informal wages (Kar and Khattar, 2023).

[30] Chaurey (2017) estimates that the tax incentives created 33,000 jobs and that the (upper bound of the) fiscal cost to taxpayers were INR 66 billion. This yields fiscal costs per newly created job of INR 2 million. The SEZ policies assessed in our study hence created jobs at less than a quarter of the tax costs.

[31] We also do not observe workers' wages but have to rely on the approximation by the minimum wage. Moreover note that our estimates on the aggregate employment gain comprise the SEZ-related employment responses in larger urban areas, which are challenging to estimate and are more likely to include a margin of error (see Appendix 2.B.3 for a more detailed discussion).

2.8 Conclusion

This chapter has studied a highly prevalent type of place-based policy in less-developed countries: the establishment of Special Economic Zones. While the number of SEZs in the developing world has increased steeply over recent decades, there is hardly any evidence on their effectiveness in fostering local economic development. A notable exception are studies on SEZs in China. But given the particularities of the Chinese institutional context, there is scepticism in the policy domain that the Chinese experience extends to SEZs in other countries (see e.g. African Development Bank (2016) and World Bank (2017)).

We add to the literature by studying the local economic impact of SEZs in India. The empirical analysis relies on granular census information and on hand-collected data on the location and characteristics of SEZs. We use a transparent empirical identification design to document that the SEZ Act stimulated quantitatively important non-agricultural employment growth in SEZ-hosting municipalities and their close neighbors. Additional analyses suggest that genuinely new non-agricultural jobs were created (rather than jobs being relocated in space). We furthermore shed light on the anatomy of the response: We present evidence consistent with workers migrating towards SEZ areas to take up the new jobs. And we document that SEZ establishment stimulated sectoral transition from the primary sector to manufacturing and services. This sectoral shift centers around local female employment and may thus have added to the empowerment of women. Last but not least, the positive local employment effects emerge across different types of SEZs: privately and publicly run zones and SEZs with different industry denominations. Overall, we interpret our findings to dispel the general pessimism about zone programs in developing countries outside of China.

Appendix

2.A Data

This appendix complements Section 2.4 in the main paper providing more information on the data compilation process (Section 2.A.1), descriptive statistics (Section 2.A.2) and the geographic location of SEZs by industry (Section 2.A.3).

2.A.1 Data compilation procedure

Figure 2.A.1 illustrates each individual step implemented in QGIS 3.10. to arrive at the municipality sample.

Figure 2.A.1: Automated workflow in QGIS 3.10 to obtain final municipality sample

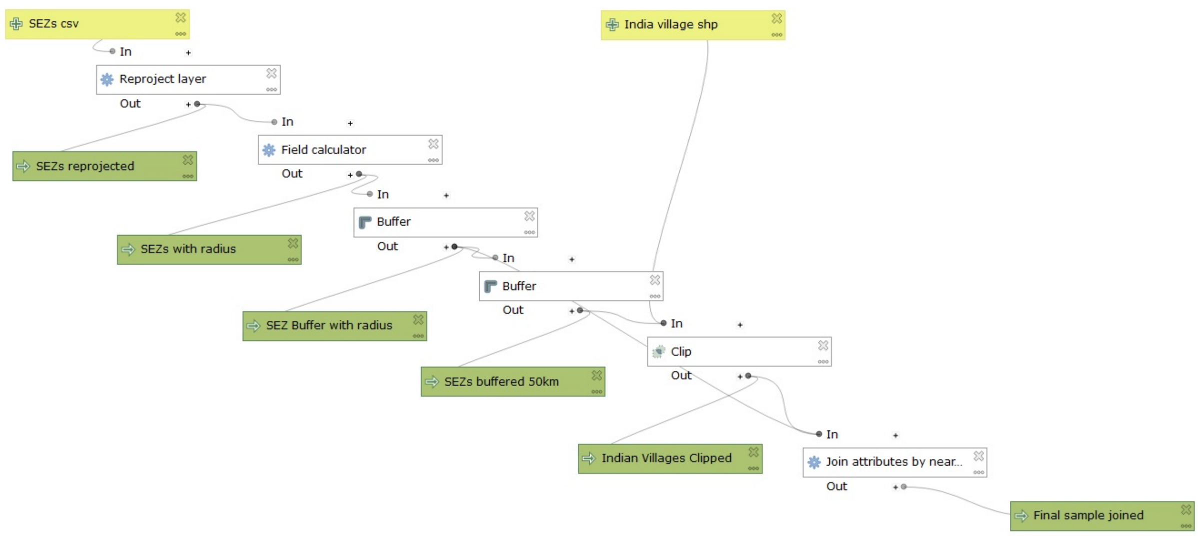

Figure 2.A.2: Mapping municipalities into distance bins around SEZs

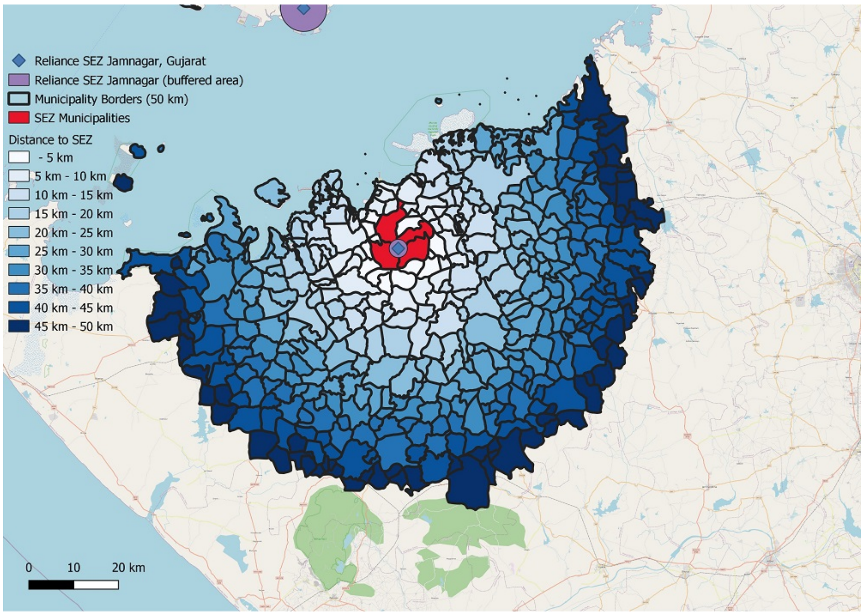

Notes: This figure illustrates the procedure of mapping municipalities into distance bins using the“Reliance SEZ" in Jamnagar (Gujarat) as an example.

Figure 2.A.2 illustrates the procedure for the Reliance SEZ in Jamnagar, where the red-colored polygons correspond to municipalities, whose administrative borders intersect with the SEZ-area. We consider these municipalities as municipalities that contain an SEZ. The blue-shaded polygons illustrate neighboring municipalities, classified by their distance to their closest SEZ ("Reliance SEZ" in the example above). The light blue color indicates municipalities which are within a 5km distance to their closest SEZ; darker blue colors indicate municipalities in a distance of 5-10km, 10-15km etc. to the closest SEZ (up to 50km).

2.A.2 Descriptive statistics

Table 2.A.1 summarizes the baseline sample, i.e. excluding large cities with a population larger than 500K.

Table 2.A.1: Descriptive statistics

	Mean	*SD*	*Median*	*# Municipalities*	*# Obs.*
Economic Census					
- Non-agricultural employment	290.0	2,457	41	49,669	140,386
- Male non-agricultural employment	220.2	1,967	30	49,669	140,386
- Female non-agricultural employment	69.77	573.2	8	49,669	140,386
- Non-agricultural employment (large firms)	87.68	1,518	0	49,669	140,386
- Non-agricultural employment (small firms)	202.3	1,337	36	49,669	140,386
- Manufacturing employment	113.0	1,307	7	48,093	96,186
- Service employment	211.3	1,637	34	48,093	96,186
- Number of firms	115.6	692.1	23	49,669	140,386
Population Census					
- Agricultural employment	520.1	793.4	303	42,910	127,868
- Male agricultural employment	330.6	501.2	194	42,910	127,868
- Female agricultural employment	189.5	333.5	93	42,910	127,868
- Main agricultural employment	433.9	706.7	240	42,654	85,308
- Marginal agricultural employment	117.9	223.8	42	42,654	85,308
- Population	3,061	15,224	1,043	42,910	127,868

Notes: Small and large firms are classified according to the 10-worker rule. Marginal workers (as opposed to main workers) work less than 183 days a year. Information on main and marginal workers is only available for the years 2001 and 2011. Information on sector employment (Manufacturing, Services) is only available for the years 2005 and 2013. The sample consists of all municipalities which are located within a 50 km radius of one of the 147 SEZs and observed for at least two consecutive rounds in the economic census. Municipalities with more than 500K inhabitants are excluded.

Table 2.A.2 summarizes additional information on SEZs.

Table 2.A.2: Descriptive statistics SEZ-level data

	Mean	*SD*	*Median*	*N*
- Year of notification	2007	1.17	2007	147
- Year of operation	2010	2.07	2010	147
- Developing time (in years)	2.67	1.76	3	147
- Area sq. km	1.76	7.40	0.27	147
- Private SEZ	0.77	0.42	1	147
- Public SEZ	0.23	0.42	1	147
- IT SEZ	0.57	0.50	1	147
- Multiproduct SEZ	0.09	0.29	1	147
- Pharma SEZ	0.09	0.29	0	147
- Engineering SEZ	0.12	0.32	0	147
- Apparel SEZ	0.05	0.23	0	147

Notes: Authors' own calculations based on sources described in the main text. Private implies that the SEZ was established by a private body. Year of operation denotes the year in which the SEZ initialized its operation. Sample includes all SEZ that became operational until 2013.

2.A.3 Geographical location of SEZs by industry and developer

The maps in Figure 2.A.3 show the geographic distribution of different types of SEZs (IT, multi-product and public/private, respectively) across India.

Figure 2.A.3: Geographical location of SEZs by industry and developer

a) SEZs by industry **b)** SEZs by developer

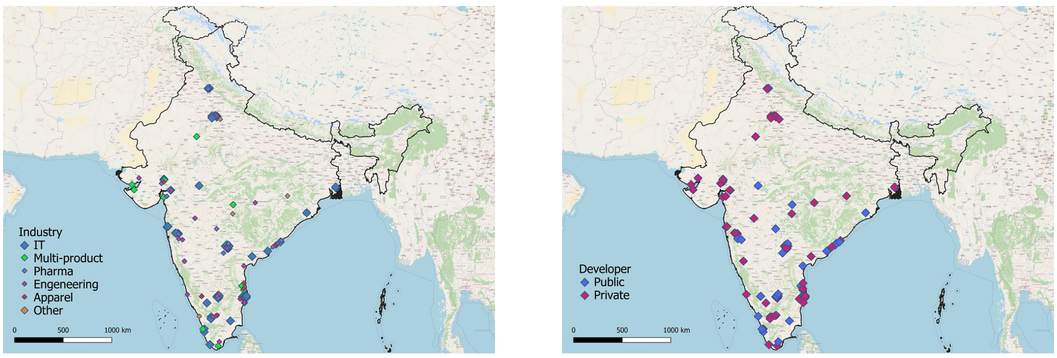

Notes: Panel (a) plots the location of all SEZs in India that were established under the SEZ Act 2005 and became operational until 2013 by their industry designation. Panel (b) plots the location of all SEZs in India that were established under the SEZ Act 2005 and became operational until 2013 by their type of developer.

2.B Results

This appendix complements Section 2.5 of the main paper. We present additional robustness checks for our baseline results (Section 2.B.1), further details on the nightlights event study (Section 2.B.2), the relocation analysis (Section 2.B.3) and the structural change analysis (Section 2.B.4) and show additional results for outcomes like infrastructure (Section 2.B.5)

2.B.1 Robustness

Baseline results. We check the robustness of our baseline results with regard to (1) alternative distance bin classifications (Figure 2.B.1), (2) alternative standard error clustering (Figure 2.B.2), (3) including municipalities up to a distance of 200km (Figure 2.B.3), (4) including large cities (Figure 2.B.4), (5) estimating our baseline model without additional controls and with CEM matching (2.B.5, panels (a)-(b) and (c)-(d) respectively).We find that none of these modifications alter the conclusions derived in Section 2.5.

Figure 2.B.1: SEZ effect on employment (10km and 2.5km distance bins)

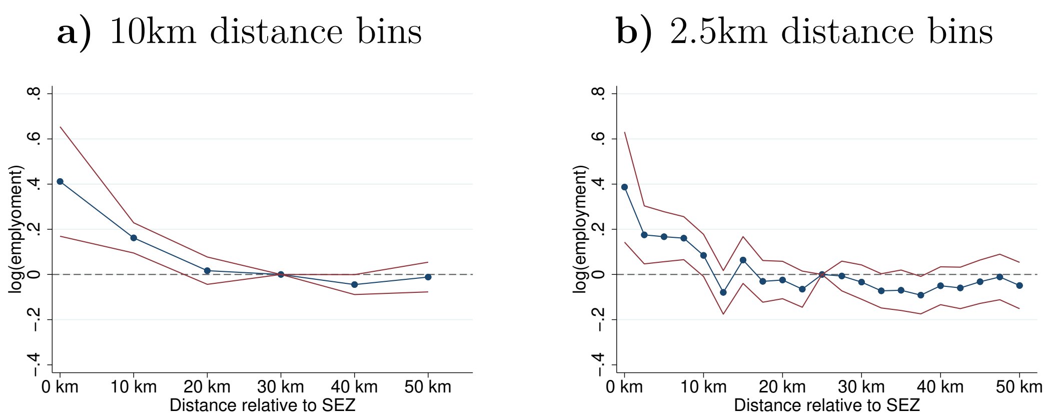

Notes: In this figure, distance bins are redefined as spreading 10km (panel (a)) and 2.5km (panel (b)). The dots indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.1). Red lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. Employment data are based on the Economic Census for the years 2005 and 2013.

Figure 2.B.2: SEZ effect on employment (SE clustered by closest SEZ and Conley)

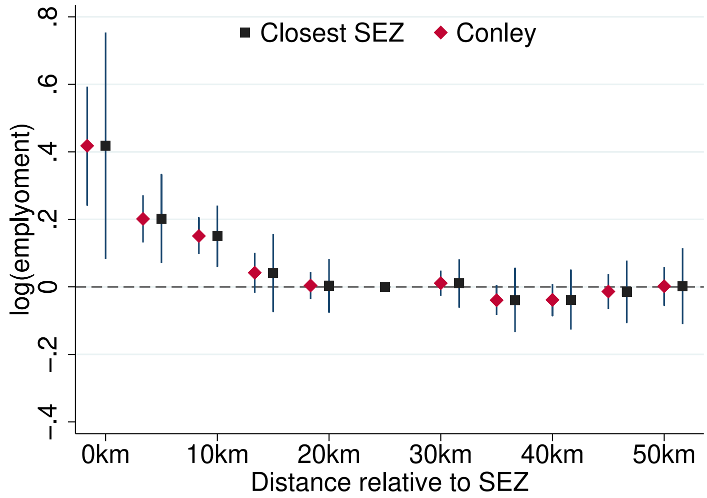

Notes: The dots indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.1). Each d refers to a distance on the horizontal axis e.g. the coefficient at 0km refers to $d = 0$. Red diamonds show the effects for when using Conley standard errors (Conley, 1999) with a distance cut-off at 30km. Black squares depict the results when clustering by closest SEZ. Red lines indicate 95%-confidence intervals. Employment data are based on the Economic Census for the years 2005 and 2013.

Figure 2.B.3: SEZ effect with 200km radius

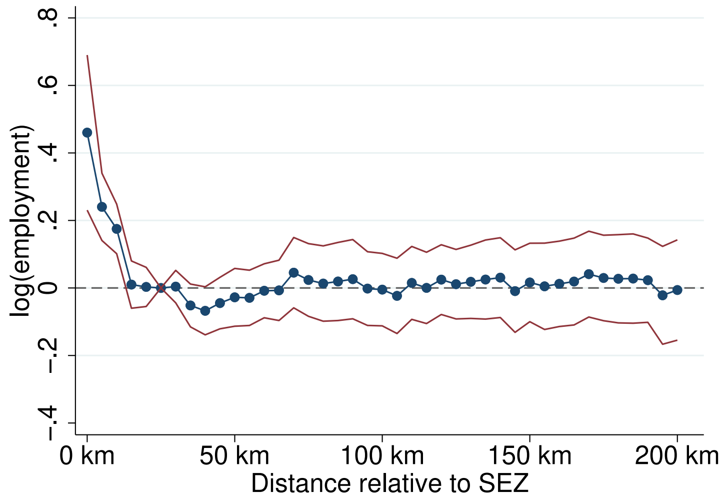

Notes: The dots indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.1). In this figure, the radius drawn around SEZs has been increased from 50km to 200km. Red lines indicate 95%-confidence intervals. The standard errors are clustered at the district level. Employment data are based on the Economic Census for the years 2005 and 2013.

Figure 2.B.4: SEZ effect on employment with and without large cities

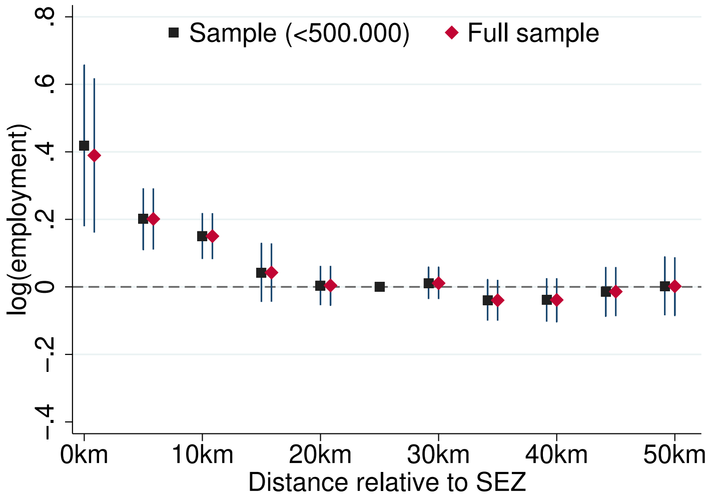

Notes: The squares and diamonds indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.1). Black squares depict the effects of SEZs on employment in small municipalities (baseline), i.e. $\leq$ 500K ($\hat{\beta}_d$). Red diamonds show the effects including large municipalities, i.e. $> 500K$ ($\hat{\beta}_d + \hat{\theta}_d$). Each subscript d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Black lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. Employment data are based on the Economic Census for the years 2005 and 2013.

Figure 2.B.5: Spatial difference-in-differences model

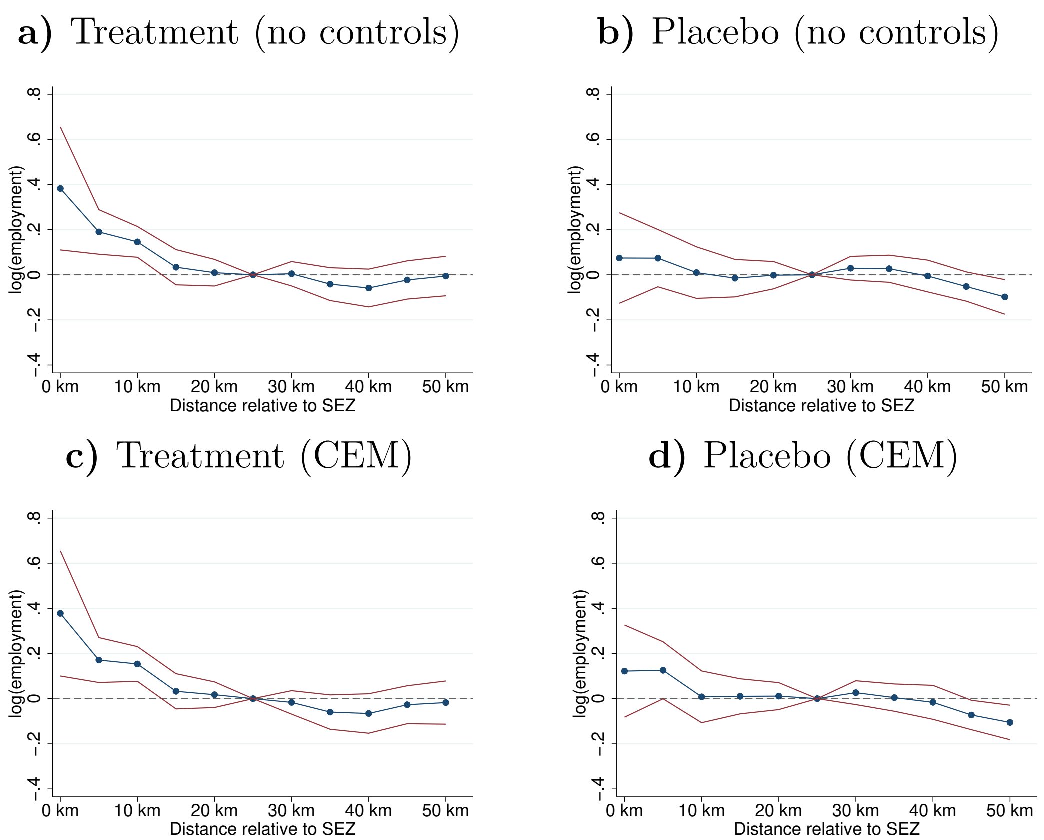

Notes: The dots indicate the estimated parameters $\hat{\beta}_d$. Each subscript d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Red lines indicate 95%-confidence intervals. Panel (a) refers to specification Eq. (2.1) without the controls $\boldsymbol{\eta}'(\mathbf{X}_i \times POST_t)$, panel (c) is based on coarsened exact matching (CEM). The panels in the right column depict the respective placebo regressions. Standard errors are clustered at the district level. Employment data based on the Economic Census for 1998, 2005 and 2013.

We, moreover, explore whether the Mahatma Gandhi National Rural Employment Guarantee Act (MGNREGA), a public work program enacted in 2005 and thus in parallel to the SEZ policy, may act as a confounder in our analysis. MGNREGA guarantees at least 100 days of wage employment per year for unskilled manual work. If no job is found, the government pays transfers to workers who applied. At least one member of every household is eligible for this program. We provide three pieces of evidence, which suggest that our findings are not affected by the MGNREGA program.

First, MGNREGA was implemented at the district level while our analysis uses variation at the municipality level. Districts in India are large spatial units. On average, a district accommodates around 900 municipalities and has a radius of 63.6km (determined based on shapefiles that capture the 641 districts from the Population Census 2011). Our main empirical identification stems from within-district variation: the distance of SEZ-municipalities to reference location is only 25km. This setup makes it unlikely that effects from the MGNREGA program bias our estimates.

Second, MGNREGA provided jobs in the public sector, with a focus on the maintenance and the construction of infrastructure (roads, wells, etc.). As documented in Figure 2.B.13, we find no indication that infrastructure construction emerged differentially between SEZ-locations and reference municipalities, again speaking against MGNREGA acting as a confounder in the analysis.

Third, we explicitly control for MGNREGA takeup in our empirical analysis.[32] Specifically, we compute the total number of person-days under MGNREGA in district k at time t relative to the total district population ($MGNREGA_{kt}$) and use this variable as an additional control.[33] The specification reads:

$$\begin{aligned}\ln(y_{it}) = &\sum_{d=0, d\neq 5}^{10} \beta_d (D_{[d_i=d]} \times POST_t) + \boldsymbol{\eta}' (\mathbf{X}_i \times POST_t) + POST_t \\ &+ \xi \times MGNREGA_{kt} + \alpha_i + \varepsilon_{it},\end{aligned} \tag{2.6}$$

where the other variables are defined as in our main specification. The results are shown in Figure 2.B.6 and resemble our baseline estimates. This corroborates that MGNREGA does not act as a confounder in our analysis.

[32] We added to each district information on MGNREGA take-up which has been made publicly available by Clement Imbert. The data was retrieved from LINK, last accessed: April 19th, 2024. Also note that, we ran additional specifications, where we estimated our baseline model separately for urban and rural municipalities, following the observation that MGNREGA was targeted at rural areas only. The results point to a negative (positive) effect of SEZ establishment on marginal agricultural (non-farm) employment in rural places, while there is no significant effect in urban areas. The latter finding reflects a lack of statistical power, however, as the number of urban municipalities in our data is small.

[33] For the year 2005, we assign the number of person-days under MGNREGA in 2006 (the first year after the MGNREGA policy was enacted). Note, moreover, that the number of person-days under MGNREGA is normalized on district population in 2001 (drawn from Census data).

Figure 2.B.6: Baseline, when controlling for $\log(MGNREGA_{kt})$

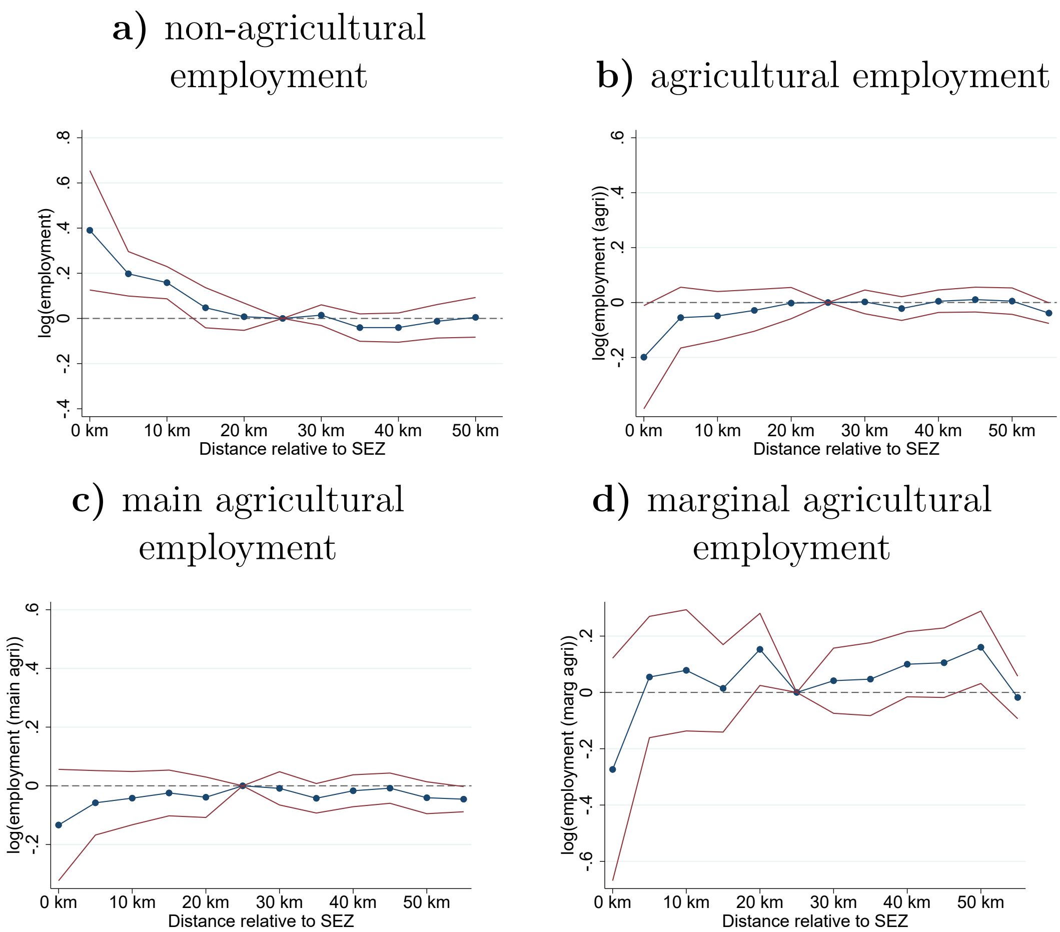

Notes: The dots indicate the estimated parameters $\hat{\beta}_d$ according to Eq. (2.6). Each d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Panel (a) depicts results for non-agricultural employment. Panels (b) - (d) show results for agricultural employment, respectively. Red lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. $MGNREGA_{kt}$ measures the ratio of person-days for in year t relative to the population at the district level in 2001.

