Enhancing Revenue Collection Capacity through Automation: Evidence from a VAT*

Bhanu Gupta† Tejaswi Velayudhan‡ Vedant Vohra§ Ankur Garg¶

March 17, 2026

Abstract

Tax administrations around the world are investing heavily in digitization, creating new opportunities for enforcement. Can low-cost enforcement tools built atop this digital infrastructure increase tax capacity? We study an Indian intervention implementing automated cross-checks to target fraudulent input tax credits. Using administrative data, we show that the policy reduced fraudulent credits and increased tax payments, resulting in a 5 percent reduction in overall evasion at minimal cost. We find little adverse impact on legitimate activity. This type of policy is readily implementable in many developing countries that already have digital filing but are yet to exploit its full potential.

Keywords: Value-Added Tax; Tax Evasion; Tax Enforcement; Firm Networks
JEL Codes: H25, H26, H32, O17

*The authors would like to thank the Delhi Department of Trade and Taxes for their cooperation in making this project possible. We thank the Center for Effective Governance of Indian States (CEGIS) team for their assistance in data procurement and management. We gratefully acknowledge financial assistance from J-PAL and CEGIS. We are grateful to David Agrawal, Julie Cullen, Itzik Fadlon, Roger Gordon, Craig McIntosh, Karthik Muralidharan, and seminar participants at Ashoka University, UC San Diego, UCI, SoCCAM, and the 7th World Bank/ODI Global/IFS Public Finance Conference for their insightful comments.

†Department of Economics, Ashoka University. Email: bhanu.gupta@ashoka.edu.in
‡Department of Economics, University of California, Irvine. Email: tvelayud@uci.edu
§Department of Economics, University of California, San Diego. Email: vohra@ucsd.edu
¶Department of Trade and Taxes, Government of NCTD.

1 Introduction

Information is central to a state's capacity to enforce a tax system.[1] Digitization of tax administration has enabled both the collection and use of vast troves of information in developed and developing countries alike. This digital architecture interacts particularly meaningfully with the Value Added Tax (VAT), whose main advantage over a retail sales tax lies in the information generated. Each transaction between firms generates dual reports—sellers report the transaction as a sale while buyers report it as a purchase—creating independently-verifiable information for the tax authority. Although digitization has spread widely—for example, at least 70% of countries across income levels mandate electronic filing of VAT returns—the use of these data for enforcement still lags.[2] Fewer than 30% of low-income countries perform any automated cross-checks. This paper demonstrates the potential of such automated checks to curb evasion with minimal additional investment, in the context of India's Goods and Services Tax (GST).

A key vulnerability of VAT systems is the provision of fraudulent input tax credits (ITCs): tax-registered sellers provide invoices that generate tax credits for their clients without making payments on those invoices themselves. Fraudulent credits erode 2–4 percent of net VAT collections in low-income countries (about 3 percent in India), with marked spillovers along production chains (Carrillo et al., 2017; Waseem, 2023; Almunia et al., 2024).[3] Beginning in August 2021, the GST online filing portal automatically blocked firms from issuing input tax credits to clients if they had not made their payments for two consecutive months. The reform neither raised existing statutory penalties nor required intervention by tax personnel; instead, it harnessed existing e-filing infrastructure to instantly prevent the generation of some credits in excess of payment. This type of reform is feasible in most developing countries that already mandate e-filing. Yet, the absence of such automated checks suggest that there may be political or administrative costs hindering their adoption, or that they are perceived to have low benefits. While we cannot assess all factors hindering their adoption, we provide evidence on two aspects of the tradeoffs involved: (1) can a feasible and low-cost form of automated enforcement reduce evasion and raise revenue, and (2) does it unintentionally hinder long-term revenue growth by penalizing short-term cash-flow shocks?

We address these questions using granular, administrative data from New Delhi combined with an event-study design.[4] We begin by adopting a data-driven approach to identify potential fake credit generators (FCGs), i.e. taxpayers who generate input tax credits without making

[1] A large literature has shown the empirical and theoretical relevance of information for tax capacity. See for example Kleven et al. (2016); Kumler et al. (2020); Jensen (2022); Pomeranz (2015); Best et al. (2021); Kleven et al. (2011).

[2] All statistics cited in this paragraph are author's calculations using IMF RA-FIT 2022 data, https://data.rafit.org, accessed April 1, 2025.

[3] Evasion and revenue estimates are reported in Lok Sabha Questions (2024) and Goods and Services Tax Network (2023b), respectively.

[4] Delhi contributed 4.4 percent of the total GST revenue in 2023 (Goods and Services Tax Network, 2023b).

associated payments. Taxpayers must report their