Heterogeneous zone characteristics. In this part, we test whether the impact of SEZs on local employment hinges on the characteristics of the SEZ: the developer (public vs. private body) and the zone's industry denomination.

Table 2.B.1 presents estimates of Eq. (2.5) – where we compare privately and publicly developed zones – with and without reverting to matching. The results are similar to the baseline findings in Section 2.6.3. If anything, the point estimates suggest that employment effects are more pronounced for publicly developed zones, but the estimated effects are not statistically different from each other.

Next, we assess whether zone's industry denomination shapes SEZ's local employment effect, again in specifications with and without matching. The point estimates in Figure 2.B.7 and Table 2.B.2 indicate that multi-product zones have a higher local employment effect than other SEZs, but we cannot rule out statistically that they are different from employment effects of SEZs with other industry denominations.

Figure 2.B.7: Employment effects by SEZ industry

a) 0-10km (incl. SEZ municipalities): **b)** 10-20km:

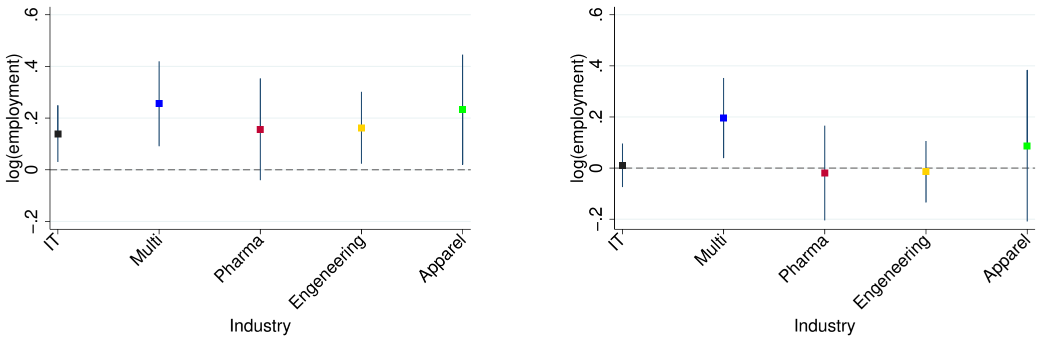

Notes: The plotted coefficients refer to $\hat{\beta}_d + \theta_d$ based on a variant of Eq.(2.5) as explained in section 2.6.3. Panel (a) depicts results for municipalities up to 10km away from their closest SEZ (incl. SEZ-municipalities). Panel (b) illustrates results for municipalities that are 10-20km away from their closest SEZ. Straight lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. For the purpose of giving a comprehensive picture of the full set of SEZ location choices across industries the industry sample includes all municipalities. Employment data are based on the Economic Census for the years 2005 and 2013.

The fact that zones with different industry denomination emerged in the wake of the 2005 SEZ Act (often developed by private investors) illustrates that firms with different underlying characteristics seem to reap benefits from locating within SEZs. SEZs in India offer various advantages to firms within their borders – most importantly, tax cuts on export income, benefits from being outside of the Indian tariff area as well as lighter regulation and less bureaucracy. The heterogeneity in zone denominations suggests that all of these factors may have pull and may be instrumental in attracting firms to SEZ areas.

To dig deeper and substantiate this claim, we draw on rich firm-level data on Indian companies, the so-called Prowess database, which is described in detail in Appendix 2.B.3. This data is matched to firms, which are active in SEZs, obtained from various publicly available lists (e.g. SEZ-developer webpages and members directory of the export promotion council), see Appendix 2.B.3 for further details. Based on this data, we determine the export intensity (measured by the value of firms' exports relative to overall sales), import intensity (measured by the value of firms' imports relative to overall sales) and the level of consultant fees (which serve as a proxy for the regulatory burden faced by the firms[34]) for firms in different industries (IT vs. manufacturing) and, within these industries, for firms in and outside of SEZ areas.

The findings are presented in Figure 2.B.8. They point to pronounced differences in underlying characteristics of IT and manufacturing firms and (within sectors) of firms within and outside of SEZ areas. Panel (a) suggests that both IT firms and manufacturing firms engage in export activity to a relevant extent, with rates being somewhat higher among IT firms. We further observe that firms in SEZs feature a higher export intensity than firms outside of these zones – consistent with firms with high-export intensity, which benefit overproportionally from tax reductions on export income, selecting into SEZ areas. Panel (b) shows that the import intensity starkly differs across IT and manufacturing firms: While imports make up less than 1% of sales for IT firms, this ratio stands at about 17% for manufacturing firms located in SEZs; and somewhat lower at 6% for manufacturing firms outside of SEZs. This suggests that manufacturing firms, contrary to their IT counterparts, find it attractive to locate in SEZ areas for tariff-related reasons (while both types of firms, to some extent, reap benefits from low tax rates on export income). And, again, consistent with firms' incentives, there seems to be selection of particularly import-intensive entities

[34]The idea is that firms subject to tighter regulation outsource a larger fraction of regulatory compliance work thus reporting higher consulting fees. It would have been ideal to combine information on external consultancy fees with information on internal costs incurred to comply with government regulations, but this information is not available, unfortunately.

into SEZ areas within the manufacturing sector. Panel (c) presents analogous evidence for consultancy fees, which serve as a proxy for firms' regulatory burden. We compare the costs between SEZ- and non-SEZ firms as well as between IT and manufacturing firms. On average, IT firms spent about five times as much on consultancy fees as their counterparts in the manufacturing sector, suggesting that they benefit overproportionally from a lighter regulatory burden within SEZs.[35]

In sum, the descriptive evidence in Figure 2.B.8 suggests that SEZ-benefits likely differ quite pronouncedly across IT and manufacturing industries. The fact that zones with different industry denomination emerged endogeneously and that the establishment of these zones is associated with broadly similar local treatment effects (i.e. impacts on local economic activity) serves as tentative evidence that different types of SEZ-benefits have pull and contribute to the attractiveness of SEZs for corporate activity.

Figure 2.B.8: Industry characteristics

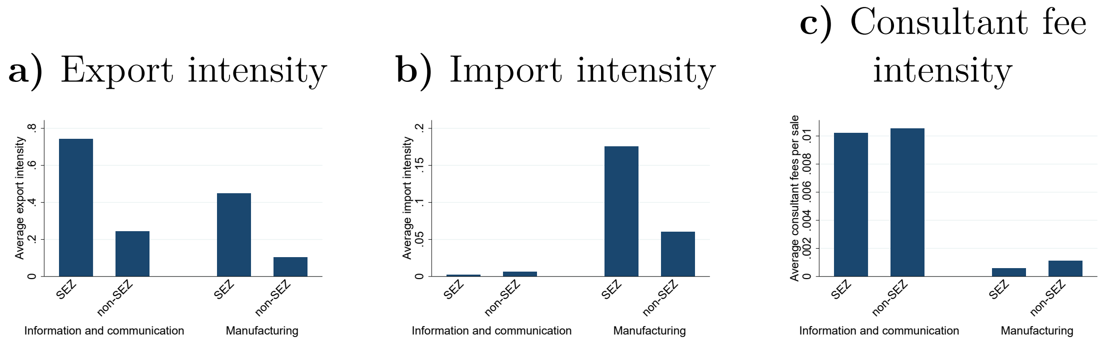

Notes: The figure depicts the export-intensity (exports over sales), import-intensity (imports over sales) and consultant fee intensity (consultant fees over sales) of firms in the IT and manufacturing industry, outside and inside of SEZs. SEZ firms are identified via fuzzy name matching. For details on the data, see Appendix 2.B.3.

2.B.2 Event study and nightlights

This appendix complements Section 2.5.1 in the main paper, where we present event study regressions based on annual nightlight data to corroborate the common-trend assumption.

[35] Similar evidence emerges from other data sources like the World Enterprise Survey, where firms are asked about their biggest obstacle to doing business. In the IT sector, 12.1% of firms state "business licensing and permits", which is significantly more than in any other sector besides construction. In manufacturing, only 4.3% state this to be the case. Also note that Panel (c) does not point to a within-industry selection of firms with particularly high regulatory costs towards SEZ areas. This may relate to a lack of intra-firm variation in regulatory burdens or to the imperfect nature of consultancy fees as a proxy for firms' regulatory and bureaucracy costs.

Figure 2.B.9: SEZ effect on nightlights by SEZ-industry

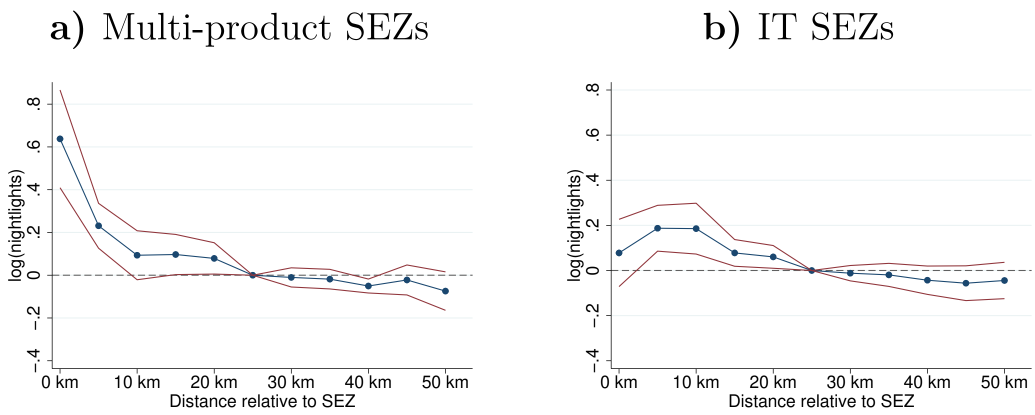

Notes: The dots indicate the parameters $\hat{\beta}_d$ as estimated by (2.1) Each subscript d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Red lines indicate 95%-confidence intervals. Panel (a) depicts the effect of multi-product SEZs on municipal nightlight intensity. Panel (b) depicts the effect of IT-SEZs on municipal nightlight intensity. Standard errors are clustered at the district level. Employment data based on the Economic Census for 1998, 2005 and 2013.

Figure 2.B.9 reestimates our baseline spatial difference-in-differences model with nighlight data, differentiating between multi-product and IT SEZs. The exercise confirms our baseline estimates and shows a positive treatment effect for SEZ hosting municipalities and municipalities in close proximity. Intuitively, the effect is particularly pronounced for multi-product SEZs, which, first, tend to be dominated by manufacturing firms with a high nightlight intensity and, second, tend to be located in more rural areas with low underlying nightlight intensity (making it easier to identify nightlight effects).

Figure 2.B.10 presents event study estimates for the impact of multiproduct SEZs on nightlight emissions. It compares municipalities treated by SEZs to reference locations as defined in the main text. The figure shows that nightlights emerged in parallel prior to SEZ establishment. After SEZ establishment, nightlight intensity increased significantly in treated relative to reference SEZs.

Figure 2.B.10: Nightlights in event study for multi-product SEZs

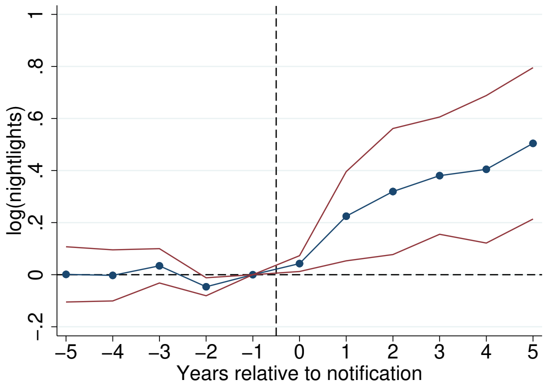

Notes: Event study estimates for municipalities hosting multi-product SEZs, municipalities in 20-25km distance serve as controls. The figure plots the $\hat{\theta}_k$ as estimated from Eq.(3.2) following Callaway and Sant'Anna, 2021. Endpoints are binned. Red lines indicate 95% confidence intervals. Standard errors are clustered at the district level.

2.B.3 Aggregate employment effects and relocation of economic activity

This appendix complements Section 2.5.2 in the main paper in two ways. First, we do a back-of-the-envelope calculation to obtain a rough idea about the magnitude of the aggregate employment effect of SEZs (in levels). Second, we offer robustness checks that explore whether and to what extent SEZs establish new economic activity or trigger relocation of economic activity in space. The two questions are intertwined as relocation of economic activity dampens the aggregate employment effect of the SEZ policy.

1. Back-of-the-envelope: Aggregate non-farm employment gain and aggregate relocation: In this subsection, we quantify the number of jobs that were established by SEZs in total within our sample frame. The analysis draws on our baseline estimates in panel (a) of Figure 2.5.1. They suggest that, for municipalities with a population below 500K, employment increased by 52%, 22%, and 16%, respectively, in SEZ-municipalities and municipalities in distance bins of 0-5km and 5-10km. Drawing on the average pre-treatment employment levels in SEZ-municipalities with less than 500K inhabitants in our sample of municipalities (3,139) and the indicated distance bins (574 and 439, respectively) and the total number of such municipalities per distance bin (152; 1,264 and 2,390), the aggregate effect of SEZs on municipalities within a 10km radius amounts to 575,598 additional workers ($= 0.52 \times 3,139 \times 152 + 0.22 \times 574 \times 1,264 + 0.16 \times 439 \times 2,390$).

We augment this number by the effects of SEZs on municipalities with a population of more than 500K, which are excluded from our baseline sample.[36] For these municipalities the estimated effect of SEZs on employment is smaller and estimated at 5% for SEZ-hosting municipalities, 7% for municipalities in a 0-5km distance and a small negative effect of -3% in municipalities in a 5-10km distance from SEZs. Again, considering the average pre-treatment employment levels in SEZ-municipalities with more than 500K inhabitants (666,796), the two closest distance bins (1,233,342 and 280,455, respectively) and the total number of such municipalities per distance bin (12; 4; 7), the aggregate effect of SEZs on municipalities within a 10km radius amounts to 680,102 additional workers.

Thus, overall employment in 10km radii around SEZs increased by about 1.25 million, which corresponds to an employment increase by 7.3% relative to the pre-treatment year 2005. Note that official statistics quantify the increase of employment within SEZs at 0.94 million over our period of study 2005-2013. Taken at face value, this suggests that 3/4 of the estimated net employment increase accrues within-SEZ municipalities and 1/4 of it reflects spillovers to surrounding regions (including SEZ municipalities themselves).[37]

In a second step, we use a back-of-the-envelope calculation to strengthen our argument in the main text that the observed estimates plausibly reflect the creation of new economic activity rather than job relocation in space. The results in Table 2.5.1 of the main text do not show any indication that the expansion of employment in SEZ areas correlates with declining employment paths in neighboring municipalities in further distance (> 10km, which would serve as 'source jurisdictions' in case of job relocation). The point estimates are small and statistically insignificant.

For distance rings smaller than 30km, the coefficient estimates nevertheless turn out negative. To obtain a notion of the quantitative relevance of these point estimates, we take the estimated 7.3% employment increase within a 10km-radius (see above), and calculate the aggregate employment decrease across municipalities in 10-30km distance rings from SEZs as implied by the point estimates in the first row of Table 2.5.1. We again evaluate the estimated coefficients at the average pre-treatment employment (624; 370; 310 and 226) and account for the number of municipalities (4,178; 4,334; 4,788 and 5,524) for the 10-15km, 15-20km, 20-25km and 25-30km distance bin, respectively. The total job loss

[36] We estimate the separate effect for large municipalities using interaction terms in a variant of Eq. (2.5).

[37] Figures are accessible via the Indian Export Promotion council: https://www.epces.in/facts-and-figures.php#hpgallery-6, last accessed: June 26th, 2022. Furthermore, one concern might be that SEZ-induced employment effects may systematically differ across municipalities of different size, which may potentially bias the specified aggregate employment response. Note that, if we split the set of municipalities with less than 500K inhabitants into four equally-sized bins and run our regression model separately in each subset of municipalities before aggregating, this yields very similar estimates for the aggregate employment effect (1.24 million new jobs).

calculated for these jurisdictions is 16,524 jobs, which is thus minuscule relative to the aggregate employment gain in SEZ areas (1.25 million workers).

As a word of caution, note, however, that the aggregate employment response calculated above hinges significantly on the response determined for SEZs in larger urban areas. This response is more difficult to determine than the response of smaller municipalities (see our discussion in the main text) and involves more uncertainty. Note that even if we abstract from SEZ-related job creation in larger urban areas altogether, the number of relocated jobs is still small relative to aggregate employment creation in SEZ-areas, namely 2.9% $(= 16,524/575,598)$. The bang-for-the-buck estimates in Section 2.7 change in turn. The costs per job created then are higher: 1,034,792 INR (17,118 USD in ppp per job) and the ratio of workers' wages to fiscal costs drops to 0.4.

2. Firm-level data. The data we use in the paper contains information about economic outcomes (e.g. employment), aggregated at the level of towns and villages. While this allows for a fine geographical resolution, information on firm-level characteristics is limited. To overcome this limitation to some extent, we constructed a firm-level dataset that offers more detailed information on the characteristics of firms located in SEZ areas.

To construct the dataset, we, first, retrieved the names of firms which are active in SEZs from various publicly available lists (e.g. SEZ-developer webpages and members directory of the export promotion council).[38] Second, we purchased the Prowess database compiled by the Centre for Monitoring of the Indian Economy (CMIE). Prowess contains detailed information on about 50,000 Indian firms, covering more than 70% of national industrial output from the organized sector, and it is widely used in empirical work (Barrows and Ollivier, 2021; Goldberg, Khandelwal, Pavcnik, and Topalova, 2010; Stiebale and Vencappa, 2018). Within this database, we identify SEZ firms through a *fuzzy matching* of firm names as appearing in the publicly available lists. With our procedure, we are able to identify 782 firms in the Prowess database that are active in SEZs. We believe that this data provides complementary information that is helpful in the context of our paper.

Prowess comes with two main drawbacks. First, it does not include disaggregated information on plant activity and plant locations (firm addresses are for entities' main location). The data is hence not well suited to model the spatial location of economic activity (contrary to the census data used in the main analysis). Second, the informal economy is not

[38] As an example, see here for the SEZs firm names in the members directors of the export promotion council on export oriented units and Special Economic Zones: `https://www.epces.in/members-directory.php`. Last accessed: June 2nd, 2024.

well represented in the Prowess data. As a convincing assignment of all (formal and informal) economic activity to a geographic entity is central for the credibility of our results, we cannot build on Prowess in the main empirical analysis, but the dataset is still helpful to generate additional insights.

Specifically, we use the data in two key ways in our analysis: First, we rely on the firm-level information to model key characteristics of SEZ firms (their export intensity, import intensity and external consultant fees), which shape firms' benefits from being active in an SEZ (see Appendix B.1 for details). Second, we use this data to augment the relocation analyses presented in the main text by assessing if firms relocate corporate activity within company groups towards SEZ areas. If relocation takes place from entities outside of the 50km radius, related relocation of corporate activity would not be captured in the relocation analysis presented in Section 2.5.2 of the main paper.

Specifically, we assess if entering an SEZ is associated with a reduction of economic activity at other locations of the same firm group outside of SEZ areas (defined as any group entity apart from the SEZ firm). The main analysis accounts for all firms, which belong to the same company group as the SEZ firm but are located outside of the 50km radius accounted for in our baseline analysis (where similar result patters emerge when we account for all non-SEZ firms in company groups with connections to an SEZ area, including entities within the 50km radius). The evolution of balance sheet items like sales, assets, profits, and the total wage bill is compared to firms in company groups without SEZ activity. Formally, we run the following regression:

$$\ln(y_{it}) = \sum_{k=2000, k\neq 2004}^{2013} \theta_k \mathbf{1}[t=k] + \gamma_t + \alpha_i + \epsilon_{it}, \tag{2.7}$$

where y_{it} denotes balance sheet information of firm i in year t. α_i and γ_t denote firm and year fixed effects, respectively. The θ_ks capture differences in the outcomes of interest across treated firms (belonging to groups that enter an SEZ) to control firms (belonging to groups without SEZ connection) across time, accounting for the time frame from 2000 to 2013, where the pre-SEZ reform year 2004 serves as base category.

The results are presented in Figure 2.B.11. The figure shows that firm outcomes emerged in parallel prior to the enactment of the SEZ policy. There is no indication for a drop in firm activity after companies entered SEZs at group locations outside of the SEZ area. This speaks against the notion that economic activity was shifted into SEZs across longer distances. On the contrary, the estimates point to an increase in economic activity at these locations after treatment, which is consistent with prior evidence, documenting that firm activity at different group locations are complements rather than substitutes (see e.g. Desai, Foley, and Jr., 2009, Becker and Riedel, 2012, Chodorow-Reich, Smith, Zidar, and Zwick, 2024).[39]

Figure 2.B.11: Intra-Group Relocation of Economic Activity (> 50km)

Notes: The figure determines the impact of SEZ establishment on non-SEZ firms within the same company group, comparing firms in groups that establish entities in an SEZ with company groups that have no SEZ activity. The sample is restricted to firms that are located beyond 50km from the closest SEZ entity within the same group. Only geocoded firms are included. 95% confidence intervals are displayed.

2.B.4 Structural change

This appendix complements Section 2.5.3 in the main paper. First, we calculate the aggregate reduction in agricultural employment and extend the back-of-the-envelope calcula-

[39]If group locations are e.g. connected through input-output-linkages, expanding investment at one group location is associated with higher investments at other group locations. Equivalently, the cost reductions may imply that it becomes more attractive for multinational firms to operate in India, which might enhance real economic activity at Indian group locations inside and outside of SEZs.

tion from Appendix 2.B.3, comparing this reduction to the observed increase in aggregate non-farm employment (in levels). This allows us to determine the fraction of the non-agricultural employment gain that is sourced from workers, who were previously employed in agriculture. The decline of agricultural employment is also a relevant ingredient for the cost-effectiveness calculation in Section 2.7. Second, we provide evidence for pronounced employment gains in low-skilled service industries in SEZ areas and surroundings, which further supports the narrative that low-skilled workers transit from agriculture to non-farm sectors.

1. Aggregate decline in agricultural employment: In our baseline sample of municipalities with a population below 500K, agricultural employment decreased by 17%, 4%, and 5%, respectively, in SEZ-municipalities and municipalities in distance bins of 0-5km and 5-10km. Drawing on the average pre-treatment agricultural employment levels in SEZ-municipalities with less than 500K inhabitants in our sample of municipalities (777), the indicated distance bins (571 and 532, respectively) and the total number of such municipalities per distance bin (152; 1,264 and 2,390), the aggregate negative effect of SEZs on municipalities within a 10km radius amounts to 112,521 less agricultural jobs ($= -0.17 \times 777 \times 152 - 0.04 \times 571 \times 1,264 - 0.05 \times 532 \times 2,390$). To account for the effect in large municipalities, too, we estimate effects separately and follow the same aggregation procedure as in the base analysis. Note, however, that the estimated effects for this sample should be interpreted with caution as they are based on a small sample of large municipalities. The decrease in agricultural jobs in large municipalities amounts to 291,740 ($= -0.45 \times 31377 \times 12 - 0.68 \times 22916 \times 4 - 0.65 \times 13181 \times 7$). Overall, our back-of-the-envelope calculation suggests that SEZs reduced agricultural employment by 404,262 jobs. Comparing this figure to the estimated increase in non-agricultural employment from Appendix 2.B.3, we find that around every third new non-agricultural job (32%) is sourced from the agricultural sector.

Note that similar findings emerge when we refine the analysis and allow for heterogeneity in the effect of SEZs on municipalities of different size. In particular, we again divide the sample of municipalities with less than 500,000 inhabitants into four equally sized subsamples according to population quartiles. We then reestimate our baseline model in these subsamples and rerun the aggregation exercise above for each of the four samples and for large municipalities in addition. We do so for non-agricultural employment as well as for agricultural employment. This leaves the ratio of the aggregate agricultural employment decrease and the non-agricultural employment increase unchanged, at 32.2%.

2. Characteristics of new employment. We also shed light on the type of jobs that emerged in and around SEZs. We turn to the Economic Census for this exercise, which contains detailed 3-digit industry codes. As the industry classification changed within our sample frame, we use the concordance tables provided by the Ministry of Statistics and Programme Implementation to harmonize the National Industry Classification (NIC) of 2008 – which is used in the Economic Census of 2013 – and the NIC codes of 2004 and 1987 – which is used in the Economic Censuses of 2005 and 1998. In cases of industry splits, we assign the industry code, which has a higher employment share according to the Economic Census of 2013. We draw on this data to classify industries in the service sector into low-skilled industries (e.g. restaurants) and high-skilled service industries (e.g. information technologies, communication, or management). The evidence is presented in Figure 2.B.12. The figure suggests that employment in industries classified as 'low-skilled' increases pronouncedly, see panel (a). Employment in industries classified as 'high-skilled' shows small, but insignificant increase, in turn, see panel (b). An exception is employment in the IT-sector, which – albeit high-skilled – rises strongly after SEZ establishment, likely reflecting the importance of this worker group for SEZ activities. In sum, the evidence supports the narrative that the decline in (marginal) agricultural employment reflects that lower-skilled workers transitioned from the agricultural sector to lower-skilled service jobs (additionally to manufacturing).

Figure 2.B.12: Service employment (high- vs. low-skilled)

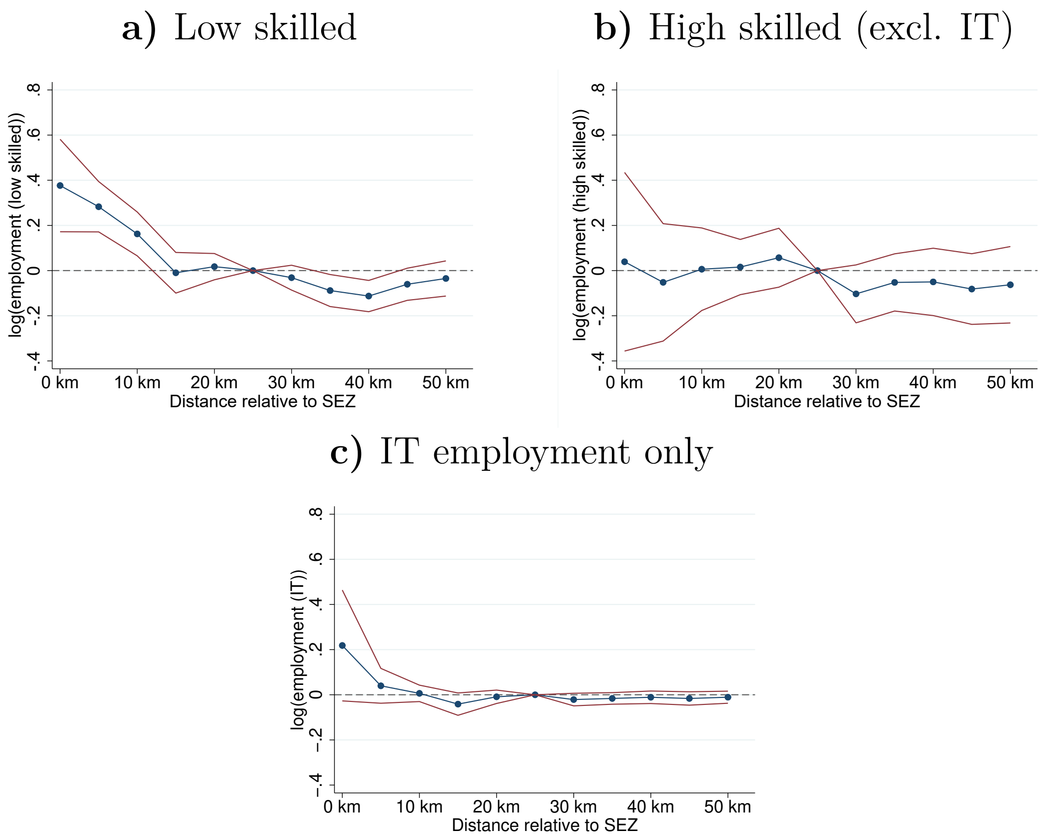

Notes: We characterize industries as low-skilled or high-skilled based on on the National Industry Classification (NIC) 2008 as follows: NIC codes IT: 581-639. NIC codes high skilled: 641-750. NIC codes low skilled: 451-563. The dots indicate the estimated parameters $\hat{\beta}_d$. Each subscript d refers to a distance on the horizontal axis, e.g. the coefficient at 0km refers to $d = 0$. Red lines indicate 95%-confidence intervals. Standard errors are clustered at the district level. Employment data are based on the Economic Census for 2005 and 2013.Red lines indicate 95%-confidence intervals.

2.B.5 Additional outcomes

Local public goods. In this section, we explore whether the SEZ Act led to higher provision of local public goods, e.g. streets or electricity infrastructure, that benefited local residents (which was one goal of the SEZ policy, see Section 2.2). The population census allows us to shed some light on local public good provision. We observe the number of schools in each municipality and whether a municipality had access to any kind of electricity or to a paved road, respectively. Re-estimating Eq. (2.1) with these different dependent variables does not point to any SEZ-induced improvements in electricity and road access. The number of schools slightly increased in treated municipalities after SEZ establishment (relative to municipalities in further distance). This positive effect vanishes, however, when we normalize the number of schools on population size. Finally, we find no effect on local literacy rates, see Panel (d).

Figure 2.B.13: SEZ effect on local infrastructure and literacy

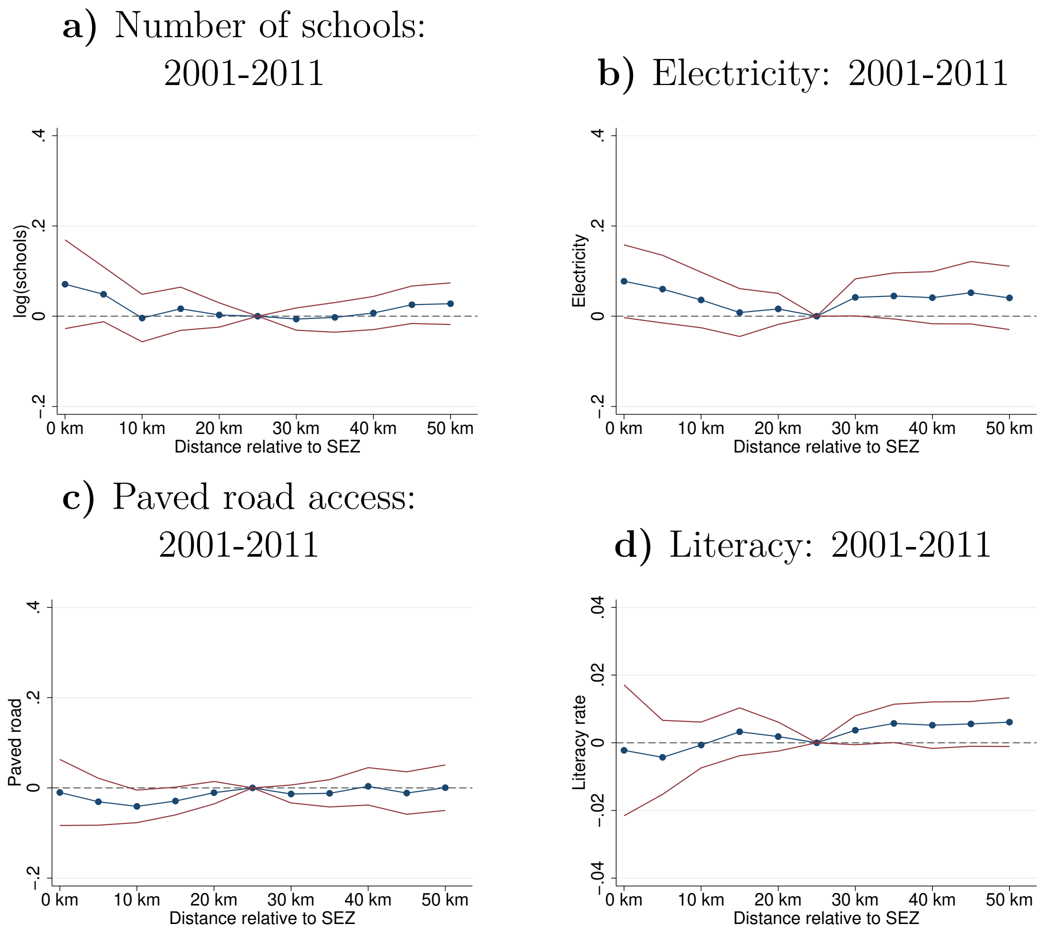

Notes: The dots indicate the estimates for $\hat{\beta}_d$ as estimated according to Eq. (2.1). Each d refers to a distance on the horizontal axis e.g. the coefficient at 0km refers to $d = 0$. Panel (a) depicts results for the number of schools. Panel (b) depicts results for electricity access. Panel (c) depicts results for paved road access. Panel (d) depicts the results for the literacy rate. Red lines indicate 95% confidence intervals. The standard errors are clustered at the district level. Data are based on the Population Census for the years 2001 and 2011. Hence, only municipalities that are within 50km of SEZs that became operational until 2011 are included.

Firm entry. We have shown in the main part of the paper that the SEZ Act led to more employment in SEZ-hosting and neighboring municipalities. This part complements these insights by exploring the extensive margin, that is the change in the number of firms through entry or exit. We show in Column (1) of Table 2.B.3 that the policy led to a strong positive response at the extensive margin in SEZ-hosting municipalities and their neighbors up to 10km. The placebo regressions in Column (2) point to no differences in pre-treatment trends. We further document in Columns (3)-(6) that the increase in the number of firms was primarily driven by male firm ownership and by small firms.

Table 2.B.3: SEZ effect on firm entry

Distance bins	(1) *Total*	(2) *Placebo*	(3) *Male*	(4) *Female*	(5) *Large*	(6) *Small*
0km	**0.296*** **	**-0.100**	**0.390*** **	**0.094**	**0.063**	**0.314*** **
	(0.112)	(0.106)	(0.114)	(0.147)	(0.219)	(0.116)
0-5km	**0.204*** **	**0.002**	**0.260*** **	**0.108**	**-0.176**	**0.210*** **
	(0.045)	(0.061)	(0.057)	(0.092)	(0.107)	(0.045)
5-10km	**0.142*** **	**-0.024**	**0.176*** **	**0.099**	**-0.049**	**0.144*** **
	(0.035)	(0.059)	(0.044)	(0.077)	(0.123)	(0.037)
10-15km	**0.056**	**0.001**	**0.078**	**0.009**	**-0.156**	**0.059**
	(0.049)	(0.040)	(0.053)	(0.052)	(0.101)	(0.051)
15-20km	**-0.009**	**-0.010**	**0.032**	**-0.001**	**-0.032**	**-0.010**
	(0.027)	(0.028)	(0.030)	(0.046)	(0.058)	(0.028)
20-25km	–	–	–	–	–	–
25-30km	**-0.005**	**0.023**	**0.006**	**-0.037**	**-0.111** **	**-0.008**
	(0.025)	(0.028)	(0.033)	(0.037)	(0.056)	(0.025)
30-35km	**-0.058* **	**0.047**	**-0.057**	**-0.009**	**-0.094**	**-0.060* **
	(0.032)	(0.034)	(0.038)	(0.048)	(0.081)	(0.033)
35-40km	**-0.056* **	**0.027**	**-0.063**	**-0.007**	**-0.144* **	**-0.058* **
	(0.032)	(0.030)	(0.039)	(0.053)	(0.079)	(0.033)
40-45km	**-0.024**	**-0.002**	**-0.034**	**0.025**	**-0.125* **	**-0.025**
	(0.035)	(0.035)	(0.042)	(0.057)	(0.068)	(0.036)
45-50km	**-0.008**	**-0.044**	**-0.015**	**0.031**	**-0.183** **	**-0.009**
	(0.037)	(0.037)	(0.042)	(0.060)	(0.072)	(0.037)
Observations	92,926	84,120	85,216	36,888	16,712	92,828
R-squared	0.905	0.900	0.883	0.841	0.842	0.904
Municipality fixed effects	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results from Eq. (2.1) with the number of different types of firms as the dependent variable. Column (1) reports the estimated effects on total firm count. Column (2) reports the placebo results. Columns (3)-(6) report the results for male owned-, female owned-, large- and small firm count. Data are based on the Economic Census for the years 1998, 2005 and 2013. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.B.1: Employment effects by developer

	(1)	(2)	(3)	(4)	(5)	(6)
			Employment			
Matching	None		Industry		Industry & size	
Distance bins	*Private*	*Public*	*Private*	*Public*	Private	*Public*
0km	**0.314****	**0.549*****	**0.382*****	**0.411****	**0.369****	**0.400****
	(0.159)	(0.209)	(0.146)	(0.180)	(0.150)	(0.158)
0-5km	**0.125****	**0.357*****	**0.130****	**0.266*****	**0.122****	**0.263*****
	(0.056)	(0.080)	(0.056)	(0.089)	(0.056)	(0.089)
5-10km	**0.120*****	**0.202*****	**0.121*****	**0.191****	**0.121*****	**0.200****
	(0.040)	(0.064)	(0.040)	(0.076)	(0.041)	(0.077)
10-15km	**0.025**	**0.061**	**0.027**	**0.021**	**0.021**	**0.017**
	(0.051)	(0.063)	(0.051)	(0.072)	(0.052)	(0.070)
15-20km	**-0.009**	**0.030**	**-0.007**	**0.012**	**-0.015**	**0.012**
	(0.039)	(0.046)	(0.038)	(0.045)	(0.042)	(0.045)
20-25km	–	–	–	–	–	–
25-30km	**-0.023**	**0.070***	**-0.022**	**0.034**	**-0.028**	**0.033**
	(0.032)	(0.038)	(0.032)	(0.043)	(0.032)	(0.043)
30-35km	**-0.091****	**0.051**	**-0.092****	**0.074***	**-0.088****	**0.080****
	(0.040)	(0.059)	(0.040)	(0.042)	(0.040)	(0.040)
35-40km	**-0.075***	**0.024**	**-0.078***	**0.024**	**-0.079***	**0.028**
	(0.044)	(0.058)	(0.044)	(0.050)	(0.044)	(0.051)
40-45km	**-0.054**	**0.055**	**-0.057**	**0.074**	**-0.057**	**0.067**
	(0.051)	(0.043)	(0.050)	(0.045)	(0.051)	(0.046)
45-50km	**-0.069**	**0.123****	**-0.073**	**0.072**	**-0.076**	**0.084**
	(0.049)	(0.054)	(0.048)	(0.067)	(0.048)	(0.065)
Observations	92,980	92,980	92,954	92,954	91,960	91,960
R-squared	0.899	0.899	0.919	0.919	0.919	0.919
Municipality fixed effects	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results based on Eq. (2.5) contrasting employment effects of public and private SEZs. Columns (1)-(2) report results without matching. In columns (3)-(4), we match on industries as in Figure 2.6.4. Columns (5)-(6) show results when municipalities are matched according to SEZ-industry and SEZ-area relative to municipality area. Employment data are based on the Economic Census for the years 2005 and 2013. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.B.2: Employment effects by SEZ industry

	(1)	(2)	(3)	(4)	(5)	(6)
			Employment			
Matching	None		Developer		Developer & size	
Distance bins	*Multi*	*IT*	*Multi*	*IT*	Multi	*IT*
0km	**0.625*** **	**0.420*** **	**0.544** **	**0.416*** **	**0.641***	**0.413*** **
	(0.186)	(0.157)	(0.210)	(0.157)	(0.348)	(0.157)
0-5km	**0.324*** **	**0.114**	**0.280** **	**0.112**	**0.239** **	**0.106**
	(0.117)	(0.075)	(0.122)	(0.075)	(0.113)	(0.076)
5-10km	**0.139***	**0.168*** **	**0.119**	**0.165*** **	**0.077**	**0.164*** **
	(0.075)	(0.057)	(0.075)	(0.057)	(0.083)	(0.057)
10-15km	**0.164**	**0.061**	**0.103**	**0.061**	**0.040**	**0.056**
	(0.124)	(0.067)	(0.088)	(0.067)	(0.083)	(0.068)
15-20km	**0.157***	**-0.011**	**0.171***	**-0.012**	**0.097**	**-0.012**
	(0.092)	(0.038)	(0.093)	(0.038)	(0.071)	(0.038)
20-25km	–	–	–	–	–	–
25-30km	**0.020**	**-0.002**	**0.039**	**-0.001**	**-0.006**	**0.000**
	(0.046)	(0.039)	(0.042)	(0.039)	(0.042)	(0.039)
30-35km	**-0.024**	**-0.016**	**-0.039**	**-0.015**	**0.007**	**-0.015**
	(0.057)	(0.053)	(0.065)	(0.053)	(0.065)	(0.053)
35-40km	**-0.044**	**-0.041**	**-0.075**	**-0.039**	**-0.111** **	**-0.038**
	(0.062)	(0.052)	(0.069)	(0.052)	(0.052)	(0.052)
40-45km	**0.030**	**-0.014**	**0.005**	**-0.012**	**-0.030**	**-0.011**
	(0.048)	(0.058)	(0.051)	(0.059)	(0.065)	(0.058)
45-50km	**0.155** **	**0.005**	**0.127***	**0.007**	**0.112**	**0.012**
	(0.073)	(0.061)	(0.076)	(0.061)	(0.078)	(0.061)
Observations	51,202	51,202	51,202	51,202	50,414	50,414
R-squared	0.898	0.898	0.898	0.898	0.899	0.899
Municipality fixed effects	✓	✓	✓	✓	✓	✓
Year fixed effects	✓	✓	✓	✓	✓	✓

Notes: Regression results based on Eq. (2.5) with $industry_i$ instead of $priv.developer_i$ as an identifier. CEM is applied with IT being the treatment category. Columns (1)-(2) report the results without matching. Columns (3)-(4) show results when municipalities are matched according to SEZ developer (public or private) as in Figure 2.6.4. Columns (5)-(6) report results when municipalities are matched according to SEZ-developer and SEZ-area relative to municipality area. The sample includes all municipalities. Employment data are based on the Economic Census for the years 2005 and 2013. Standard errors are clustered at the district level. *** p<0.01, ** p<0.05, * p<0.1.

Chapter 3

Carbon Taxation and Firm Behavior in Emerging Economies: Evidence from South Africa

Joint with Johannes Gallé, Rodrigo Oliveira, Nadine Riedel and Edson Severnini.

3.1 Introduction

Reducing carbon emissions is essential for combating climate change. Although this is widely recognized and prioritized on policy agendas, global carbon emissions continue to rise (Global Carbon Budget, 2023). Among the various policy tools available to curb emissions, carbon pricing is widely regarded as the most efficient approach (Gordon, 2023; Marron and Toder, 2014; Timilsina, 2022). Yet, implementing carbon pricing – particularly carbon taxes – poses challenges for emerging markets and less developed countries. Many low and middle income countries (LMICs) have so far been reluctant to embrace carbon taxation. One main concern is that a carbon tax may come with negative economic effects and hinder economic development (Strand, 2020). LMICs' relatively strong reliance on carbon-intensive energy, their lower levels of technological development and adaptive capacity, and oftentimes weak economic resilience may render them particularly vulnerable to negative economic effects of carbon taxation (Marron and Toder, 2014; Metcalf, 2021). This chapter provides the first comprehensive analysis of how firms in an emerging economy respond to carbon taxation. It focuses on South Africa, the 13^{th}-largest carbon emitter globally and the first African nation to introduce a nationwide carbon tax in 2019. As a potential trailblazer for other developing countries, South Africa's experience is particularly relevant given the growing global emphasis on carbon pricing, with 65 carbon pricing schemes worldwide and 46 more underway (UNFCCC, 2024). While carbon pricing can in general take the form of carbon taxes or emissions trading schemes, carbon taxes are easier to implement, which may make them relatively more attractive for LMICs.[1] To make it politically feasible, South Africa introduced its carbon tax with substantial allowances during the initial phase to ease the transition.

Using novel and comprehensive administrative tax data, we analyze the dynamic effects of the policy on firm behavior. Our findings show that the carbon tax did not negatively affect key firm outcomes such as sales, profits, or employment. Notably, the announcement of the tax – made four years before its implementation – appears to have spurred an increase in firm activity, suggesting anticipatory behavior. The anticipation led firms to increase (potentially emission intensive) economic activity prior to the implementation of the tax, after which these activities become more costly.

To quantify the economic impacts of the carbon tax, we examine firm responses to both its announcement and implementation. Our analysis focuses on manufacturing and mining firms, which are key sectors in the context of the carbon tax. Using matching techniques, we address imbalances in observable pre-treatment characteristics between taxed and non-taxed firms. Our matched sample performs well in terms of trend comparisons between treated and untreated firms, and achieves a matching ratio comparable to a recent study evaluating the EU ETS (Colmer, Martin, Muûls, and Wagner, 2024). Based on this matched sample, we estimate event-study regressions to capture the dynamic effects of the tax announcement (draft bill released in 2015) and its implementation (enacted in 2019).

The analysis leverages the universe of corporate income tax returns filed from 2011 to 2021, allowing us to track firms for six years after the release of the draft carbon tax bill and two years following the tax's implementation in 2019.[2] Beginning in 2019, we integrate detailed data on emissions and tax payments with other firm-level tax records. This enables an in-depth analysis of the carbon tax base, which we find covers over 80% of nationwide

[1] When relying on a carbon tax, countries can leverage existing tax authorities and target large emitters. But they may face political resistance due to public aversion to new taxes. ETS, on the other hand, are more politically feasible as they allow free distribution of emissions permits to firms, but require establishing and regulating a market, often necessitating a new agency, which can be challenging for developing countries due to limited state capacity.

[2] The fiscal year in South Africa begins in March and ends in February of the following year. Therefore, the 2020 fiscal year started in March 2019 and ended in February 2020. This timing means that our estimated effects for 2020 are largely influenced by the implementation of the carbon tax (introduced in June 2019) but remain unaffected by the COVID-19 pandemic, which impacted the subsequent fiscal year.

emissions. Although the statutory tax rate is uniform across industries, we identify significant variation in *effective* tax rates – the amount firms actually pay per tonne of CO_2. Our descriptive analysis reveals a non-linear relationship between emissions and tax payments, driven by the tax's design, which includes allowances, exemptions, and special provisions for certain industries.

The main results indicate that the announcement of the Carbon Tax Bill in 2015 led to increases in sales, capital, employment, and profits. The effects over time are consistent across manufacturing and mining sectors, except for capital, which exhibits a more muted reaction in the mining sector. These positive effects grew gradually and remained significant even after the carbon tax was implemented in 2019. For example, by 2019, sales, capital, employment, and profits had increased by approximately 20%. Contrary to concerns that carbon taxes may hinder economic growth, our findings suggest no negative impacts on firm performance on average within the first four years of the tax's implementation. Using a synthetic control approach, however, we provide suggestive evidence that the increased firm activity following the tax announcement coincided with a temporary rise in carbon emissions.

To examine how differences in exposure to the carbon tax influence firm responses, we conduct a heterogeneity analysis. Our findings reveal that firms with fewer allowances – and therefore facing higher *effective* tax rates – experienced significant increases in sales and employment following the tax announcement. In contrast, firms with more allowances showed more muted responses, with estimated effects on sales and employment that are less pronounced and, in some cases, statistically indistinguishable from zero after the tax's implementation. These results might sound counterintuitive because one would typically expect that firms facing higher effective tax rates (due to fewer allowances) would experience negative economic outcomes, given the increased cost burden.

We explore two mechanisms that could explain these seemingly counterintuitive results. First, the four-year gap between the release of the carbon tax bill and its implementation allowed firms to manage uncertainty, not only by gaining clarity on the future cost of production but also by confirming that there would be no cap on emissions, as long as the tax was paid. This resolution of uncertainty may have enabled firms to plan strategically for the transition. Our findings show that firms with higher exposure to the carbon tax experienced increased sales and significantly higher depreciation following the announcement of the bill. These results suggest that these firms not only used their capital more intensively but also accelerated the depreciation of their emission-intensive machinery in anticipation of the tax. The accelerated depreciation, in particular, likely reflects an effort to mitigate the risk of stranded assets – emission-intensive machinery that could become uneconomical or obsolete under the new tax regime.

As a second potential mechanism, we explore whether the carbon tax incentivized firms to upgrade their production technology. While we find no evidence of increased investments in R&D, treated firms did slightly increase imports between 2018 and 2020. This could suggest they sourced more production inputs from abroad, either to replace domestically taxed inputs or to upgrade their technology. Notably, our analysis shows that, on the intensive margin, treated firms primarily imported products within the same categories as in previous years. However, on the extensive margin, they also introduced new products, indicating some level of technology or product innovation. While we cannot completely rule out technology upgrading, our evidence suggests that this channel is less significant in explaining the substantial effects of the tax announcement and implementation.

This study makes three main contributions to the literature. First, it contributes to the literature on the economic impacts of carbon pricing, which has predominantly focused on cap-and-trade schemes and/or developed world settings (Andersson, 2019; Bushnell, Chong, and Mansur, 2013; Calel and Dechezlepretre, 2016; Colmer, Martin, Muûls, and Wagner, 2024; Cui, Wang, Zhang, and Zheng, 2021; Cui, Zhang, and Zheng, 2023; Dechezlepretre, Nachtigall, and Venmans, 2023; Martin, de Preux, and Wagner, 2014; Martin, Muuls, Preux, and Wagner, 2014; Yamazaki, 2017, 2022). In contrast, we estimate the impacts of a carbon tax in an emerging economy context, where carbon taxes may be more favorable than cap-and-trade schemes due to limited state capacity. While in both cases the state needs to know firms' emissions, carbon tax administration can rely on the existing tax system while cap-and-trade systems require establishing and regulating a market, often necessitating a new agency.

Second, it contributes to the literature on environmental policy and labor market outcomes, particularly in the context of politically contentious carbon taxes. Market-based approaches, such as carbon pricing, are widely regarded as more cost-effective than command-and-control regulations (Carlson, Burtraw, Cropper, and Palmer, 2000; Fowlie, Holland, and Mansur, 2012).[3] While command-and-control policies have been shown to reduce employment and earnings in developed economies (Greenstone, 2002; Walker, 2013), market-

[3] Yet carbon taxes often face significant public resistance (Anderson, Marinescu, and Shor, 2023; Douenne and Fabre, 2022; Ewald, Sterner, and Sterner, 2022). This resistance is often tied to the tax's distributional implications, which can be politically challenging to defend (Fried, Novan, and Peterman, 2022; Känzig, 2023; Steckel, Dorband, Montrone, Ward, Missbach, Hafner, Jakob, and Renner, 2021).

based approaches, including carbon taxes, generally do not appear to adversely affect employment (Martin, de Preux, and Wagner, 2014; Yamazaki, 2017). Our study aligns with this developed world evidence by demonstrating that flexible environmental policies, such as carbon taxes, do not seem to harm economic activity and employment, even in the context of a less developed country with a higher degree of informality.

Third, it contributes to the literature on anticipatory firm behavior, particularly the so-called "green paradox," where environmental policies may temporarily worsen environmental outcomes due to pre-emptive actions by firms (Di Maria, Lange, and van der Werf, 2014; Lemoine, 2017; Lueck and Michael, 2003), unless the output is not storable (Clay, Jha, Lewis, and Severnini, 2024). We provide evidence that a carbon tax, even in a setting without emissions caps in the short run, may induce such behavior. This appears to arise from firms resolving uncertainty or seeking to recover costs from stranded assets, offering new insights into how carbon taxes might influence firm behavior in the absence of emissions caps.

The remainder of the paper is organized as follows: Section 3.2 provides an overview of the institutional background. Section 3.3 introduces the data sources used in the empirical analysis. Section 3.4 details the empirical methodology, while Section 3.5 presents and discusses results. Finally, Section 3.7 offers concluding remarks.

3.2 Institutional background

South Africa is a middle-income country with a GDP per capita of USD 7,055 and a tax-to-GDP ratio of 21% in 2023. It is one of the few countries in the world to have adopted a carbon tax and, to date, the only African country to have implemented any sort of carbon pricing scheme (World Bank, 2024). Appendix Figure 3.A.1 shows that, in 2023, the carbon tax increased the government revenue by about ZAR 1.5 billion.[4] This section explains the institutional setting and the legal framework of the carbon tax.

Policy process. The implementation of the South African carbon tax in 2019 has been preceded by a long political process in which a variety of private and public stakeholders discussed and pushed their interests. Figure 3.2.1 shows the timeline of the main events until the Carbon Tax Act in 2019. In 2010, the Carbon Tax Discussion Paper marks the first

[4] This amounts to 0.1 % of total tax revenues in Soth Africa. To put this figure into perspective, it is comparable to the increase in the early childhood development grant. For details, see `https://shorturl.at/taDvq`. Last accessed: March 11th, 2025.

official mentioning of a ‘carbon tax’. The paper outlined potential ways a carbon policy could be designed and already signaled the clear intention to regulate carbon emissions in South Africa. Three years later, the Carbon Tax Policy Paper was a more concrete proposal, still, however without detailing the final framework. It was not until the end of 2015 that the first draft of the Carbon Tax Bill was published. The draft detailed the various allowances for different sectors and activities. Importantly, it also included the schedule of GHG emission factors, i.e., if unaware before, firms learned at that point how much emissions they produced and consequently how much they were going to pay. Finally, the Carbon Tax Act was implemented in 2019 and became effective in June 2019. The empirical analysis will focus on these latter two events, which we interpret as the *announcement* and the *implementation* of the carbon tax.

The process of the carbon tax enactment was accompanied by much resistance from affected industries (Baker, 2022). In addition, the policy process was open for public comments, allowing affected industries or interest groups to actively engage and participate in the discussions around the design of the carbon tax policy. The opportunity for public comment was heavily used – Appendix 3.C shows the distribution of comments from various stakeholders, ranging from industry associations to individuals. The content of the comments ranged from demanding further clarifications to substantial criticism and contesting of the the carbon tax. While some comments were taken into consideration, the final Carbon Tax Act still featured all the main elements proposed in the initial Carbon Tax Bill in 2015. Nevertheless, this heavy public engagement illustrates the pressure the legislation was facing in pushing the carbon tax through. It may have led to the initial design of granting various allowances and exemptions meant to attenuate the alleged negative consequences of the tax.

Exemptions and allowances. In principle, the carbon tax applies to all firms emitting CO_2 if the extent of their polluting activity exceeds a certain threshold as defined in the Carbon Tax Act, 2019. These thresholds can be industry- or activity-specific. For example, the threshold for firms with combustion activities was set at 10 MW of installed thermal input capacity. This means that regardless of utilization or fuel type, if a firm has the capacity to combust 10 MW(th), then its emissions will be subject to the carbon tax. If the firm has a smaller capacity, carbon taxes do not apply. This ensures that the smallest firms are exempted altogether.

Initially, the first phase of the carbon tax was set to run from June 1, 2019, to December

Figure 3.2.1: Timeline of the main events until the Carbon Tax Act in 2019

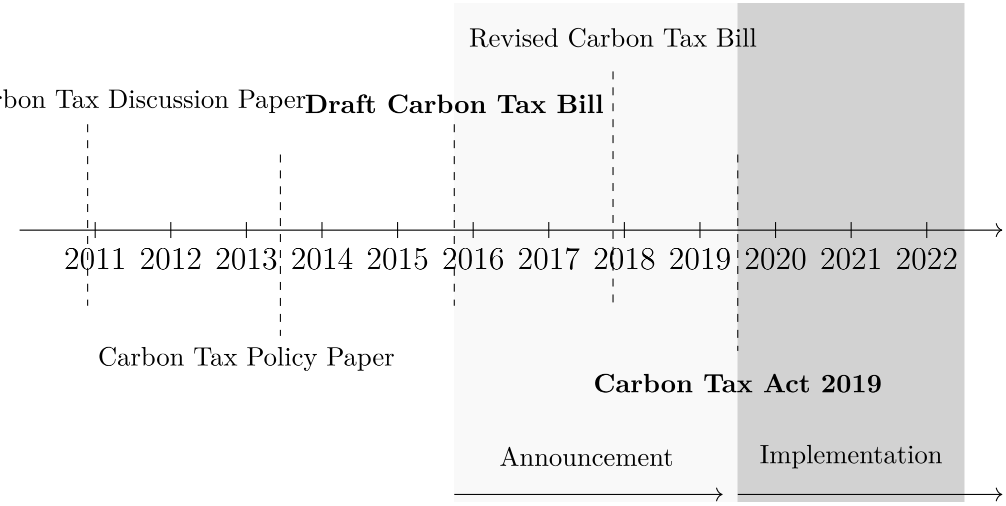

Notes: This figure plots the main events between the Carbon Tax Discussion Paper and the Carbon Tax Act in 2019.

31, 2022. However, the government extended this phase by three years to support economic recovery in the wake of the COVID-19 pandemic. At the time of the introduction of the carbon tax in 2019, the statutory tax rate amounted to \$120 ZAR ($\sim$ \$7 USD) per tonne of CO_2. Despite being relatively low compared to other carbon pricing policies (Timilsina, 2022), the statutory tax rate increases annually by inflation plus 2 percent for the first few years. As of 2024, the statutory tax rate amounts to \$190 ZAR ($\sim$ \$11 USD). This low rate is expected to increase after the end of the transition phase, in December 31, 2025. The government plans to raise the carbon tax rate to at least US\$ 20/tCO_2 by 2026, to US\$ 30/tCO_2 by 2030, and accelerating to higher levels up to US\$ 120/tCO2 beyond 2050 (Qu, Suphachalasai, Thube, and Walker, 2023).[5]

Having said that, in the first few years following the introduction of the carbon tax, the effective tax rate could be reduced further through various allowances, as part of the transition from the initial phase of the carbon tax program. First of all, for the majority of sectors, there is a basic tax-free allowance of at least 60% of emissions, which means that only 40% of firms' emissions are taxed. Additionally, there are specific allowances for fugitive emissions, for the extent of trade exposure, performance allowances (for firms emitting less than their industry-specific standard) as well as carbon budget and offset

[5]See more details at `https://tinyurl.com/54xwt25z`. Last accessed: March 6th, 2025.

allowances. These last two types of allowances refer to credits for voluntarily participating in the carbon budget program or for purchasing carbon offsets, provided the offsets are generated within the country. As a result of the tax-free allowances, which can total up to 95%, the effective rate could be as small as $6 ZAR (US$ 0.40) per ton of CO_2 emissions (Steenkamp, 2022). The low rate aimed at allowing large emitters enough time to transition to clean technologies, but are set to be phased out at the end of the transition phase in December 31, 2025.

In addition to the carbon tax allowances, firms in South Africa can deduct the petroleum and diesel levies, as well as the electricity levy, from their carbon tax liabilities. These levies are pre-existing fiscal tools that interact with the carbon tax policy, all aiming at managing environmental impacts and promote energy efficiency. The petroleum and diesel levies, introduced in 2003, are charged on petrol and diesel to fund road infrastructure and public transport systems, while also incentivizing fuel efficiency. The electricity levy, implemented in 2009, is applied to electricity consumption to reduce reliance on coal-fired power generation and support the transition to cleaner energy sources. Alongside the carbon tax allowances, these levies are deductible, reducing the overall carbon tax liabilities.[6]

As a result, the Net Emissions Equivalent (NEE), which represents the emissions subject to the carbon tax after accounting for allowances and deductions, is calculated as:

$$NEE = ((E - S) \times (1 - C)) - (D \times (1 - M)) + P \times (1 - J) + F \times (1 - K), \tag{3.1}$$

where E refers to all fuel combustion-related emissions of a taxpayer, from which sequestrated emissions S may be subtracted. C represents the sum of allowances applicable to fuel combustion activities. D corresponds to the CO_2 emissions from petrol and diesel. Since petrol and diesel are already subject to a fuel levy, their emissions are multiplied by their respective allowances M and subtracted, effectively exempting them from carbon taxation to avoid double taxation. P represents industrial process-related CO_2 emissions, and J denotes the corresponding allowances that can be deducted. Lastly, F refers to fugitive emissions, and K represents the applicable allowances for these emissions.

The net emissions, as calculated above, still do not directly correspond to actual carbon tax payments, as certain expenditures on other taxes can be credited against the carbon tax liability. For instance, as mentioned above, electricity providers can deduct the costs incurred from complying with the electricity levy. Thus, electricity providers are effectively

[6] As stated by South Africa Revenue Services, the fuel and electricity levies are excluded from the carbon tax calculation to avoid double taxation (South African Revenue Service, 2021).

“exempt" from the carbon tax during its current introductory phase.

3.3 Data

Our analysis uses detailed administrative tax data from the South African Revenue Service (SARS), accessed confidentially through the South African National Treasury. The individual data components are described below, but all three data sources can be linked using anonymized tax reference numbers provided by SARS.

Carbon tax data. SARS provides information on all individual carbon tax filings since the implementation of the carbon tax in 2019. Overall, about 300 South African firms are subject to the carbon tax, reporting detailed information on their emissions inventories and tax payments. The reported emissions are further broken down by the components as shown in equation (3.1). This includes detailed information on fuel usage and various types of industrial process and fugitive emissions.[7] In addition, we observe the extent of allowances that are claimed by firms.

Financial and customs data. As a second source of information, we rely on the CIT-IRP5 firm panel, which harmonizes and combines corporate income tax (CIT) information, value-added-tax (VAT) as well as customs tax data on imports and exports (Ebrahim, Kreuser, and Kilumelume, 2021; Pieterse, Gavin, and Kreuser, 2018). The CIT-IRP5 firm panel encompasses the entire population of South African firms, totaling over 600,000 annual firm observations. Only 0.05% of these firms are subject to the carbon tax, leaving the vast majority outside its scope. These non-taxed firms provide a large pool for identifying a suitable control group. The CIT data is based on corporate tax returns submitted to SARS and comprises information on total sales and profits, capital stock, and the total wage bill. The information is based on the South African tax year which runs from March to February of the following year.[8] Customs information is obtained from transaction-level customs declaration forms containing information on the value of the transaction, product code as well as information on the partner country. From the customs derived data we mainly use firm-level aggregates of imports and exports. Finally, the data contain information of various transaction-level VAT-forms, which are aggregated to the firm-level.

[7] See Figure 3.A.2 in the Appendix for details.

[8] In contrast to the fiscal year, the reporting period for the carbon tax follows the calendar year and runs from January until December.

Employment Data. The final source of information we use is the individual panel, which contains individual-level data submitted by employers registered under the pay-as-you-earn (PAYE) scheme (Ebrahim and Axelson, 2019). Since the 2010/2011 tax year it is mandatory for employers to be part of PAYE. As the data is still incomplete for the early years up to 2013, we only use the individual panel starting in 2013. Most importantly, the data allow us to calculate the number of employees and the distribution of wages within each firm by linking it with the CIT-IRP5 firm panel.

Descriptive Statistics. Compared to the total population of firms in South Africa, only a small fraction is in principle liable for the carbon tax. Only about 300 firms have declared emissions associated with the carbon tax. This implies that the emissions from many firms are not covered. However, Panel (a) of Figure 3.3.1 shows that CO_2 are highly concentrated among firms. The gross emissions reported in carbon tax returns represent approximately 80% of total nationwide emissions (colored in red) (Crippa, Guizzardi, Pagani, Schiavina, Melchiorri, Pisoni, Graziosi, Muntean, Maes, Dijkstra, et al., 2024).[9] This suggests that carbon emissions in South Africa are concentrated among a relatively small number of firms, but also that the carbon tax would theoretically be able to cover the majority of emissions despite exempting the majority of firms.

Turning to net emissions (colored in green), which are calculated using equation (3.1) and represent the emissions subject to the carbon tax after applying various allowances, it becomes evident that this potential remains largely untapped. Over the years, only about 18% of nationwide emissions have been effectively taxed. The gap between gross emissions (total firm emissions) and net emissions (those subject to the carbon tax) is primarily due to the substantial allowances that firms can claim, which reduce their carbon tax base. This discrepancy is more pronounced when excluding firms that report emissions but effectively pay no tax (shown in orange). In most cases, this is because firms can offset their carbon tax payments by crediting the electricity levy, as detailed in Section 3.2. As a result, the proportion of taxed *and paid for* emissions out of total nationwide emissions drops to around 4%.

This discrepancy can be illustrated at the firm level by examining the *effective tax rates*, which reflect how much firms in different sectors are actually paying per ton of emitted

[9] Country-wide emissions are obtained from the Emissions Database for Global Atmospheric Research (EDGAR), which can be accessed under the following link: `https://edgar.jrc.ec.europa.eu/`. Last accessed: March 11th, 2025.

Figure 3.3.1: Emissions covered by carbon tax

a) Gross and net emissions **b)** Effective tax rate

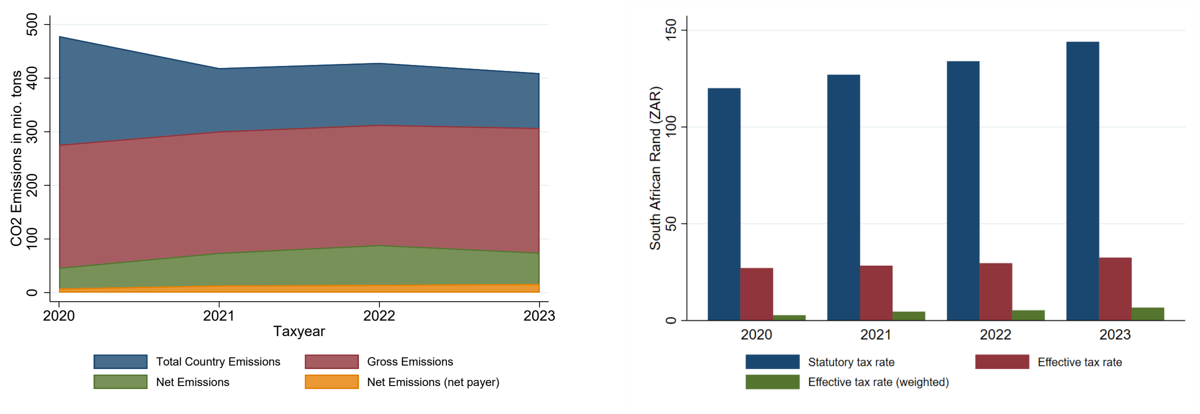

Notes: This figure plots in Panel (a) the nationwide emissions of South Africa (blue) from the Emissions Database for Global Atmospheric Research (EDGAR), the total emissions from the carbon tax firms (red), the net-of-allowances emissions (green), and net-of-allowances emissions of firms that are effectively paying the carbon tax (yellow). Panel (b) shows the effective tax rate filed by each company. Blue bars represent the statutory tax rate, and red and green bars are the effective tax rate paid after allowances. The tax year always refers to the reporting period in the previous calendar year. Hence, the tax year 2021 refers to the reporting period of January 2020 until December 2020. Data Source: SARS and EDGAR.

CO_2. The statutory tax rate, as defined by law, started at \$120 ZAR per ton of CO_2 and gradually increased to \$134 ZAR in 2023, as shown in Panel (b) of Figure 3.3.1. However, the average effective tax rate is significantly lower, as indicated by the red bar. When the average effective tax rate is weighted by firm-level emissions, it becomes clear that the effective tax rate per ton of CO_2 decreases even further. In other words, the green bars indicate that, on average, larger emitters are able to claim more allowances, leading to a reduced effective tax rate.

Naturally, the extent of emissions and, consequently, carbon tax payments vary across industries. Panel (a) of Figure 3.3.2 shows that the composition of net emissions differs by industry, which affects tax liability. Petrol and diesel are taxed through the fuel levy. To avoid double taxation, emissions from petrol and diesel combustion can be deducted from fuel combustion emissions and are thus exempt from the carbon tax. Fuel combustion accounts for about two-thirds of emissions in the mining sector, while it represents only half of the emissions in the manufacturing sector. The remaining emissions in both sectors are due to fugitive and industrial process emissions. Panel (b) displays the distribution of carbon tax firms by industry. Nearly half of all firms operate in the manufacturing sector, while another fifth are in the mining sector.

Next, we investigate heterogeneity in tax payments. Panel (c) of Figure 3.3.2 plots the distribution of carbon tax revenues across sectors. More than 80% of the tax revenues come from manufacturing firms and another 10% from mining firms. Panel (d) suggests, however, that the largest emitters do not necessarily pay the most. In fact, while we see manufacturing firms contributing over 80% of carbon tax payments but only about 25% of emissions, the electricity sector is responsible for about 64% of emissions but contributes only less than 2% of carbon tax payments. This is due to the deduction of the electricity levy. A similar pattern is observed for wholesale traders, where the share of emissions far exceeds their carbon tax payments. In the mining sector, although on a smaller scale, the share of carbon tax payments is three times larger than its share of emissions.

It is noteworthy that a substantial share of firms filing carbon tax returns report no tax liability, effectively categorizing them as non-payers. In our sample, this applies to 17% of the filings (see Appendix 3.A). There are several reasons why firms may report no carbon tax liability. First, some firms may have no emissions to declare for a given tax period. Second, electricity providers can offset the carbon tax by crediting the electricity levy. Third, firms might have zero net emissions, either because they only have emissions from petrol and fuel or because their emissions are fully sequestrated elsewhere (see Eq. (3.1)). In fact, most firms with zero emissions only report emissions from petrol and diesel, meaning they are effectively exempt from explicit carbon taxation.[10]

[10] The fuel levy, however, was aligned with the carbon tax rate so that their emissions are equivalent taxed via the fuel levy.

Figure 3.3.2: Industry heterogeneity

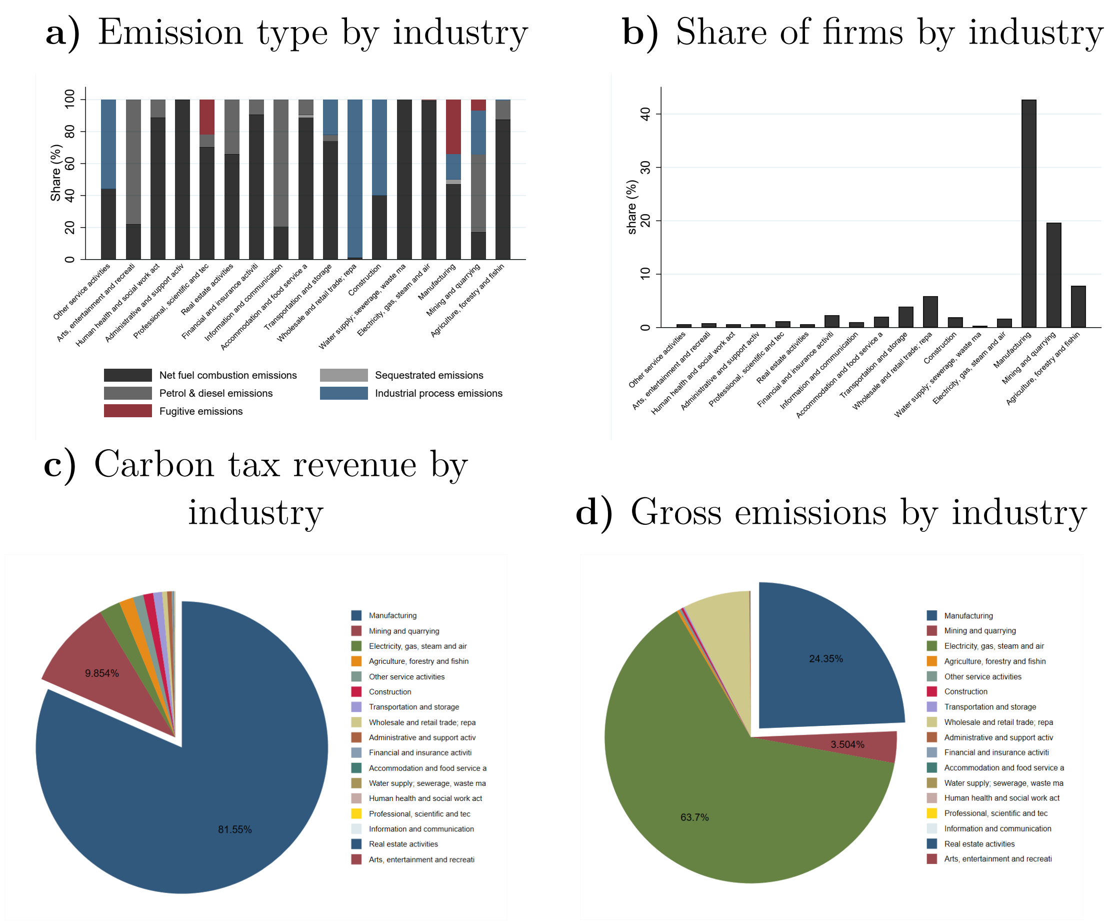

Notes: This figure plots various descriptive statistics based on the carbon tax returns filed by South African firms. Panel (a) depicts the share of emissions types by industry disaggregated according to Equation (3.1). Panel (b) plots the share of treated carbon tax firms by industry. Panel (c) displays the aggregate carbon tax revenue for the tax years 2020-2023 by industry. Panel (d) displays the aggregate emissions of treated carbon tax firms for the tax years 2020-2023 by industry. Data Source: SARS.

3.4 Empirical approach

In this section, we describe our empirical approach. The analysis relies on a matched difference-and-differences strategy, where matching is based on firms' observable characteristics before the announcement of the carbon tax.

Coarsened Exact Matching. In principle, firms that are liable for carbon tax might be different from those that are not. Due to various exemptions and size thresholds (cf. Section 3.2), simple comparisons of treated and non-treated firms are prone to selection bias. We overcome this issue by relying on matching techniques that reduce imbalances in the characteristics of treated and non-treated firms. Matching has become the leading approach to finding an appropriate control group in the field of environmental economics (Colmer, Martin, Muûls, and Wagner, 2024; Dechezlepretre, Nachtigall, and Venmans, 2023).

In our analysis, we use coarsened exact matching (CEM). This involves temporarily coarsening the data based on observed firm characteristics before the treatment period. The coarsening process applies a predefined common binning strategy, creating unique observations within the coarsened data. Each of these unique observations constitutes a stratum. Treated and untreated firms are then exactly matched on these strata. Observations whose strata do not contain at least one treated and one untreated observation are dropped, and weights are used to compensate for the different strata sizes (Iacus, King, and Porro, 2012). Importantly, and contrary to many other matching strategies, coarsened exact matching does not only account for imbalances in means but also for imbalances in higher moments and interactions (Blackwell, Iacus, King, and Porro, 2009; Iacus, King, and Porro, 2012). We match firms exactly at the 3-digit industry level. Additionally, firms are matched on the basis of profits, sales, import and export volume, number of employees, total wage bill, and capital stock. Each of these variables is divided into five equally sized bins. The baseline year is chosen to be the tax year of 2014-15, just before the first Carbon Tax Bill was drafted and made public. We excluded treated firms with sales and capital equal to zero in 2015, i.e., we kept only firms that were operating in 2015.

Estimation strategy. After creating a comparison group through matching, we estimate the causal effect of carbon taxation on firm behavior. In particular, we estimate an event-study type model as follows:

$$\ln(y_{it}) = \sum_{k=2011,k\neq 2015}^{2021} \theta_k \mathbf{1}[t=k] \times \mathbf{1}[i=CarbonTax] + \alpha_i + \gamma_{pt} + \epsilon_{it}, \tag{3.2}$$

where y_{it} is an outcome of interest for firm i in year t. $\mathbf{1}[i=CarbonTax]$ is an indicator for whether a firm ever files a carbon tax return.[11] α_i and γ_{pt} are firm and province-by-year fixed effects, respectively. Our coefficients of interest are the θ's. We exclude 2015 as

[11] Note, however, that due to the deduction of allowances, not all firms that file a carbon tax return have an actual carbon tax liability. However, we classify them as treated since they will begin paying the tax after the end of the transitional period. Therefore, our estimates reflect intent-to-treat effects. See Figure 3.A.3 in the Appendix for details.

the baseline year to accommodate anticipatory effects upon the release of the draft of the Carbon Tax Bill but before the actual implementation of the tax. Further, the year 2015 corresponds to the year for which we apply coarsened exact matching. In order to estimate the effect of the carbon taxes on firm behavior, it is required to establish a meaningful comparison group. To this end, we run eq. (3.2) with the weights obtained from the coarsened exact matching. The baseline estimation sample comprises 168 treated firms, i.e., about two thirds of all treated mining and manufacturing firms, and to 2,465 comparison (non-treated) firms from the same sectors.[12]

Identifying assumptions. Our empirical strategy for isolating the causal impact of the carbon tax on firm activity is based on two key identifying assumption. The first is the assumption of parallel trends of treatment and comparison groups. Put differently, we assume that in the absence of the carbon tax both treated and untreated firms would have followed the same economic trends over time. Although it is impossible to observe these counterfactual trends, we can scrutinize the plausibility of this assumption by inspecting whether the treatment and comparison groups have followed similar trends prior to the treatment. Given that a lengthy policy process preceded the implementation of the carbon tax, we explicitly allow for anticipatory behavior and define our first treatment as the announcement of the Carbon Tax Bill in 2015. Hence, we would need to assume parallel trends prior to the announcement in 2015.

Figure 3.4.1 plots the averages over time for our variables of interest such as sales, capital, employment, and profits by treatment status. While the dashed lines display the unmatched sample, the solid lines indicate the averages for the matched sample. The visual co-movement prior to 2015 supports the parallel trends assumption needed for a causal interpretation of the treatment effects. Although the parallel trends assumption relies on changes over time, the pre-treatment matching helps to alleviate further validity concerns by reducing cross-sectional differences between the treatment and comparison groups. The reasoning behind this is that carbon tax firms in South Africa are particularly large and could be exposed to systematically different types of contemporaneous shocks. As shown in Figure 3.4.1, the matching reduces the differences in pre-treatment characteristics substantially and thereby the likelihood of an estimation bias stemming from idiosyncratic differences in contemporaneous shocks.

The second identifying assumption is the Stable Unit Treatment Value Assumption (SUTVA).

[12] This ratio of matched to unmatched treated firms, as well as the absolute number of matched firms, is comparable to other studies e.g., Colmer, Martin, Muûls, and Wagner, 2024.

This assumption states that the observed trajectories of the treatment and comparison groups depend only on their respective treatment statuses. This assumption would be violated if the matched comparison firms would be indirectly treated by the carbon tax policy as well. For instance, if firms strategically adjust economic activity to stay below the liability thresholds. Depending on the direction of the effects, this can either lead to an upward or downward bias in our estimates. In order to alleviate this concern, we run our estimation with different matching stringencies. If we assume that the violation of SUTVA is increasing in the similarity of treatment and comparison firms, the estimation bias arising from SUTVA should become more prevalent the more similar treatment and comparison firms are. On the other hand, allowing for larger cross-sectional differences between treatment and comparison firms would attenuate potential biases arising from SUTVA violation. Appendix 3.B shows our estimation results when using different forms of matching. The reported results are comparable across different matching approaches, indicating that SUTVA violations might not be a major concern in our empirical setting. Moreover, it is important to note that the thresholds defining carbon tax liability are based on thermal input capacity, making them less susceptible to manipulation compared to metrics like sales. This limits the scope for firms to strategically bunch below the treatment threshold.

Figure 3.4.1: Coarsened Exact Matching Means

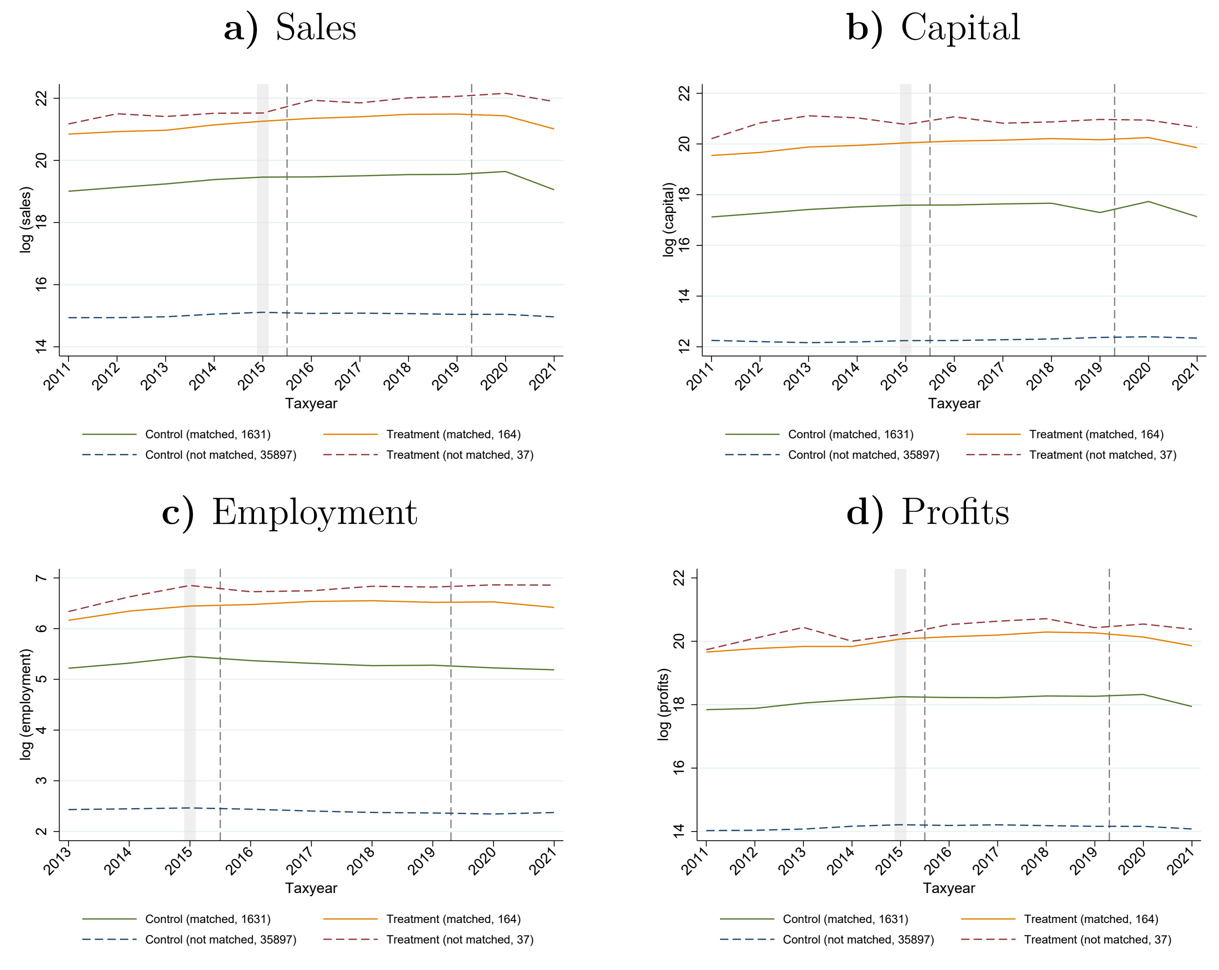

Notes: This figure plots the raw means of observable firm characteristics differentiated by treatment and matching status. The dashed red line depicts unmatched treated firms that are subject to the carbon tax. The yellow line depicts matched treated firms, and the green line matched comparison firms that are not subject to the carbon tax. The dashed blue line depicts all remaining untreated firms that were not matched. Firms were matched exactly on the 3-digit industry and coarsened with 3 cutpoints for sales, capital, number of employees, and profits based on the year 2015. Data Source: SARS.

3.5 Results

In this section, we report and discuss our estimation results. Based on the matched sample, we estimate equation (3.2). First, we present the results of the effects of the carbon tax on sales, capital, employment, and profits for all sectors. Then, we explore the heterogeneity by sector and allowance take-up. Finally, we provide evidence of potential mechanisms that explain our results.

3.5.1 Baseline results

Figure 3.5.1 presents the results from the dynamic event-study specification for four main outcomes across all firms: sales, capital, employment, and profits. The dashed vertical lines indicate two key events. First, at the end of 2015, the first Carbon Tax Draft Bill was published. To account for anticipation effects that may have begun at that time, we use 2015 as our baseline year (which is omitted from the graph). Second, the actual implementation of the Carbon Tax Act occurred in 2019. It is important to note that the fiscal year in South Africa runs from March to February. Therefore, the fiscal year 2020 started in March 2019 and ended in February 2020. As a result, the estimated effects for 2020 largely reflect the impact of the carbon tax (which was implemented in June 2019), without being influenced by the COVID-19 pandemic, which began in the following fiscal year (Burger and Calitz, 2021).

In all panels, we observe that prior to 2015, there were no significant differences in the evolution of outcomes between taxed and non-taxed firms. It is only after 2016, following the carbon tax announcement, that the outcomes begin to show positive effects on sales, capital, employment, and profits. This suggests that the announcement did not hinder firm growth; instead, firms may have had an anticipatory incentive to expand operations *before* the carbon tax was implemented. After the actual implementation in 2019, these positive effects persist. Appendix 3.B shows that these results remain robust to different matching strategies.

3.5.2 Environmental effects

A limitation of the administrative records is that we cannot observe firms' emissions before the carbon tax policy was implemented, making it impossible to directly measure the policy's impact on emissions at the firm level. To address this, we turn to more aggregate measures and estimate the environmental effects at the country level using a synthetic control method. This approach involves constructing a "synthetic" version of South Africa by combining data from countries that did not implement a carbon tax or similar carbon pricing schemes. The synthetic control is essentially a weighted combination of these countries, designed to resemble South Africa in terms of key characteristics such as emissions patterns before the policy was introduced (Abadie, Diamond, and Hainmueller, 2010; Andersson, 2019). This method is widely used in the empirical literature to overcome the challenge of not having a well-defined comparison group (Abadie, 2021).

Figure 3.5.1: The effects of the carbon tax policy on firm outcomes

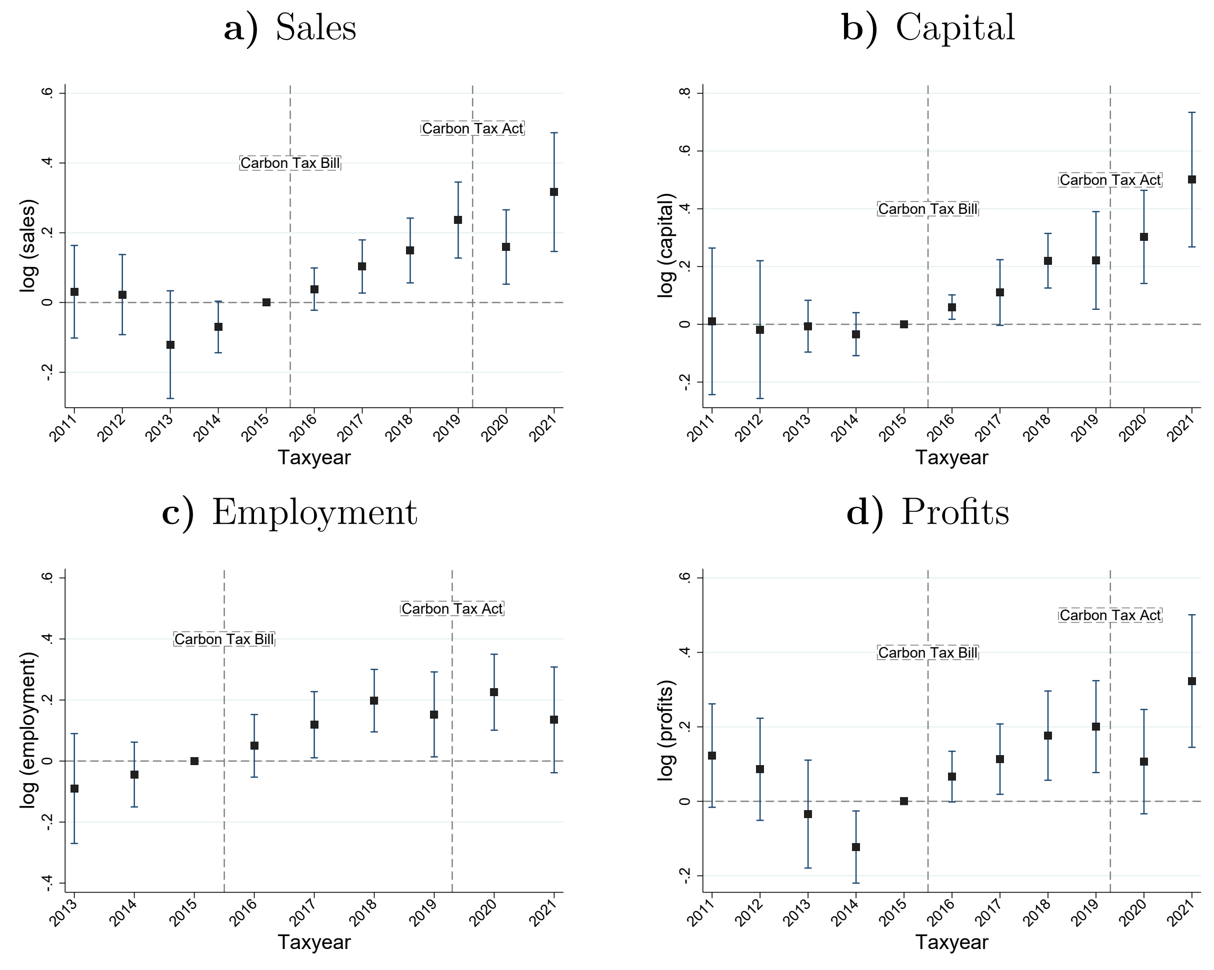

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The vertical bars around the estimates, represented by the squares, show the 95% confidence intervals. The “tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

For this analysis, we use country-level emissions data from the Emissions Database for Global Atmospheric Research (EDGAR) (Crippa, Guizzardi, Pagani, Schiavina, Melchiorri, Pisoni, Graziosi, Muntean, Maes, Dijkstra, et al., 2024), combined with data from the Penn World Table for population and GDP figures (Feenstra, Inklaar, and Timmer, 2015). The counterfactual in this case is a “synthetic South Africa,” which is constructed using average of annual CO_2 emissions data averaged from 2011 until 2015. To more closely match the composition of our firm-level sample, we exclude emissions from the energy and transport sectors. Therefore, the focus of the environmental analysis is on country-level emissions primarily arising from fuel combustion in the manufacturing sector, as well as fugitive and industrial process emissions.

Additionally, we match countries based on the average GDP as well as their average emission intensity (measured as the ratio of emissions and GDP) between 2011 and 2015. The resulting "synthetic South Africa" is a weighted average consisting primarily of the following countries: United Arab Emirates (66.7%), India (3%), Trinidad and Tobago (2.9%), Vietnam (2%) and Saudi Arabia (0.5%).[13] This synthetic counterfactual is used to estimate what emissions in South Africa would have looked like in the absence of the carbon tax policy.

Figure 3.5.2 compares the emissions of South Africa to those of the "synthetic South Africa". The latter is supposed to represent the counterfactual scenario in which no carbon tax would have been introduced. Panel (a) plots the evolution of both while Panel (b) illustrates the absolute differences over time. The results show a temporary increase in CO_2 emissions between the announcement of the carbon tax and its actual implementation with a drop in emissions after the actual policy implementation. Although only suggestive in nature, the emission pattern aligns well with the positive economic impacts on firm outcomes. They suggest that the positive effect of the tax announcement on firm activity has been accompanied by a temporary increase in emissions. As the firms we observe under the carbon tax account for about 80% of nationwide emissions (cf. Figure 3.3.1), it is reasonable to assume that their behavior affects overall country emissions.

3.5.3 Heterogeneity by sector and allowance take-up

Figure 3.5.3 presents the results for the two types of firms in the sample. First, the patterns differ by sector depending on the outcome. The carbon policy announcement has positive effects on both sales and profits for firms in both sectors, with stronger effects observed for mining firms. However, these effects are more imprecise for mining firms due to the smaller sample size. In contrast, the effects on capital and employment vary by sector: there is little to no effect on capital for the mining sector. Regarding employment, the impact is immediate for the mining sector following the Carbon Tax Draft Bill, while it takes two years for the effect to materialize in the manufacturing sector.

These findings align with a South African National Treasury's 2019 memorandum, which

[13] Overall "synthetic South Africa" comprises 130 countries, of which 125 countries have a weighted contribution of less than 0.5%. Appendix Table 3.B.1 features the full list of countries comprising "synthetic South Africa".

Figure 3.5.2: The effect of the carbon tax policy on CO_2 emissions

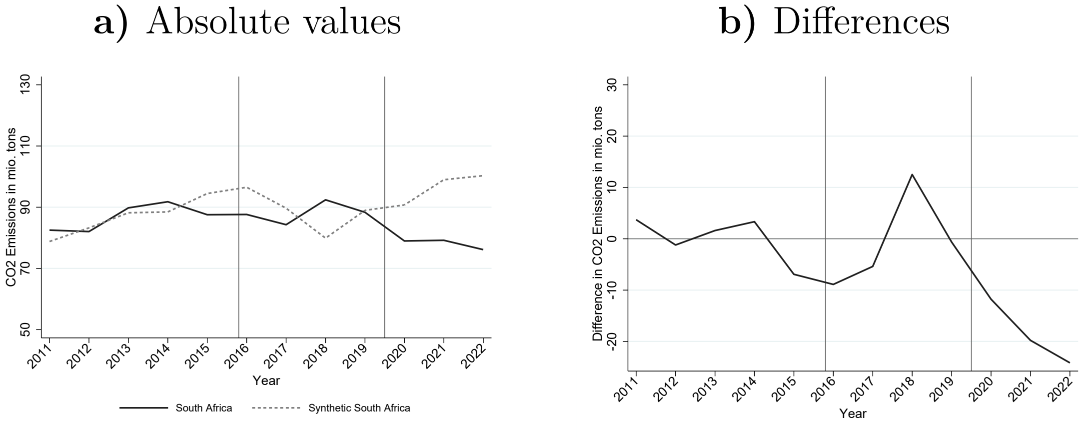

Notes: This figure plots the effect carbon tax policy CO_2 emissions using the synthetic control method. In Panel a) absolute values of CO_2 emissions are shown. Panel b) depicts the difference in CO_2 emissions between South Africa and its synthetic counterpart. The plotted emission values cover emissions from the following IPCC codes: 1.A.2, 1.B., 2.A., 2.B., 2.C. and 2.D.. Synthetic South Africa is constructed based on South Africa's average annual emissions in these sectors as well as the average GDP and emissions intensity (emissions/GDP) between 2011-2015. Countries that have implemented a carbon pricing policy are removed from the pool of potential donor countries. Data Source: EDGAR and Penn World Tables.

highlighted that the carbon tax implementation would be complemented by transitional tax incentives designed to minimize the policy's initial impact. Specifically, the government aimed to cushion potential adverse effects on energy-intensive sectors, such as mining. This could explain why mining firms, despite being heavily impacted by the policy, experienced quicker adjustments in sales and employment following the policy announcement.

As explained earlier, allowances are a key feature of the South African carbon tax policy, as they determine the effective tax rate and, consequently, the treatment intensity. Figure 3.5.4 presents the results separately for firms that benefited the least from allowances (1st quintile) and the most (5th quintile). The results suggest that firms in the 1st quintile, which are limited in claiming allowances, respond quickly to the policy announcement. In contrast, the observable effects for firms in the 5th quintile, which benefit the most from allowances, are much smaller and often statistically indistinguishable from zero.

These results might sound counterintuitive because one would typically expect that firms facing higher effective tax rates (due to fewer allowances) would experience negative economic outcomes, given the increased cost burden. However, the quicker response of firms in the 1st quintile could be due to their anticipation of higher tax liabilities, prompting them to adjust their behavior more rapidly in response to the policy announcement. On the other hand, firms in the 5th quintile, which benefit more from allowances, might feel less pressure to alter their operations, leading to a more muted response. This suggests that the financial cushioning provided by allowances could reduce the urgency for firms to make immediate adjustments, despite facing lower effective tax rates. We will explore these and other potential mechanisms further in the next section.

Figure 3.5.3: Heterogeneous effects of the carbon tax policy by sector (manufacturing vs. mining)

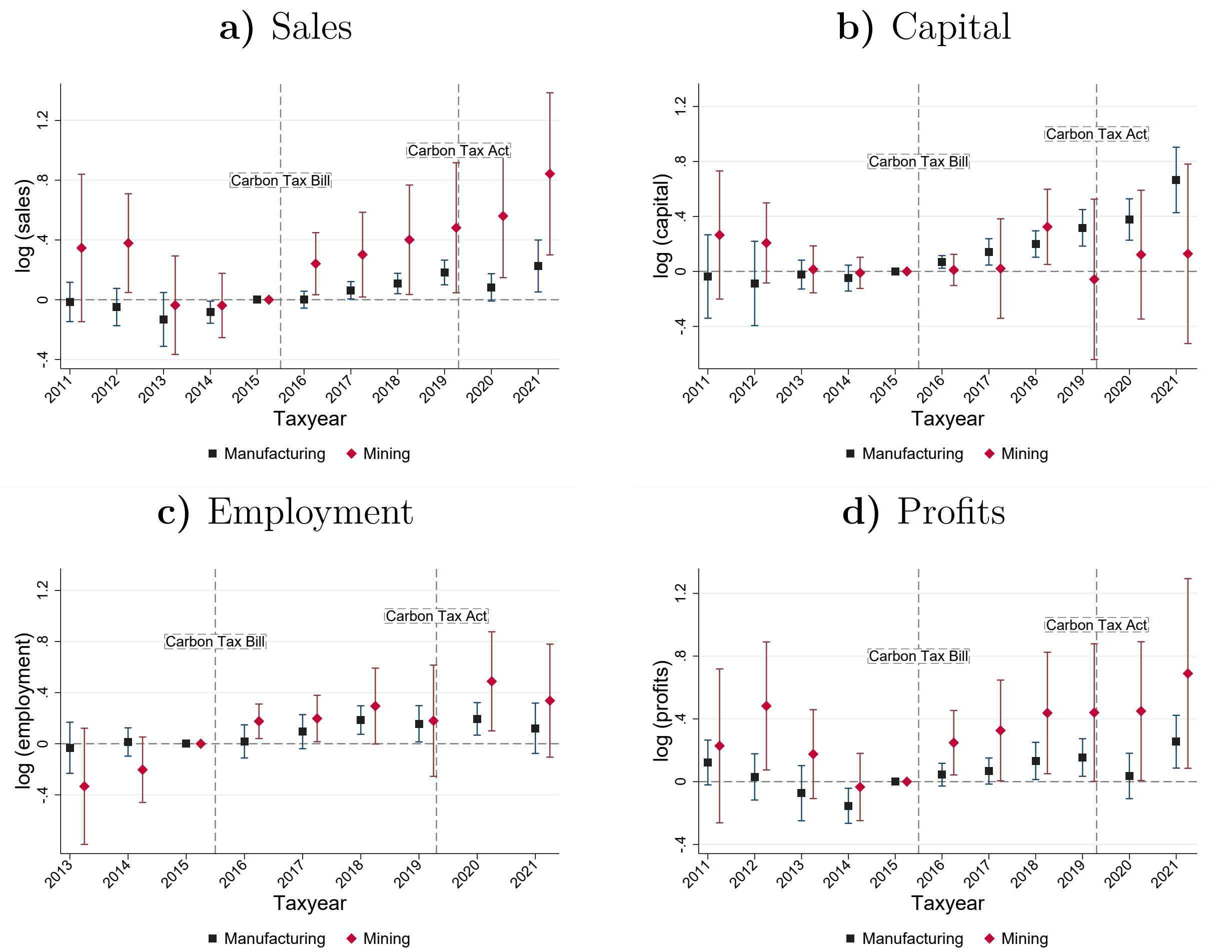

Notes: This figure plots θ_k coefficients estimated from equation (3.2). The blue squares depict the point estimates for the manufacturing sector. The red diamonds depict the point estimates for the mining sector. The blue and red vertical bars around the point estimates represent the 95% confidence intervals. The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

Figure 3.5.4: Heterogeneous effects of the carbon tax policy by allowance take-up

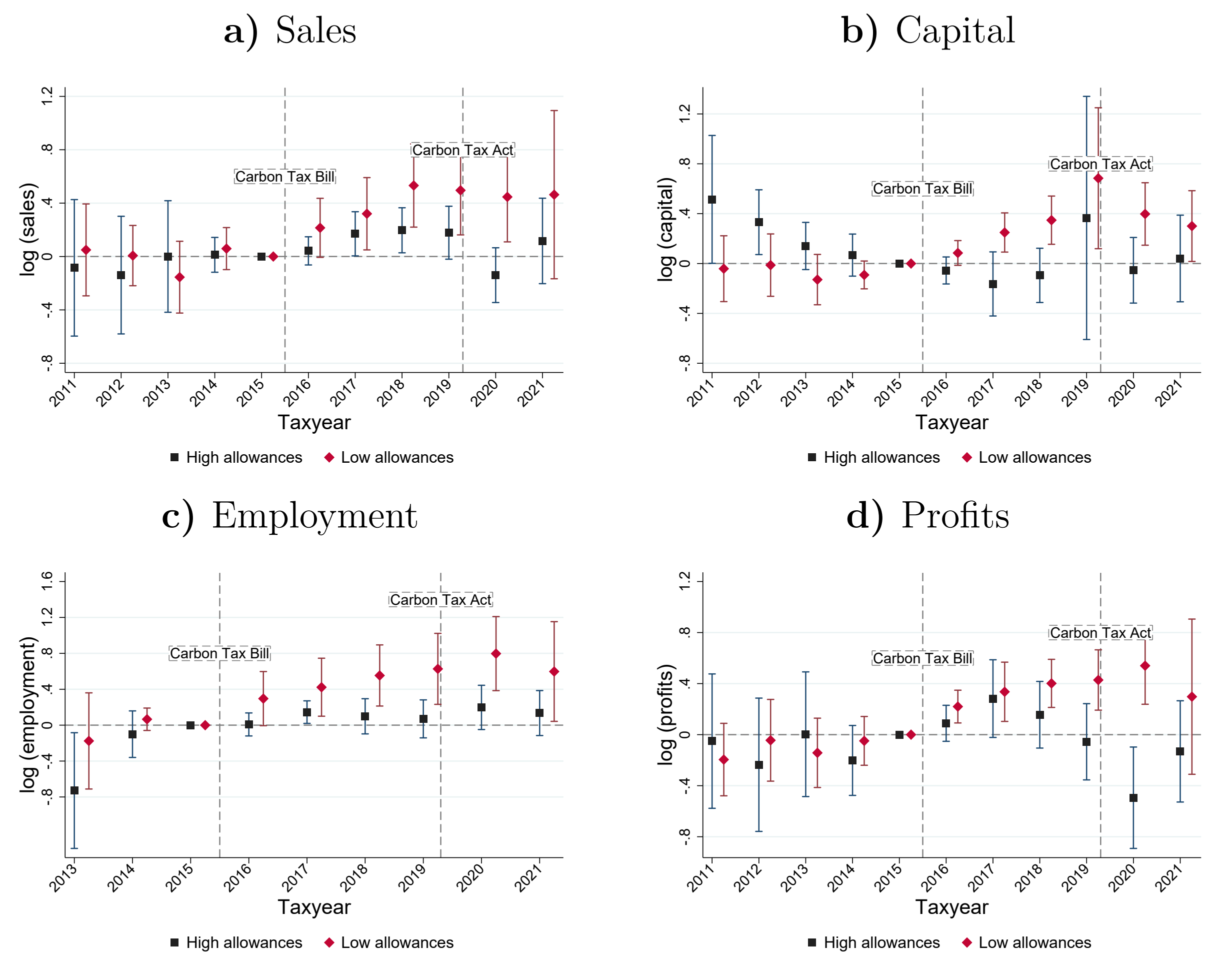

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The blue squares represent the point estimates for firms with high allowance take-up (5th quintile), while the red diamonds represent the point estimates for firms with low allowance take-up (1st quintile). The blue and red vertical bars around the point estimates represent the 95% confidence intervals. The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

3.6 Mechanisms

In previous sections, we showed that the announcement and implementation of the carbon tax led to positive effects on sales, capital, employment, and profits. In this section, we explore two potential channels that could explain this pattern. Specifically, we focus on how uncertainty resolution, firms' anticipation of future costs, and technology upgrading may have driven firms to increase their activity following the tax announcement.

3.6.1 Uncertainty and anticipation

Firms must make decisions under uncertainty across many dimensions, including the regulatory framework in which they operate, particularly regarding taxation. Environmental accountability through taxation had been part of the South African policy agenda well before its formal implementation in 2019, with discussions about its structure and purpose taking place as early as 2010 (cf. Figure 3.2.1). However, the release of the Carbon Tax Bill in 2015 represents a key milestone in this process, as it outlined the core elements of what would later become the Carbon Tax Act.

The Carbon Tax Bill provided critical details, such as the emission factors – how much CO_2 a specific production activity would generate – as well as the applicable tax rates and the allowance structure. These disclosures allowed firms to calculate their likely tax liability with a greater degree of certainty. Furthermore, the tax design itself, which does not impose an emissions cap as a cap-and-trade system would, likely eased concerns about production limits. Under the carbon tax framework, firms retained the flexibility to produce as much as they desired, as long as they found it economical to pay the associated tax. This aspect reduced a significant layer of uncertainty by ensuring that firms' production decisions would not be constrained by a rigid emissions cap. Overall, the 2015 Carbon Tax Bill clarified several previously uncertain aspects of the policy, enabling firms to better anticipate its financial implications and adapt their strategies accordingly.

This resolution of uncertainty might thus have improved the efficiency of firms that were uncertain of how costly their production in the near future would be. It is reasonable to assume that the Carbon Tax Bill in 2015 constitutes a positive information update to firms' expectations as it already incorporated many aspects to attenuate concerns that firms had at the outset (see Appendix 3.C for details). If a firm's activity is correlated with its expectation of the future, this channel can rationalize the observed positive effects.

Uncertainty about carbon pricing has been shown to be high and relevant for firm decisions in more developed settings (Fuchs, Stroebel, and Terstegge, 2024). Arguably, uncertainty might be even higher in less developed economies and thus relevant for our setting. To investigate this channel, we focus on the response of the mining sector, in which the resolution of uncertainty was especially prevalent. In fact, they were explicitly mentioned in a press release accompanying the release of the Carbon Tax Bill in 2015, stating that ' *taking into account the current state of the mining and other distressed sectors [the carbon tax] will be designed to ensure that such sectors are not adversely affected when the tax is implemented.*" This suggests that these sectors were uncertain and concerned about the future cost of production before 2015, and thus updated their expectations positively afterwards. In line with these arguments, subsection 3.5.3 reports that mining firms responded more strongly than manufacturing firms, which supports this hypothesis.

The resolution of uncertainty clarified the future costs for firms subject to the carbon tax. One clear example of how these updated expectations influence firm behavior is through strategic anticipation of future production costs. In the sense of Sinn, 2012, firms might shift polluting activities from the expensive future (under carbon tax) to the cheaper present (without carbon tax). As firms decide on how to use their capital and how to extend production across time, it could be that firms increase production upon policy announcement and stop increasing or even decrease it after the carbon tax is implemented. While this idea primarily stems from the mining sector, manufacturing firms might have a similar incentive by using up their existing " dirty" capital for the production of goods. Our analysis has revealed that firms were affected heterogeneously by the carbon tax. In particular, the allowances varied across sectors but also across firms within sectors. These allowances and, therefore, approximately the effective tax rates became known to firms upon the announcement in 2015. Considering the anticipation of future costs as an explanation, firms with higher expected tax rates should react more strongly, which is corroborated in Figure 3.5.4.

We present two additional observations in Figure 3.6.1 that support this channel. First, Panel (a) shows the dynamic effects of the announcement on the inventory levels of treated firms, distinguishing by their allowance take-up. The results indicate a modest increase in inventory following the 2015 announcement for firms with low allowance take-up (and therefore facing relatively higher effective tax rates). This upward trend continues until 2019, with an observed inventory increase of approximately 30% in 2018, significant at the 5% level, but statistically indistinguishable from firms with high allowance take-up. This pattern suggests that some firms accumulated inventory in anticipation of the tax implementation. After the tax came into effect in 2019, the estimated coefficients decline, consistent with a reduction in production.

Second, Panel (b) illustrates a sharp rise in capital depreciation following the announcement, with effects concentrated among firms with low allowance take-up. Together with the positive effects on sales reported in Panel (a) of Figure 3.5.4, these findings collectively imply that firms expecting higher tax rates may have intensified the use of emission-intensive machinery before the tax implementation. The accelerated depreciation likely reflects an effort to mitigate the risk of stranded assets – machinery that could become uneconomical or obsolete under the new tax regime. We cannot rule out, however, that firms may have also employed bookkeeping strategies or asset management practices to speed up depreciation, potentially to reduce taxable income or to adjust the valuation of assets on their balance sheets in response to the anticipated tax burden. This could have been a way to optimize their financial position under the carbon tax.

Figure 3.6.1: Heterogeneous effects of the carbon tax policy by allowance take-up: inventory and depreciation

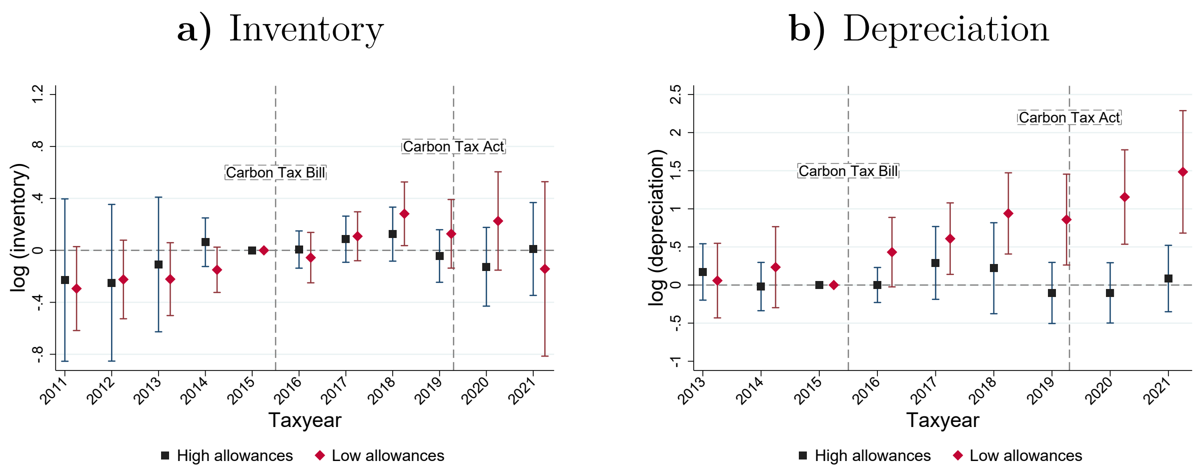

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The blue squares represent the point estimates for firms with high allowance take-up (5th quintile), while the red diamonds represent the point estimates for firms with low allowance take-up (1st quintile). The blue and red vertical bars around the point estimates represent the 95% confidence intervals. The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

3.6.2 Technology upgrading

Another explanation for the increased firm activity upon the tax announcement is technology upgrading. Firms could change production technologies to lower emissions and thus avoid high tax payments. Ultimately, this could increase the competitiveness and profitability of firms and, therefore, explain the positive effects we observe (Acemoglu, Aghion, Bursztyn, and Hemous, 2012; Porter and Linde, 1995). In other contexts, it has been shown that the introduction of an Emissions Trading System (ETS), in fact, triggered innovation and investments in cleaner technology (Calel, 2020; Calel and Dechezlepretre, 2016). For the case of China, for instance, Cui, Zhang, and Zheng, 2023 show that the ETS induced innovation by firms leading to a comparative advantage over non-ETS firms and thus higher productivity. In South Africa, we find only limited evidence pointing toward innovation or investments in cleaner technology as the main drivers of the positive effects of the carbon tax on firm outcomes.

In Figure 3.6.2, we re-estimate our baseline specification using three indicators of innovation or technology upgrading as dependent variables. It is important to note that we did not have access to patenting data for these firms, which is why we use these alternative proxies instead. Panels (a) and (b) show that there is generally no significant effect on R&D expenditures – if anything, the point estimates are often negative. This suggests that firms subject to the carbon tax did not increase innovation. Panels (c) and (d) display the dynamic effects on total imports, which can serve as an indicator of whether firms are purchasing clean technology or new machinery from abroad. We observe positive effects on import volume following the carbon tax announcement, but these are only statistically significant on the extensive margin. When focusing on *new imports* in Panels (e) and (f), we find a positive effect on the extensive margin after the tax announcement, but no effect on the intensive margin. While we cannot rule out some impact on innovation and technology upgrading, these analyses collectively suggest that these factors are unlikely to play a major role in explaining the positive baseline results.

Figure 3.6.2: Effects of the carbon tax policy on R&D and imports (total vs. new products)

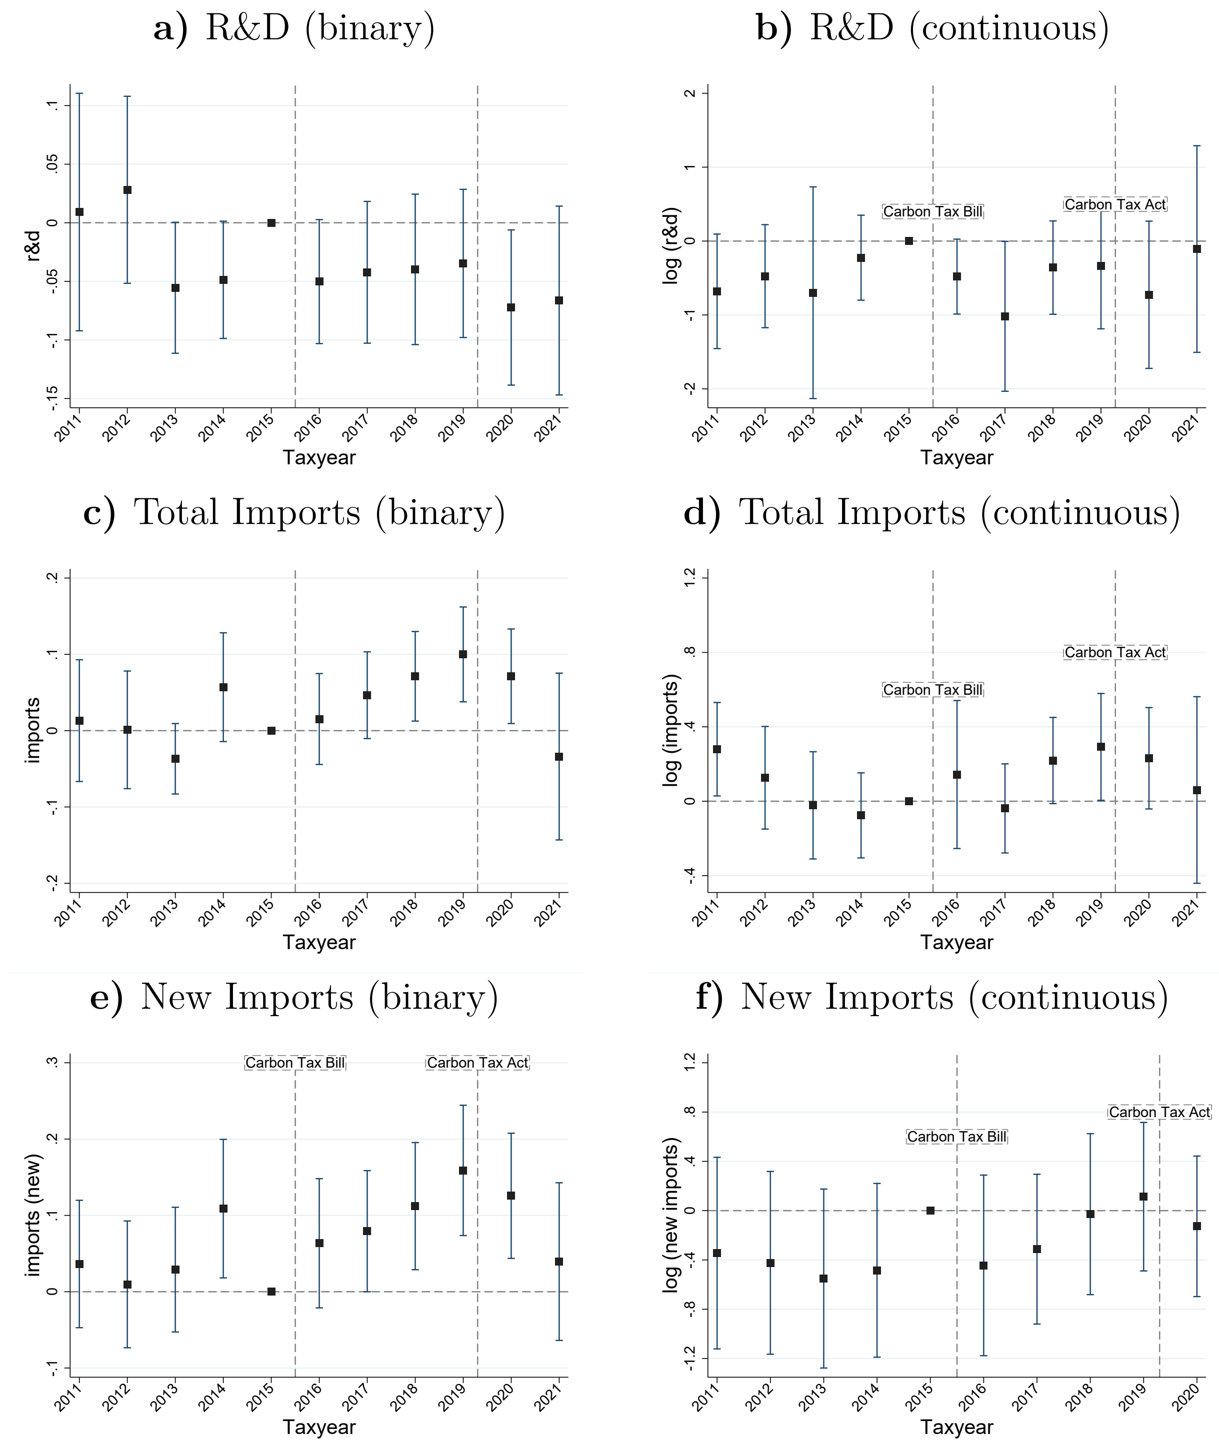

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The vertical bars around the estimates, represented by the squares, show the 95% confidence intervals. "Binary" refers to an outcome that is a dummy variable (extensive margin), while "continuous" refers to an outcome that is measured on a continuous scale (intensive margin). The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

3.7 Concluding Remarks

This chapter provides the first comprehensive analysis of how carbon taxation affects firm performance in a large emerging economy. Using detailed administrative data from South Africa, we quantify the impact of the introduction of a nationwide carbon tax on various firm outcomes. Contrary to many expectations, our findings suggest that the carbon tax does not have negative effects on sales, capital, employment, and profits. In fact, treated firms seem to experience greater growth in these outcomes compared to their matched counterparts. Upon the announcement of the tax – four years prior to its implementation – treated firms begin to increase activity, and this positive trend continues even after the tax is implemented.

Additional analyses suggest that these positive effects may result from firms anticipating the tax and adjusting their economic activities in the short term to shift polluting behavior. In particular, firms may have intensified their use of emission-intensive machinery or accelerated capital depreciation to avoid the risk of stranded assets – machinery that could become uneconomical or obsolete under the new tax regime. While we cannot fully rule out alternative explanations, we do not find robust evidence that firms are investing in cleaner production methods through R&D or importing new, cleaner equipment.

From a policy perspective, these findings suggest that the introduction of a carbon tax may not necessarily harm firm performance in the short term, even in emerging economies. The ability of firms to adapt and adjust may help mitigate the anticipated costs. However, this raises important questions about the carbon tax's effectiveness in driving long-term environmental improvements, as we find no robust evidence of firms investing in cleaner technologies or practices in response to the tax.

Looking ahead, the end of the transition period on December 31, 2025, when allowances are supposed to be phased out and the carbon tax rate increases to its statutory value, could provide firms with a stronger financial incentive to transition to cleaner production methods. As the tax burden becomes more substantial, firms may be more inclined to invest in greener technologies or adopt more sustainable practices to mitigate higher tax liabilities. However, this transition could be gradual and challenging for some firms, especially in sectors heavily reliant on emission-intensive processes. Given these dynamics, policymakers may consider complementary measures to accelerate this transition. This could include subsidies for green technology, direct incentives for R&D in sustainable practices, or additional regulatory measures that encourage firms to adopt cleaner production methods. Without such support, the carbon tax at the current statutory rate alone may not be sufficient to drive the necessary long-term environmental shifts.

Appendix

3.A Additional descriptive results

This appendix provides additional descriptive results on the anatomy of the carbon tax, complementing Section 3.3. Figure 3.A.1 shows the annual carbon tax revenue as depicted by the official budget as well as by aggregating the raw carbon tax micro data. Figure 3.A.2 depicts the different emissions sources for fuel combustion, industrial process emissions as well as for fugitive emissions. Figure 3.A.3 breaks down carbon tax firms that eventually do not pay any carbon tax.

3.A. *Additional descriptive results*

Figure 3.A.1: Aggregate carbon tax revenue

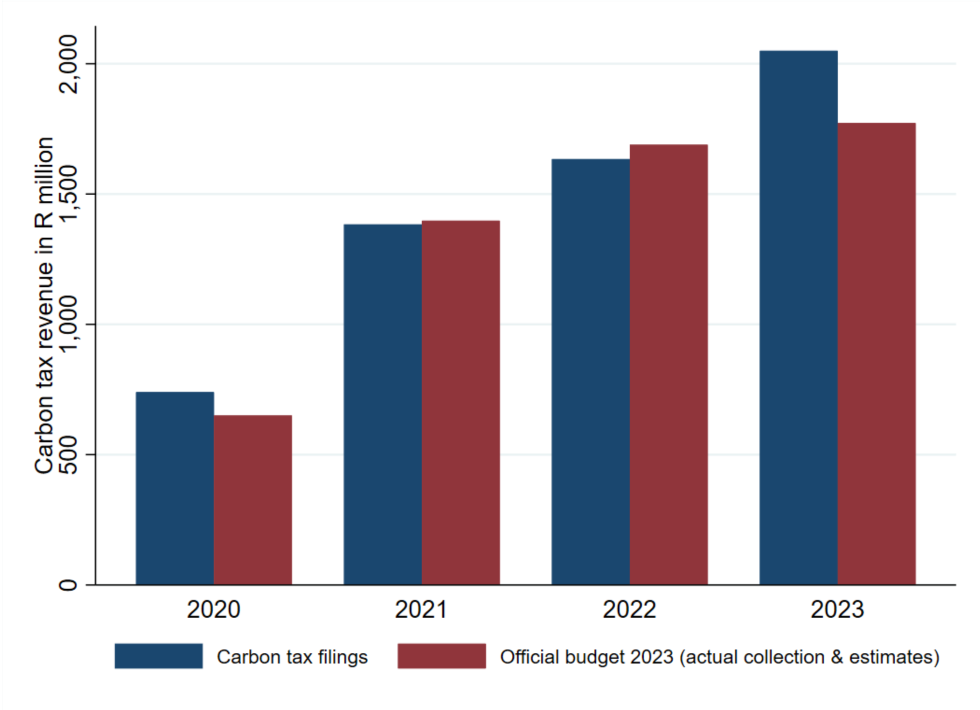

Notes: This figure depicts the annually aggregated carbon tax payments. Blue bars show the aggregate numbers for the carbon tax filings provided by SARS. Red bars indicate the aggregate figures by the Treasury in their budget review for 2023. The red bars for 2022 and 2023 are estimations by the Treasury, while 2020 and 2021 refer to actual revenue. Data Source: SARS.

Figure 3.A.2: Emissions sources

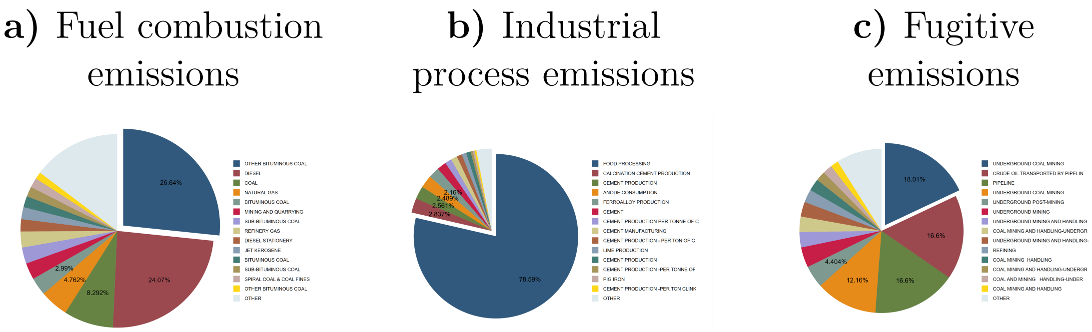

Notes: This figure plots various descriptive statistics based on the carbon tax returns filed by South African firms. Panel (a) depicts the share of fuel combustion emissions disaggregated by fuel type. Panel (b) plots industrial process emissions disaggregated by activity. Panel (c) plots fugitive emissions disaggregated by activity. All data is based on reported emissions of carbon tax firms for the taxyears 2020-2023. Data Source: SARS.

Figure 3.A.3: Anatomy of non-payers

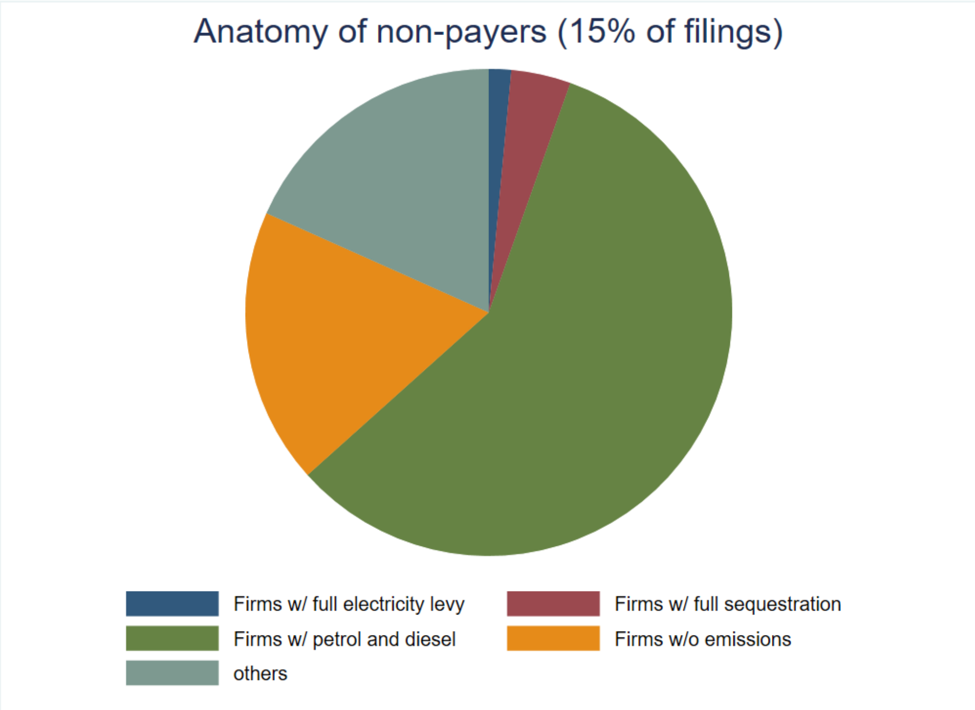

Notes: This figure plots the anatomy of carbon tax firms that eventually do not pay any carbon tax during the period 2020-2023. The green area depicts firms that only use petrol and diesel. The red area refers to firms that can fully sequestrate their emissions. The blue area depicts electricity firms that can fully deduct the electricity levy. The orange area depicts firms that report zero gross emissions. The remaining firms are all other firms that can not be exclusively classified in one of those categories. Data Source: SARS.

3.B Robustness

This appendix complements Section 3.5 by providing robustness checks to the main results. Figures 3.B.1 and 3.B.2 plot the raw means, when matching less or more restrictively on observable than in the main specification (see figure notes for details). Figures 3.B.3 and 3.B.4 show the corresponding event study results. In figure 3.B.5, we demonstrate that the results are insensitive to the choice of the baseline year.

Figure 3.B.1: Coarsened Exact Matching Means (less stringent / wide matching)

a) Sales

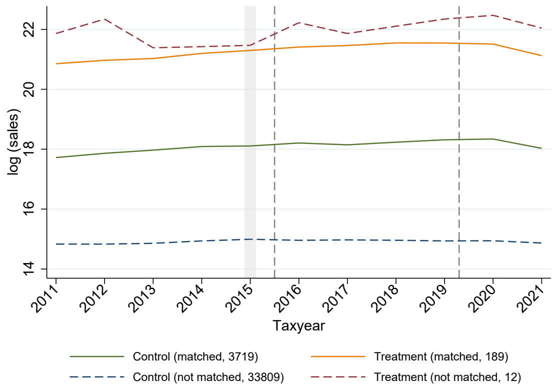

b) Capital

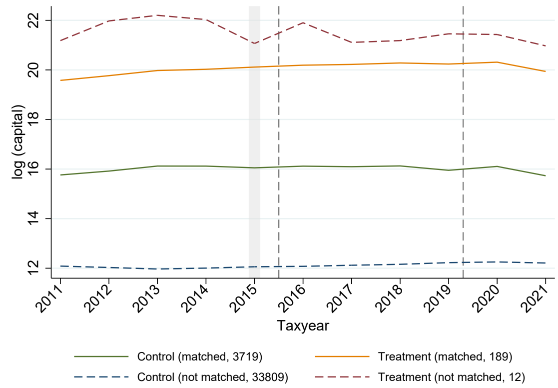

c) Employment

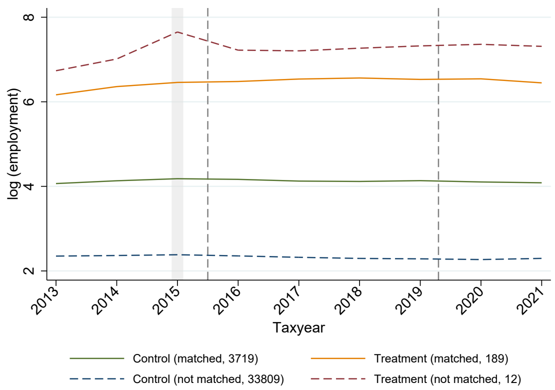

d) Profits

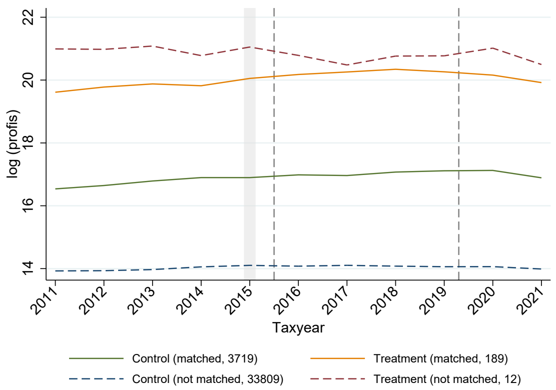

Notes: This figure plots the raw means of observable firm characteristics differentiated by treatment and matching status. The dashed red line depicts unmatched treated firms that are subject to the carbon tax. The yellow line depicts matched treated firms, and the green line matched comparison firms that are not subject to the carbon tax. The dashed blue line depicts all remaining untreated firms that were not matched. *Firms were matched exactly on the 3-digit industry and coarsened with only 1 cutpoint* for sales, capital, number of employees, and profits based on the year 2015. Data Source: SARS.

Figure 3.B.2: Coarsened Exact Matching Means (more stringent / narrow matching)

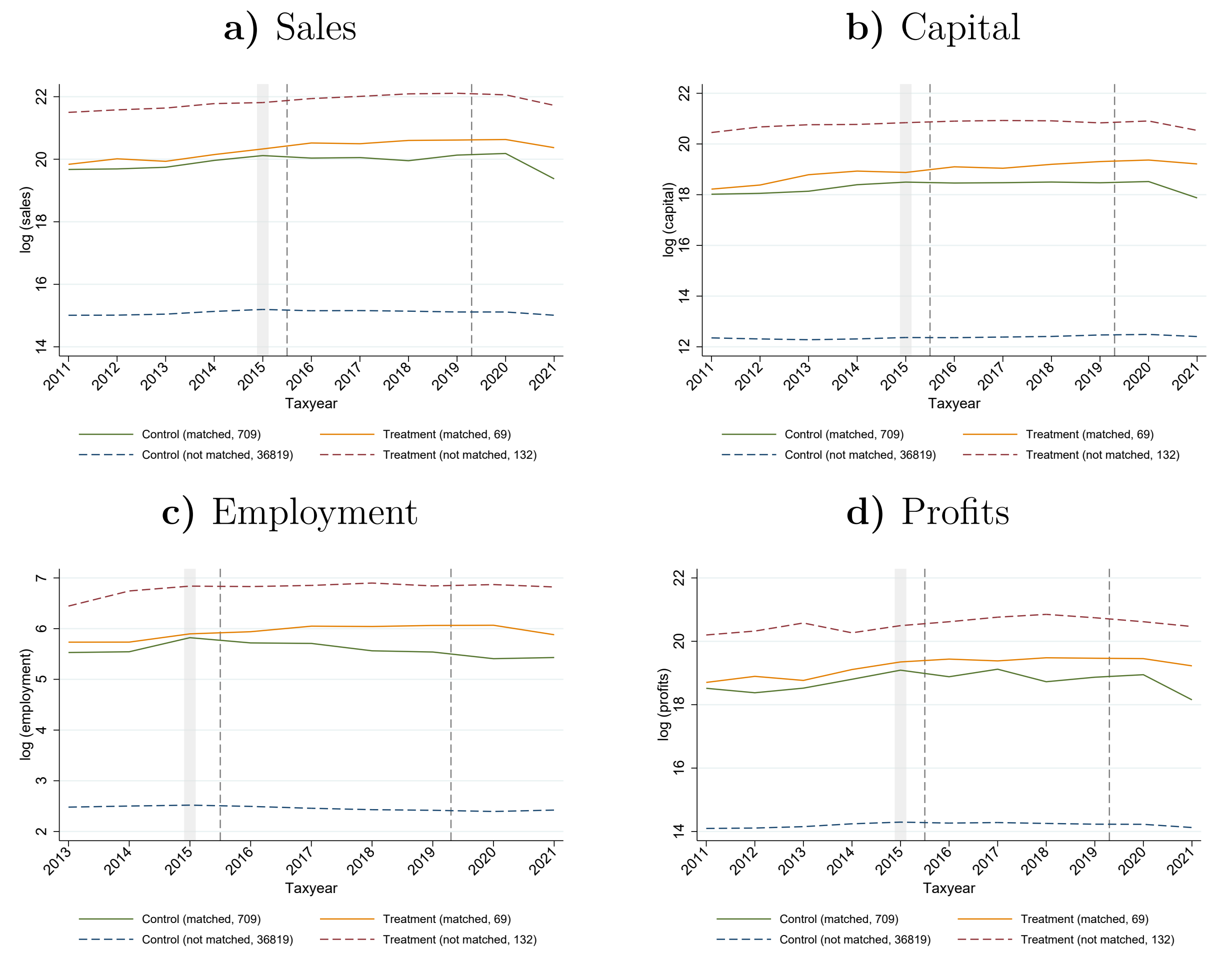

Notes: This figure plots the raw means of observable firm characteristics differentiated by treatment and matching status. The dashed red line depicts unmatched treated firms that are subject to the carbon tax. The yellow line depicts matched treated firms, and the green line matched comparison firms that are not subject to the carbon tax. The dashed blue line depicts all remaining control firms that were not matched. *Firms were matched exactly on the 1-digit industry and coarsened with 10 cutpoints* for sales, capital, number of employees, and profits based on the year 2015. Data Source: SARS.

Figure 3.B.3: The effects of the carbon tax policy on firm outcomes (less stringent / wide matching)

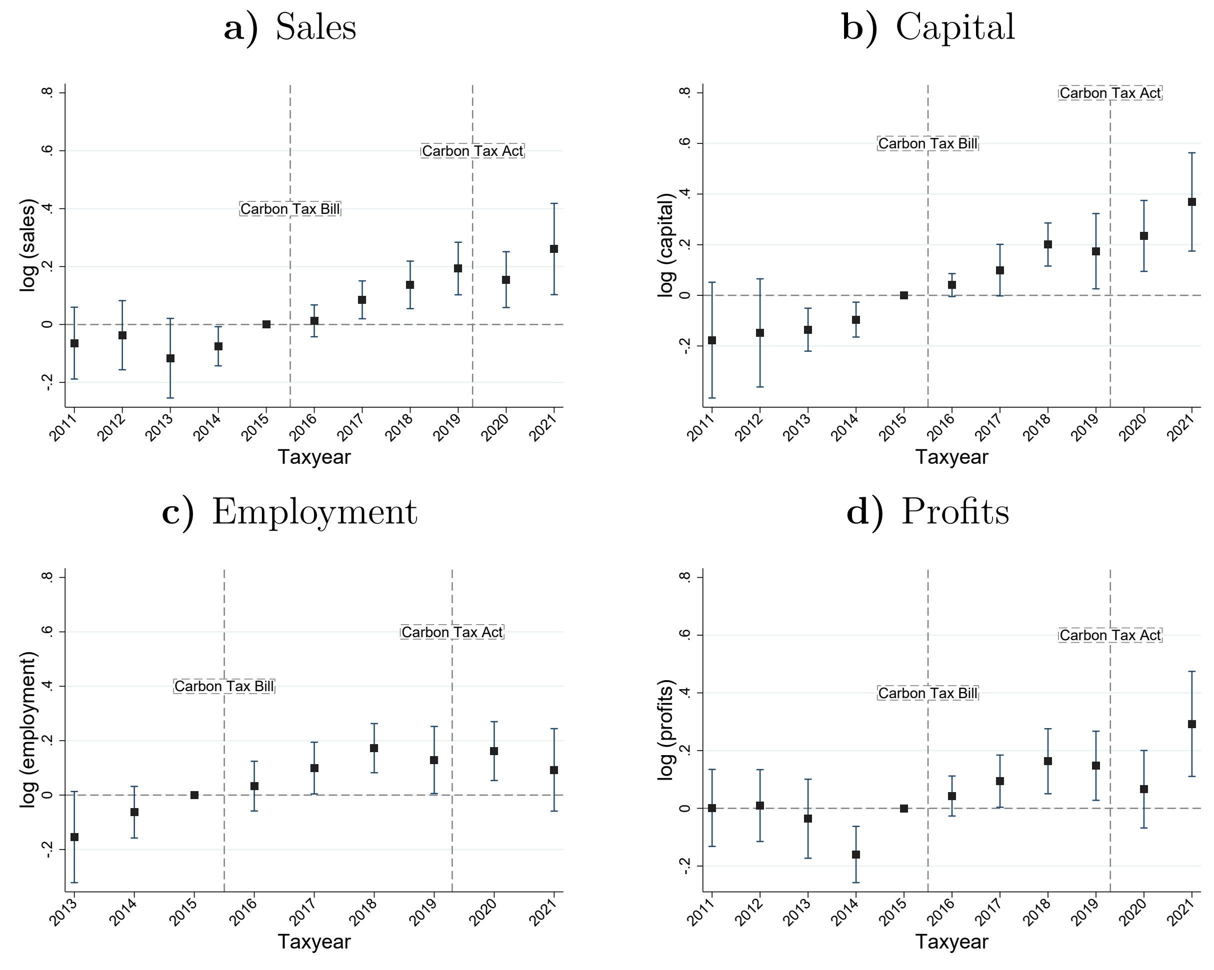

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The vertical bars around the estimates, represented by the squares, show the 95% confidence intervals. The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

Figure 3.B.4: The effects of the carbon tax policy on firm outcomes (more stringent / narrow matching)

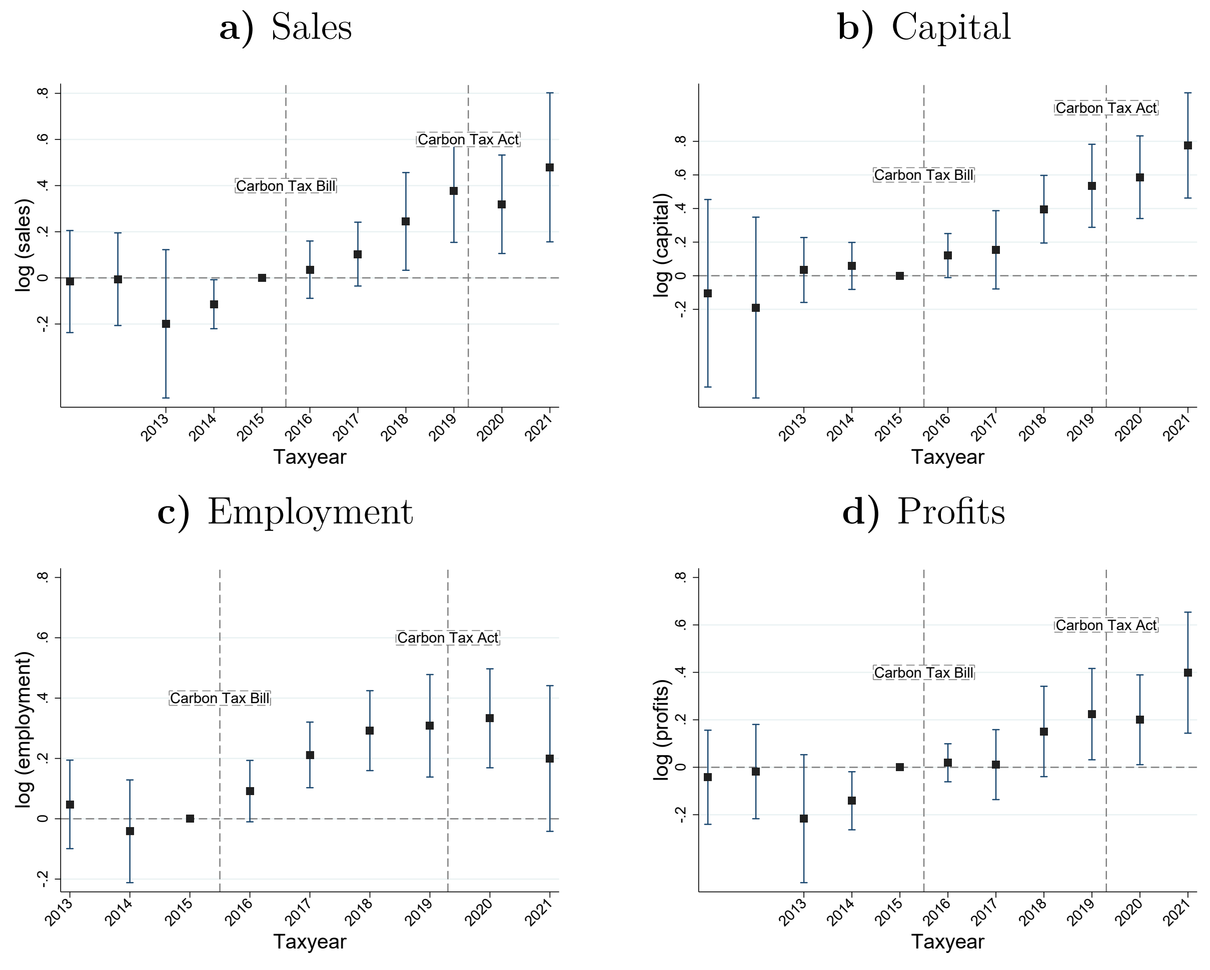

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The vertical bars around the estimates, represented by the squares, show the 95% confidence intervals. The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

Figure 3.B.5: Main event study results using 2013 instead of 2015 as the baseline year

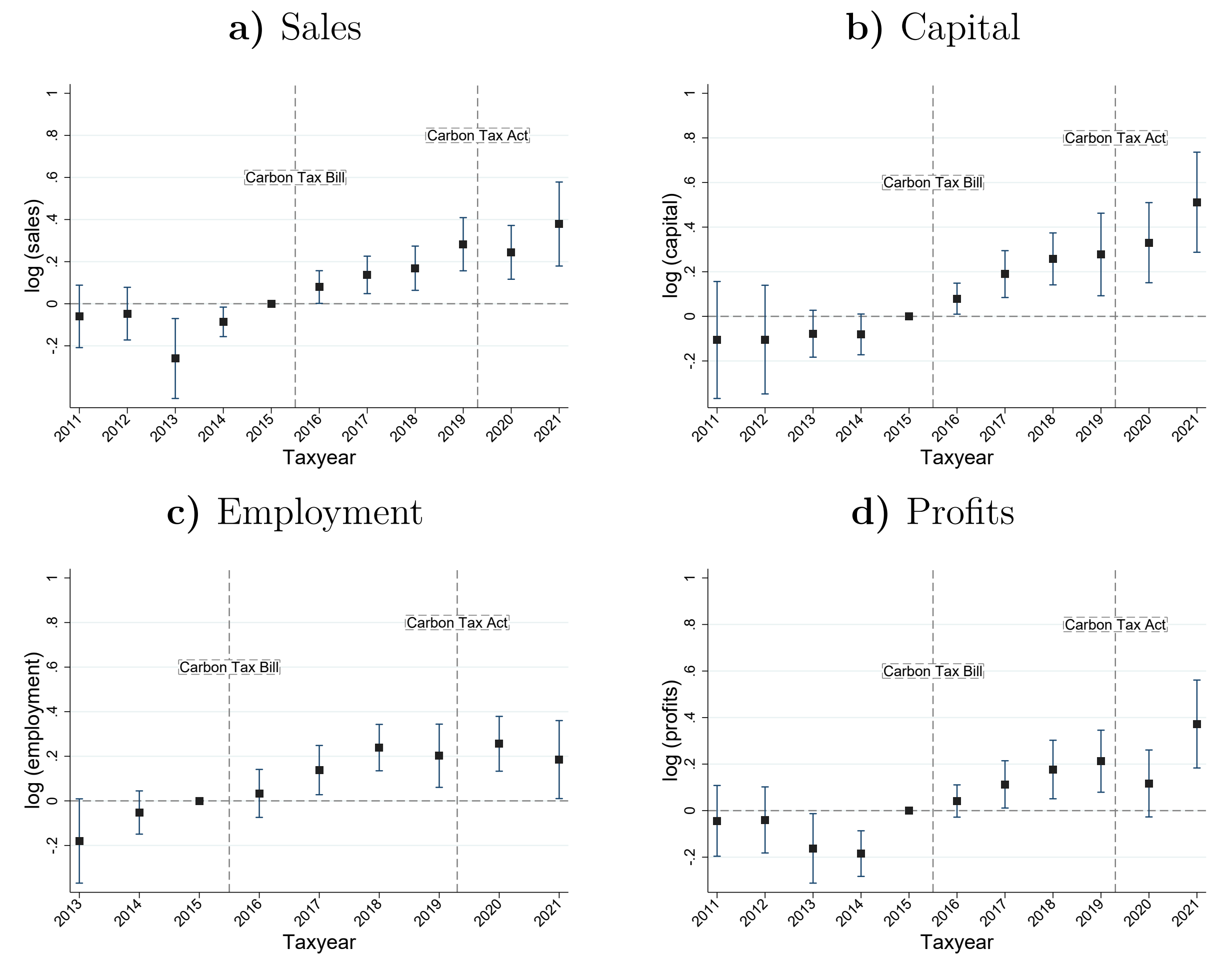

Notes: This figure plots the θ_k coefficients estimated from equation (3.2). The vertical bars around the estimates, represented by the squares, show the 95% confidence intervals. The "tax year" always refers to the previous fiscal year; for example, 2020 corresponds to the period from March 2019 to February 2020. The first dashed line marks the publication of the first draft of the Carbon Tax Bill, and the second dashed line marks the implementation of the Carbon Tax Act in June 2019. Data Source: SARS.

Table 3.B.1: Composition of “synthetic South Africa”

Percent (%)	Countries
66.7	United Arab Emirates
3.0	India
2.9	Trinidad and Tobago
2.0	Vietnam
0.5	Saudi Arabia
0.4	Oman
0.2	Lebanon Haiti Suriname Madagascar Macao Zimbabwe Yemen Kuwait Lesotho Grenada
0.2	Cameroon Burundi Jamaica Aruba Comoros Brunei Darussalam Montserrat Anguilla Burkina Faso Turkmenistan
0.2	Cape Verde Dominica Nicaragua Bahamas Qatar Tajikistan Macedonia, the former Yugoslav Republic of Costa Rica Guinea Seychelles Congo
0.2	Niger Gabon Paraguay Egypt Chad Azerbaijan Mozambique Bahrain Botswana Guinea-Bissau
0.2	Kyrgyzstan Uganda Bermuda Guyana Guatemala Ghana Mali Gambia Iraq Cote d'Ivoire
0.2	Mauritania Rwanda British Virgin Islands Senegal Cambodia Togo Sierra Leone Laos Ethiopia Moldova, Republic of
0.2	Angola Namibia Turks and Caicos Islands Bosnia and Herzegovina Peru Mongolia Uruguay Cayman Islands Sri Lanka Benin
0.2	Belarus Kenya Saint Kitts and Nevis Bolivia Armenia Belize Thailand Ecuador Israel Djibouti
0.2	Panama Algeria Liberia El Salvador Bhutan Zambia Sudan Maldives Mauritius Tanzania
0.2	Malaysia Fiji Tunisia Equatorial Guinea Cyprus Saint Vincent and the Grenadines Barbados Central African Republic Pakistan Nepal
0.2	Georgia Swaziland Sao Tome and Principe Dominican Republic Albania Myanmar Morocco Saint Lucia Honduras Jordan
0.2	Antigua and Barbuda Malawi Uzbekistan Democratic Republic Congo Indonesia Brazil Taiwan Australia Hong Kong Turkey
0.2	Nigeria Philippines Bangladesh

Notes: This table features the contributed weights of each country in order to construct a “synthetic South Africa". “Synthetic South Africa" is constructed based on South Africa's average annual emissions in the following IPCC sectors: 1.A.2, 1.B., 2.A., 2.B., 2.C. and 2.D.. In addition average GDP and emissions intensity (emissions/GDP) between 2011-2015 are used to construct a “synthetic South Africa".

3.C Public comments

Figure 3.C.1: Type of commentators

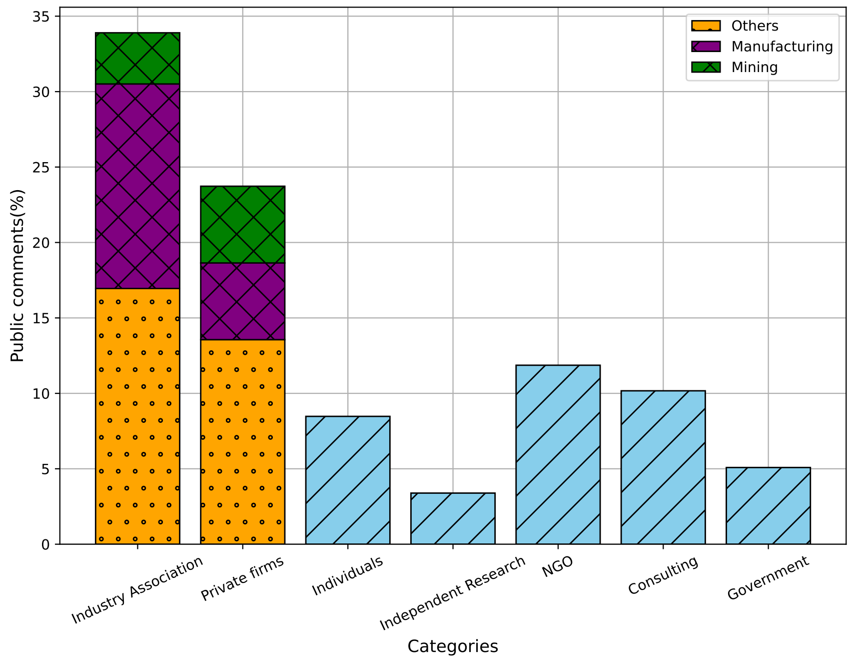

Notes: This figure plots categorization of public comments from the Carbon Tax Bill 2018. Industry associations and private firms are further categorized by their main industry denomination.

Some examples of public comments include:

- It is suggested that the formula to calculate the tax liability as in equation (3.1) is amended to allow sequestration to be deducted not only from fuel combustion emissions but also for process and fugitive emissions (*Not accepted.* Carbon Tax Bill 2018).

- Inconsistency of the tax treatment of waste management activities in the bill, where the provision of the 100 percent allowance for GHG emissions needs to be applied consistently across all sectors and provision should be made accordingly in the Bill (*Accepted. The bill has been amended to allow for a 90 percent tax-free threshold for waste incineration activities.* Carbon Tax Bill 2018).

- For performance allowances, developing an industry benchmark for the lime industry in South Africa may be challenging as there are currently only two large lime manufacturers in the country and three smaller producers. (*Noted.* Carbon Tax Bill 2018).

- The recognition of a renewable energy premium included in the electricity tariff is welcomed. It is proposed that this rebate should be extended to include renewable energy allowed as a recovery of cost by Eskom as well and not limited to the Independent Power Producers only (*Accepted. The bill has been amended to provide the credit for all renewable energy producers.* Carbon Tax Bill 2015).

- Nampak estimates that the emission factor should be closer to 0.1500 to 0.1700 tonnes CO2 per tonne of glass excluding cullet production. This factor is overstated by 25 per cent. The 2006 IPCC Guidelines emissions factor of 0.200 tonnes CO2 per tonnes of glass does not necessarily hold true for all glass production. (*Noted. The emissions factors provided in the Schedule 1 of the carbon tax bill are default emissions factors based on the 2006 IPCC Guidelines and are aligned with the Mandatory reporting regulations and Technical Guidelines. A process to submit alternative emission factors is clearly stated in the NGERs and associated technical guidelines of the DEA.* Carbon Tax Bill 2015).

Bibliography

Abadie, Alberto (2021). “Using synthetic controls: Feasibility, data requirements, and methodological aspects.” In: *Journal of Economic Literature* 59.2, pp. 391–425.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010). “Synthetic control methods for comparative case studies: Estimating the effect of California’s tobacco control program.” In: *Journal of the American Statistical Association* 105.490, pp. 493–505.

Abeberese, Ama Baafra, Ritam Chaurey, and Radhika Menon (2024). “Place-based policies and migration: Evidence from india.” In: *Available at SSRN 4697765*.

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012). “The environment and directed technical change.” In: *American Economic Review* 102.1, pp. 131–166.

African Development Bank (2016). “Chinese Special Economic Zones: Lessons for Africa.” In: *Africa Economic Brief* 7.6.

Agarwal, Bina, Pervesh Anthwal, and Mahesh. Malvika (2021). “How many and which women own land in India? Inter-gender and intra-gender gaps.” In: *The Journal of Development Studies* 57.11, pp. 1807–1829.

Albers, Wulf and Gisela Albers (1983). “On the prominence structure of the decimal system.” In: *Advances in Psychology*. Vol. 16. Elsevier, pp. 271–287.

Alder, Simon, Lin Shao, and Fabrizio Zilibotti (Dec. 2016). “Economic reforms and industrial policy in a panel of Chinese cities.” In: *Journal of Economic Growth* 21.4, pp. 305–349.

Alfaro-Urena, Alonso, Isabela Manelici, and Jose Vasquez (2022). “The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages.” In: *The Quarterly Journal of Economics* 137(3), pp. 1495–1552.

Alkon, Meir (2018). “Do special economic zones induce developmental spillovers? Evidence from India’s states.” In: *World Development* 107, pp. 396–409.

Allingham, Michael G and Agnar Sandmo (1972). “Income tax evasion: A theoretical analysis.” In: *Taxation: critical perspectives on the world economy* 3.1, pp. 323–338.

Almunia, Miguel, Jonas Hjort, Justine Knebelmann, and Lin Tian (2023). “Strategic or confused firms? evidence from “missing” transactions in Uganda.” In: *Review of Economics and Statistics*, pp. 1–10.

Altonji, Joseph, Todd Elder, and Christopher Taber (2005). “Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools.” In: *Journal of Political Economy* 113.1, pp. 151–184.

Aman-Rana, Shan and Clement Minaudier (2024). “Spillovers in State Capacity Building: Evidence from the Digitization of Land Records in Pakistan.” In: *Working Paper*.

Aman-Rana, Shan, Clement Minaudier, and Sandip Sukhtankar (2023). *Informal fiscal systems in developing countries*. Tech. rep. National Bureau of Economic Research.

Amirapu, Amrit and Michael Gechter (2020). “Labor regulations and the cost of corruption: Evidence from the Indian firm size distribution.” In: *Review of Economics and Statistics* 102.1, pp. 34–48.

Anagol, Santosh, Allan Davids, Benjamin B Lockwood, and Tarun Ramadorai (2022). *Diffuse Bunching with Frictions: Theory and Estimation*. Centre for Economic Policy Research.

Anderson, Soren, Ioana Marinescu, and Boris Shor (2023). “Can Pigou at the Polls Stop Us Melting the Poles?” In: *Journal of the Association of Environmental and Resource Economists* 10.4, pp. 903–945.

Andersson, Julius J. (2019). “Carbon Taxes and CO2 Emissions: Sweden as a Case Study.” In: *American Economic Journal: Economic Policy* 11.4, pp. 1–30.

Angelucci, Manuela and Giacomo De Giorgi (2009). “Indirect effects of an aid program: how do cash transfers affect ineligibles’ consumption?” In: *American Economic Review* 99.1, pp. 486–508.

Asher, Sam, Tobias Lunt, Ryu Matsuura, and Paul Novosad (2021). “Development research at high geographic resolution: An analysis of night lights, firms, and poverty in India using the SHRUG Open Data Platform.” In: *World Bank Policy Research Paper*. Policy Research Working Papers 9540.

Asher, Sam and Paul Novosad (2020a). “Rural roads and local economic development.” In: *American Economic Review* 110.3, pp. 797–823.

— (2020b). “Rural roads and local economic development.” In: *American Economic Review* 110(3), pp. 797–823.

ATAF, African Tax Administration Forum (2018). *African Tax Outlook 2018*. Tech. rep. African Tax Administration Forum.

Bacchetta, Marc, Ekkehard Ernst, and Juna P. Bustamante (2009). *Globalization and informal jobs in developing countries*. Tech. rep. Geneva: ILO & WTO.

Bachas, Pierre and Mauricio Soto (2021). “Corporate taxation under weak enforcement.” In: *American Economic Journal: Economic Policy* 13.4, pp. 36–71.

Baker, Lucy (2022). *The Political Economy of South Africa’s Carbon Tax*. Institute of Development Studies.

Balan, Pablo, Augustin Bergeron, Gabriel Tourek, and Jonathan L Weigel (2022). “Local elites as state capacity: How city chiefs use local information to increase tax compliance in the democratic republic of the Congo.” In: *American Economic Review* 112.3, pp. 762–797.

Barrows, Geoffrey and Hélène Ollivier (2021). “Foreign demand, developing country exports, and CO2 emissions: Firm-level evidence from India.” In: *Journal of Development Economics* 149, p. 102587.

Basurto, Maria Pia, Pascaline Dupas, and Jonathan Robinson (2020). “Decentralization and efficiency of subsidy targeting: Evidence from chiefs in rural Malawi.” In: *Journal of Public Economics* 185, p. 104047.

Becker, Johannes and Nadine Riedel (2012). "Cross-border tax effects on affiliate investment— evidence from European multinationals." In: *European Economic Review* 56(3), pp. 436–450.

Becker, Sascha O., Peter H. Egger, and Maximilian Von Ehrlich (2013). "Absorptive capacity and the growth and investment effects of regional transfers: A regression discontinuity design with heterogeneous treatment effects." In: *American Economic Journal: Economic Policy* 5.4, pp. 29–77.

Beegle, Kathleen, Joachim De Weerdt, and Stefan Dercon (2011). "Migration and economic mobility in Tanzania: Evidence from a tracking survey." In: *The Review of Economics and Statistics* 93(3), pp. 1010–1033.

Besley, Timothy and Torsten Persson (2013). "Taxation and development." In: *Handbook of Public Economics*. Vol. 5. Elsevier, pp. 51–110.

— (2014). "Why do developing countries tax so little?" In: *Journal of Economic Perspectives* 28.4, pp. 99–120.

Best, Michael Carlos, Anne Brockmeyer, Henrik Jacobsen Kleven, Johannes Spinnewijn, and Mazhar Waseem (2015). "Production versus revenue efficiency with limited tax capacity: theory and evidence from Pakistan." In: *Journal of Political Economy* 123.6, pp. 1311–1355.

Binswanger-Mkhize, Hans P (2013). "The stunted structural transformation of the Indian economy: Agriculture, manufacturing and the rural non-farm sector." In: *Economic and Political weekly* 48.26-27, pp. 5–13.

Blackwell, Matthew, Stefano M Iacus, Gary King, and Giuseppe Porro (2009). "Coarsened exact matching in Stata." In: *The Stata Journal* 9.4, pp. 524–546.

Blakeslee, David, Ritam Chaurey, Ram Fishman, and Samreen Malik (2022). "Land rezoning and structural transformation in rural India: Evidence from the industrial areas program." In: *World Bank Economic Review*. Policy Research Working Papers 36.3, pp. 488–513.

Blattman, Christopher, Julian Jamison, Tricia Koroknay-Palicz, Katherine Rodrigues, and Margaret Sheridan (2016). "Measuring the measurement error: A method to qualitatively validate survey data." In: *Journal of Development Economics* 120, pp. 99–112.

Bodemann, Balz and Kay Axhausen (2012). "Destination choice for relocating firms: A discrete choice model for the St. Gallen region, Switzerland." In: *Papers in Regional Science* 91.2, pp. 319–341.

Bosworth, Barry and Susan M Collins (2008). "Accounting for Growth: Comparing China and India." In: *Journal of Economic Perspectives* 22.1, pp. 45–66.

Brussevich, Mariya and Era Dabla-Norris (2020). "Socio-economic spillovers from Special Economic Zones: Evidence from Cambodia." In: *IMF Working Papers* 170.

Burger, Philippe and Estian Calitz (2021). "Covid-19, economic growth and South African fiscal policy." In: *South African Journal of Economics* 89.1, pp. 3–24.

Bushnell, James B., Howard Chong, and Erin T. Mansur (2013). "Profiting from Regulation: Evidence from the European Carbon Market." In: *American Economic Journal: Economic Policy* 5.4, pp. 78–106.

Bussmann, Margit (2009). "The effect of trade openness on women's welfare and work life." In: *World Development* 37.6, pp. 1027–1038.

Busso, Matias, Jesse Gregory, and Patrick Kline (2013). "Assessing the incidence and efficiency of a prominent place based policy." In: *American Economic Review* 103.2, pp. 897–947.

Calel, Raphael (2020). "Adopt or Innovate: Understanding Technological Responses to Cap-and-Trade." In: *American Economic Journal: Economic Policy* 12.3, pp. 170–201.

Calel, Raphael and Antoine Dechezlepretre (2016). "Environmental Policy and Directed Technological Change: Evidence from the European Carbon Market." In: *Review of Economics and Statistics* 98.1, pp. 173–191.

Callaway, Brantly and Pedro HC Sant'Anna (2021). "Difference-in-differences with multiple time periods." In: *Journal of Econometrics* 225.2, pp. 200–230.

Carbon Tax Act (2019). *Carbon Tax Act*. Pretoria: Government of the Republic of South Africa.

Caritas (2025). *Press Release: Closure of USAID Foreign Aid will kill millions*. https://www.caritas.org/2025/02/closure-of-usaid-foreign-aid-will-kill-millions/. (Last Accessed: March 11th, 2025).

Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen L. Palmer (2000). “Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?” In: *Journal of Political Economy* 108.6, pp. 1292–1326.

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal (2022). *Ghosting the Tax Authority: Fake Firms and Tax Fraud*. Tech. rep. National Bureau of Economic Research.

Chaurey, Ritam (2017). “Location-based tax incentives: Evidence from India.” In: *Journal of Public Economics* 156, pp. 101–120.

Chen, Binkai, Ming Lu, Christopher Timmins, and Kuanhu Xiang (2019). “Spatial misallocation: Evaluating place-based policies using a natural experiment in China.” In: *NBER Working Paper No. 26148*.

Chodorow-Reich, Gabriel (May 2019). “Geographic Cross-Sectional Fiscal Spending Multipliers: What Have We Learned?” In: *American Economic Journal: Economic Policy* 11.2, pp. 1–34.

Chodorow-Reich, Gabriel, Matthew Smith, Owen M. Zidar, and Eric Zwick (2024). “Tax Policy and Investment in a Global Economy.” In: *NBER Working Paper* 32180.

Clay, Karen, Akshaya Jha, Joshua A. Lewis, and Edson R. Severnini (2024). “Impacts of the Clean Air Act on the Power Sector from 1938-1994: Anticipation and Adaptation.” In: *NBER Working Paper No. 28962*.

Colmer, Jonathan, Ralf Martin, Mirabelle Muûls, and Ulrich J Wagner (2024). “Does pricing carbon mitigate climate change? Firm-level evidence from the European Union Emissions Trading System.” In: *Review of Economic Studies*, rdae055.

Combes, Pierre-Philippe and Laurent Gobillon (2015). "The empirics of agglomeration economies." In: *Handbook of regional and urban economics.* Vol. 5. Elsevier, pp. 247–348.

Conley, Timothy G. (1999). "GMM estimation with cross sectional dependence." In: *Journal of Econometrics* 92.1, pp. 1–45.

Crippa, Monica, Diego Guizzardi, Federico Pagani, Marcello Schiavina, Michele Melchiorri, Enrico Pisoni, Francesco Graziosi, Marilena Muntean, Joachim Maes, Lewis Dijkstra, et al. (2024). "Insights into the spatial distribution of global, national, and subnational greenhouse gas emissions in the Emissions Database for Global Atmospheric Research (EDGAR v8. 0)." In: *Earth System Science Data* 16.6, pp. 2811–2830.

Criscuolo, Chiara, Ralf Martin, Henry G. Overman, and John Van Reenen (2019). "Some causal effects of an industrial policy." In: *American Economic Review* 109.1, pp. 48–85.

Cui, Jingbo, Chunhua Wang, Junjie Zhang, and Yang Zheng (2021). "The effectiveness of China's regional carbon market pilots in reducing firm emissions." In: *Proceedings of the National Academy of Sciences* 118.52, e2109912118.

Cui, Jingbo, Junjie Zhang, and Yang Zheng (2023). "Carbon Price, Innovation, and Firm Competitiveness." In: *Innovation, and Firm Competitiveness (February 20, 2023).*

Das, Sonali, Sonali Jain-Chandra, Kalpana Kochhar, and Naresh Kumar (2015). "Women workers in India: why so few among so many?" In: *IMF Working Papers* 055.

Dechezlepretre, Antoine, Daniel Nachtigall, and Frank Venmans (2023). "The joint impact of the European Union emissions trading system on carbon emissions and economic performance." In: *Journal of Environmental Economics and Management* 118, p. 102758.

Desai, Mihir, Fritz Foley, and James Hines Jr. (2009). "Domestic effects of the foreign activities of US multinationals." In: *American Economic Journal: Economic Policy* 1.1, pp. 181–203.

Di Maria, Corrado, Ian Lange, and Edwin van der Werf (2014). "Should we be worried about the green paradox? Announcement effects of the Acid Rain Program." In: *European Economic Review* 69, pp. 143–162.

Douenne, Thomas and Adrien Fabre (2022). “Yellow vests, pessimistic beliefs, and carbon tax aversion.” In: *American Economic Journal: Economic Policy* 14.1, pp. 81–110.

Duflo, Esther (2012). “Women empowerment and economic development.” In: *Journal of Economic Literature* 50.4, pp. 1051–1079.

Duranton, Gilles and Anthony J Venables (2018). “Place-based policies for development.” In: *World Bank Policy Research Paper* 8410.

Ebrahim, Amina and Chris Axelson (2019). *The creation of an individual panel using administrative tax microdata in South Africa.* 2019/27. WIDER Working Paper.

Ebrahim, Amina, Friedrich Kreuser, and Michael Kilumelume (2021). *The guide to the CIT-IRP5 panel version 4.0.* Tech. rep. WIDER Working Paper.

Ehrlich, Maximilian V. and Tobias Seidel (2018). “The persistent effects of place-based policy: Evidence from the West-German Zonenrandgebiet.” In: *American Economic Journal: Economic Policy* 10.4, pp. 344–374.

Eichengreen, Barry and Poonam Gupta (2011). “The service sector as India’s road to economic growth.” In: *NBER Working Paper No. 16757*.

Einiö, Elias and Henry G. Overman (2020). “The effects of supporting local business: Evidence from the UK.” In: *Regional Science and Urban Economics* 83.

Ewald, Jens, Thomas Sterner, and Erik Sterner (2022). “Understanding the resistance to carbon taxes: Drivers and barriers among the general public and fuel-tax protesters.” In: *Resource and Energy Economics* 70, p. 101331.

Farole, Thomas and Lotta Moberg (2014). “It worked in China, so why not in Africa? The political economy challenge of Special Economic Zones.” In: *UNU-WIDER Working Paper*. WIDER Working Paper 152.

Feenstra, Robert C, Robert Inklaar, and Marcel P Timmer (2015). “The next generation of the Penn World Table.” In: *American economic review* 105.10, pp. 3150–3182.

Fellner, Gerlinde, Rupert Sausgruber, and Christian Traxler (2013). “Testing enforcement strategies in the field: Threat, moral appeal and social information.” In: *Journal of the European Economic Association* 11.3, pp. 634–660.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur (2012). “What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NOx Trading Program.” In: *American Economic Review* 102.2, pp. 965–93.

Fried, Stephie, Kevin Novan, and William B Peterman (2022). *Understanding the Inequality and Welfare Impacts of Carbon Tax Policies*. Tech. rep. Working Paper.

Fuchs, Maximilian, Johannes Stroebel, and Julian Terstegge (2024). *Carbon vix: Carbon price uncertainty and decarbonization investments*. Tech. rep. National Bureau of Economic Research.

Gadenne, Lucie and Monica Singhal (2014). “Decentralization in developing economies.” In: *Annu. Rev. Econ.* 6.1, pp. 581–604.

Gallé, Johannes, Daniel Overbeck, Nadine Riedel, and Tobias Seidel (2024). “Place-based policies, structural change and female labor: Evidence from India’s Special Economic Zones.” In: *Journal of Public Economics* 240, p. 105259.

Gaspar, Vitor, Laura Jaramillo, and Philippe Wingender (2016). “Tax Capacity and Growth: Is there a Tipping Point?” In: *IMF Working Paper* 16.

Global Carbon Budget (2023). “Global Carbon Budget 2023.” In: *Global Carbon Budget Project*.

Gobillon, Laurent, Thierry Magnac, and Harris Selod (2012). “Do unemployed workers benefit from enterprise zones? The French experience.” In: *Journal of Public Economics* 96.9-10, pp. 881–892.

Goldberg, Pinelopi Koujianou, Amit Kumar Khandelwal, Nina Pavcnik, and Petia Topalova (2010). “Imported intermediate inputs and domestic product growth: Evidence from India.” In: *The Quarterly Journal of Economics* 125.4, pp. 1727–1767.

Gollin, Douglas, David Lagakos, and Michael Waugh (2009). “The agricultural productivity gap.” In: *Quarterly Journal of Economics* 129(2), pp. 939–993.

Gollin, Douglas, David Lagakos, and Michael E. Waugh (2014). "The agricultural productivity gap." In: *The Quarterly Journal of Economics* 129.2, pp. 939–993.

Gordon, Roger and Wei Li (2009a). "Tax structures in developing countries: Many puzzles and a possible explanation." In: *Journal of Public Economics* 93.7-8, pp. 855–866.

— (2009b). "Tax structures in developing countries: Many puzzles and a possible explanation." In: *Journal of Public Economics* 93.7-8, pp. 855–866.

Gordon, Roger H (2023). *Carbon Taxes: Many Strengths but Key Weaknesses.* Tech. rep. National Bureau of Economic Research.

Görg, Holger and Alina Mulyukova (2024). "Place-based policies and firm performance: Evidence from Special Economic Zones in India." In: *European Economic Review* 165, p. 104752.

Grant, Matthew (2020). "Why special economic zones? Using trade policy to discriminate across importers." In: *American Economic Review* 110.5, pp. 1540–1571.

Greenstone, Michael (2002). "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufacturers." In: *Journal of Political Economy* 110.6, pp. 1175–1219.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti (2010). "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." In: *Journal of Political Economy* 118.3, pp. 536–598.

Hasan, Rana, Yi Jiang, and Radine Michelle Rafols (2021). "Place-based preferential tax policy and industrial development: Evidence from India's program on industrially backward districts." In: *Journal of Development Economics* 150, p. 102621.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil (2012). "Measuring economic growth from outer space." In: *American Economic Review* 102.2, pp. 994–1028.

Hindriks, Jean, Michael Keen, and Abhinay Muthoo (1999). "Corruption, extortion and evasion." In: *Journal of Public Economics* 74.3, pp. 395–430.

Hoy, Christopher, Thiago Scot, Alex Oguso, Anna Custers, Daniel Zalo, Ruggero Doino, Jonathan Karver, and Orgeira Nicolas Pillai (2024). "Trade-offs in the Design of Simplified Tax Regimes: Evidence from Sub-Saharan Africa." In: *World Bank Policy Research Paper No. 10909*. Policy Research Working Papers.

Hyun, Yeseul and Shree Ravi (2018). "Place-based development: Evidence from special economic zones in India." In: *IED Working Paper No. 306*.

Iacus, Stefano M, Gary King, and Giuseppe Porro (2012). "Causal inference without balance checking: Coarsened exact matching." In: *Political analysis* 20.1, pp. 1–24.

Idson, Todd L. and Walter Y. Oi (1999). "Workers Are More Productive in Large Firms." In: *American Economic Review* 89.2, pp. 104–108.

International Labour Organization (2013). *India Labour Market Update*. Tech. rep. Geneva: ILO.

International Monetary Fund (2024). "Stepping up Domestic Resource Mobilization: a new joint initiative from the IMF and World Bank." In: *https://www.imf.org/-/media/Files/Research/imf-and-g20/2024/domestic-resource-mobilization.ashx (Last accessed: October 24th, 2024)*.

Janssen, Maarten CW (2001). "Rationalizing focal points." In: *Theory and Decision* 50, pp. 119–148.

— (2006). "On the strategic use of focal points in bargaining situations." In: *Journal of Economic Psychology* 27.5, pp. 622–634.

Jia, Junxue, Guangrong Ma, Cong Qin, and Liyan Wang (2020). "Place-based policies, state-led industrialisation, and regional development: Evidence from China's Great Western Development Programme." In: *European Economic Review* 123, p. 103398.

Känzig, DR (2023). "The unequal economic consequences of carbon pricing (Working Paper No. 31221)." In: *National Bureau of Economic Research. https://doi. org/10* 3386, w31221.

Kar, Saibal and Yashika Khattar (2023). "Does Minimum Wage Affect Informal Jobs across States in India?" In: *Economic and Political Weekly* 58.7.

Khan, Adnan Q, Asim I Khwaja, and Benjamin A Olken (2016). “Tax farming redux: Experimental evidence on performance pay for tax collectors.” In: *The Quarterly Journal of Economics* 131.1, pp. 219–271.

Klasen, Stephan and Janneke Pieters (2015). “What Explains the Stagnation of Female Labor Force Participation in Urban India?” In: *The World Bank Economic Review* 29.3, pp. 449–478.

Kleven, Henrik Jacobsen (2016). “Bunching.” In: *Annual Review of Economics* 8, pp. 435–464.

Kleven, Henrik Jacobsen, Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez (2011). “Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark.” In: *Econometrica* 79.3, pp. 651–692.

Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez (2016). “Why can modern governments tax so much? An agency model of firms as fiscal intermediaries.” In: *Economica* 83.330, pp. 219–246.

Kleven, Henrik Jacobsen and Mazhar Waseem (2013). “Using notches to uncover optimization frictions and structural elasticities: Theory and evidence from Pakistan.” In: *The Quarterly Journal of Economics* 128.2, pp. 669–723.

Kline, Patrick and Enrico Moretti (2014). “Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority.” In: *The Quarterly Journal of Economics* 129.1, pp. 275–331.

Koster, Hans, Fang Fang Cheng, Michiel Gerritse, and Frank van Oort (2019). “Place-based policies, firm productivity, and displacement effects: Evidence from Shenzhen, China.” In: *Journal of Regional Science* 59, pp. 187–213.

Laitner, John (2000). “Structural change and economic growth.” In: *The Review of Economic Studies* 67.3, pp. 545–561.

LaPorta, Rafael and Andrei Shleifer (2014). “Informality and development.” In: *Journal of Economic Perspectives* 28.3, pp. 109–126.

Lemoine, Derek (2017). "Green Expectations: Current Effects of Anticipated Carbon Pricing." In: *Review of Economics and Statistics* 99.3, pp. 499–513.

Levien, Michael (2012). "The Land Question: Special Economic Zones and the Political Economy of Dispossession in India." In: *The Journal of Peasant Studies* 39.3, pp. 933–969.

Lu, Fangwen, Weizeng Sun, and Jianfeng Wu (2022). "Special Economic Zones and Human Capital Investment: 30 Years of Evidence from China." In: *American Economic Journal: Economic Policy*.

Lu, Yi, Jin Wang, and Lianming Zhu (2019). "Place-based policies, creation, and agglomeration economies: Evidence from China's economic zone program." In: *American Economic Journal: Economic Policy* 11.3, pp. 325–360.

Lueck, Dean and Jeffrey A. Michael (2003). "Preemptive Habitat Destruction under the Endangered Species Act." In: *Journal of Law & Economics* 46.1, pp. 27–60.

Marron, Donald B and Eric J Toder (2014). "Tax policy issues in designing a carbon tax." In: *American Economic Review* 104.5, pp. 563–568.

Martin, Ralf, Laure B. de Preux, and Ulrich J. Wagner (2014). "The impact of a carbon tax on manufacturing: Evidence from microdata." In: *Journal of Public Economics* 117, pp. 1–14.

Martin, Ralf, Mirabelle Muuls, Laure B. de Preux, and Ulrich J. Wagner (2014). "Industry Compensation under Relocation Risk: A Firm-Level Analysis of the EU Emissions Trading Scheme." In: *American Economic Review* 104.8, pp. 2482–2508.

Mascagni, Giulia and Christopher Nell (2022). "Tax compliance in Rwanda: Evidence from a message field experiment." In: *Economic Development and Cultural Change* 70.2, pp. 587–623.

Mascagni, Giulia, Fabrizio Santoro, Denis Mukama, John Karangwa, and Napthal Hakizimana (2022). "Active ghosts: Nil-filing in Rwanda." In: *World Development* 152, p. 105806.

McCaig, Brian and Nina Pavcnik (2013). "Moving out of agriculture: structural change in Vietnam." In: *NBER Working Paper No. 19616*.

McCaig, Brian and Nina Pavcnik (2021). "Entry and exit of informal firms and development." In: *IMF Economic Review* 69.3, pp. 540–575.

McMillan, Margaret, Dani Rodrik, Dani, and Inigo Verduzco-Gallo (2014). "Globalization, structural change, and productivity growth, with an update on Africa." In: *World Development* 63, pp. 11–32.

Megginson, William L and Jeffry M Netter (2001). "From state to market: A survey of empirical studies on privatization." In: *Journal of Economic Literature* 39.2, pp. 321–389.

Mehrotra, Santosh (2019). "Informal employment trends in the Indian economy: persistent informality, but growing positive development." In: *ILO Working Paper No. 254*.

Meiyappan, P., P. S. Roy, A. Soliman, T. Li, P. Mondal, S. Wang, and A.K. Jain (2018). *India Village-Level Geospatial Socio-Economic Data Set: 1991, 2001.* Tech. rep. Palisades, NY: NASA Socioeconomic Data and Applications Center (SEDAC).

Metcalf, Gilbert E (2021). "Carbon taxes in theory and practice." In: *Annual Review of Resource Economics* 13, pp. 245–265.

Ministry of Finance (2015). *Revenue foregone under the Central Tax System: Financial Years 2012-13 and 2013-14.* Tech. rep. New Delhi: India Budget.

Ministry of Labour and Social Security (2018). *An Analysis of the informal economy in Zambia.* Tech. rep. Lusaka: Central Statistical Office.

Mukherjee, Arpita, Parthapratim Pal, Saubhik Deb, Subhobrota Ray, and Tanu M. Goyal (2016). *Special economic zones in India. Status, issues and potential.* Vol. 43. 28. Springer.

Munshi, Kaivan and Mark Rosenzweig (2016). "Networks and misallocation: Insurance, migration, and the rural-urban wage gap." In: *American Economic Review* 106(1), pp. 46–98.

NCEUS (2009). *The challenge of employment in India: An informal economy perspective.* Tech. rep. New Delhi: National Comission for Enterprises in the Unorganised Sector.

Neumark, David and Jed Kolko (2010). "Do enterprise zones create jobs? Evidence from California's enterprise zone program." In: *Journal of Urban Economics* 68.1, pp. 1–19.

Neumark, David and Helen Simpson (2015). "Place-based policies." In: *Handbook of Regional and Urban Economics*. Vol. 5. Elsevier B.V., pp. 1197–1287.

NOAA (2013). "Version 4 DMSP-OLS Nighttime Lights Time Series." In: *Data retrieved from National Oceanic and Atmospheric Administration, URL: https://www.ngdc.noaa.gov/eog*. (Last Accessed: June 24th, 2023).

Oi, Walter Y and Todd L Idson (1999). "Firm size and wages." In: *Handbook of labor economics* 3, pp. 2165–2214.

Okunogbe, Oyebola and Victor Pouliquen (2022). "Technology, taxation, and corruption: evidence from the introduction of electronic tax filing." In: *American Economic Journal: Economic Policy* 14.1, pp. 341–372.

Okunogbe, Oyebola and Gabriel Tourek (2024). "How Can Lower-Income Countries Collect More Taxes? The Role of Technology, Tax Agents, and Politics." In: *Journal of Economic Perspectives* 38.1, pp. 81–106.

Olken, Benjamin A and Monica Singhal (2011). "Informal taxation." In: *American Economic Journal: Applied Economics* 3.4, pp. 1–28.

Ozler, Sule (2000). "Export orientation and female share of employment: Evidence from Turkey." In: *World Development* 28.7, pp. 1239–1248.

Parwez, Sazzad and Vinod Sen (2016). "Special Economic Zone, Land Acquisition, and Impact on Rural India." In: *Emerging Economy Studies* 2.2, pp. 223–239.

Pieterse, Duncan, Elizabeth Gavin, and C Friedrich Kreuser (2018). "Introduction to the South African Revenue Service and National Treasury Firm-Level Panel." In: *South African Journal of Economics* 86, pp. 6–39.

Pope, Devin G, Jaren C Pope, and Justin R Sydnor (2015). "Focal points and bargaining in housing markets." In: *Games and Economic Behavior* 93, pp. 89–107.

Porter, Michael E and Claas van der Linde (1995). "Toward a new conception of the environment-competitiveness relationship." In: *Journal of Economic Perspectives* 9.4, pp. 97–118.

Qu, H, S Suphachalasai, S Thube, and S Walker (2023). “South Africa Carbon Pricing and Climate Mitigation Policy.” In: *Selected Issues Papers*.

Rama, Martin (2003). “Globalization and workers in developing countries.” In: *World Bank Policy Research Paper No. 2958*.

Rossi, Federica and Magdalena Dej (2020). “Where do firms relocate? Location optimisation within and between Polish metropolitan areas.” In: *The Annals of Regional Science* 64.3, pp. 615–640.

Saez, Emmanuel (2010). “Do taxpayers bunch at kink points?” In: *American economic Journal: economic policy* 2.3, pp. 180–212.

Saini, Shweta, Ashok Gulati, Joachim von Braun, and Lukas Kornher (2020). “Indian farm wages: Trends, growth drivers and linkages with food prices.” In: *ZEF Discussion Papers on Development Policy* 301.

Schelling, Thomas (1960). *The Theory of Conflict*. Harvard University Press.

SEZ Act (2005). *The Special Economic Zones Act*. New Delhi: Ministry of Law and Justice.

Shenoy, Ajay (2018). “Regional development through place-based policies: Evidence from a spatial discontinuity.” In: *Journal of Development Economics* 130, pp. 173–189.

Shimeles, Abebe, Daniel Zerfu Gurara, and Firew Woldeyes (2017). “Taxman’s dilemma: coercion or persuasion? Evidence from a randomized field experiment in Ethiopia.” In: *American Economic Review* 107.5, pp. 420–424.

Sinn, Hans-Werner (2012). *The green paradox: a supply-side approach to global warming*. MIT press.

Slemrod, Joel (2013). “Buenas notches: lines and notches in tax system design.” In: *eJTR* 11, p. 259.

Slemrod, Joel, Marsha Blumenthal, and Charles Christian (2001). “Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota.” In: *Journal of Public Economics* 79.3, pp. 455–483.

South African Revenue Service (2021). *Frequently Asked Questions – Carbon Tax*. `https://www.sars.gov.za/wp-content/uploads/Docs/CarbonTax/Frequently-Asked-Questions-Carbon-Tax_3.0.pdf`. (Last accessed: January 11th, 2025).

Srivastava, Nisha and Ravi Srivastava (2015). "Women, Work, and Employment Outcomes in Rural India." In: *Economic and Political weekly* 45.28, pp. 7–8.

Steckel, Jan C, Ira I Dorband, Lorenzo Montrone, Hauke Ward, Leonard Missbach, Fabian Hafner, Michael Jakob, and Sebastian Renner (2021). "Distributional impacts of carbon pricing in developing Asia." In: *Nature Sustainability* 4.11, pp. 1005–1014.

Steenkamp, Lee-Ann (2022). "South Africa's carbon tax rate goes up but emitters get more time to clean up." In: *The conversation*.

Stiebale, Joel and Dev Vencappa (2018). "Acquisitions, markups, efficiency, and product quality: Evidence from India." In: *Journal of International Economics* 112, pp. 70–87.

Strand, Jon (2020). *Supporting Carbon Tax Implementation in Developing Countries through Results-Based Payments for Emissions Reductions*. The World Bank.

Sud, Nikita (2014). "Governing India's land." In: *World Development* 60, pp. 43–56.

Sun, Liyang and Sarah Abraham (2021). "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." In: *Journal of Econometrics* 225.2, pp. 175–199.

The Economist (2015). "Not So Special." In: *April 4th edition. Available online at `https://www.economist.com/leaders/2015/04/04/not-so-special` (Last Accessed: August 20th, 2024).*

Timilsina, Govinda R (2022). "Carbon taxes." In: *Journal of Economic Literature* 60.4, pp. 1456–1502.

Tourek, Gabriel (2022). "Targeting in tax behavior: Evidence from Rwandan firms." In: *Journal of Development Economics* 158, p. 102911.

Udry, Christopher (1994). "Risk and insurance in a rural credit market: An empirical investigation in northern Nigeria." In: *The Review of Economic Studies* 61.3, pp. 495–526.

UNCTAD (2019). *World investment report: Special economic zones*. Tech. rep. New York: United Nations Conference on Trade and Development.

UNFCCC (2024). *About Carbon Pricing*. `https://unfccc.int/about-us/regional-collaboration-centres/the-ciaca/about-carbon-pricing`. (Last Accessed: January 11th, 2025).

United Nations (2023). *The Least Developed Countries Report*. Tech. rep. geneva: United Nations Conference on Trade and Development.

Uy, Timothy, Kei-Mu Yi, and Jing Zhang (2013). "Structural change in an open economy." In: *Journal of Monetary Economics* 60.6, pp. 667–682.

Walker, W. Reed (2013). "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce." In: *Quarterly Journal of Economics* 128.4, pp. 1787–1835.

Wang, Jin (2013). "The economic impact of special economic zones: Evidence from Chinese municipalities." In: *Journal of Development Economics* 101.1, pp. 133–147.

World Bank (2005). *Doing business in 2005: Removing obstacles to growth*. Tech. rep. Washington D.C.: The World Bank.

— (2008). *Special economic zones: Performance, lessons learned, and implications for zone development*. Tech. rep. Washington D.C.: The World Bank.

— (2011a). *Fostering women's economic empowerment through Special Economic Zones*. Tech. rep. Washington D.C.: The World Bank.

— (2011b). *Risk-based tax audits: Approaches and country experiences*. Tech. rep. World Bank Publications.

— (2012). *World development report 2012: Gender equality and development*. Tech. rep. Washington D.C.: The World Bank.

— (2017). *Special economic zones: An operational review of their impacts*. Tech. rep. Washington D.C.: The World Bank.

— (2019). "Small and Medium Enterprises (SMEs) Finance." In: *https://www.worldbank.org/en/topic/smefinance (Last Accessed: July 27th, 2024)*.

— (2023a). "Agriculture, forestry, and fishing, value added (% of GDP) - India." In: *https://data.worldbank.org/indicator/NV.AGR.TOTL.ZS?locations=IN (Last Accessed: March 20th, 2023)*.

— (2023b). "Employment in agriculture (% of total employment) (modeled ILO estimate) - India:" in: *https://data.worldbank.org/indicator/SL.AGR.EMPL.ZS?locations=IN (Last accessed: June 23rd, 2023)*.

— (2024). *State and Trends of Carbon Pricing*. Tech. rep. Washington D.C.: The World Bank.

— (2025). "Tax revenues as share of GDP." In: *https://data.worldbank.org/indicator/GC.TAX.TOTL.GD.ZS?locations=ZG-8S (Last Accessed: March 10th, 2025)*.

Yamazaki, Akio (2017). "Jobs and climate policy: Evidence from British Columbia's revenue-neutral carbon tax." In: *Journal of Environmental Economics and Management* 83, pp. 197–216.

— (2022). "Environmental taxes and productivity: Lessons from Canadian manufacturing." In: *Journal of Public Economics* 205, p. 104560.

Zambia Revenue Authority (2022). *Tax Statistics in Zambia 2021*. Tech. rep. Lusaka: Central Statistical Office.

Curriculum Vitae

2019–2025	University of Mannheim (Germany)
	Ph.D. in Economics
2017–2019	University of Bonn (Germany)
	M.Sc. in Economics
2013–2017	University of Heidelberg (Germany)
	B.Sc. in Economics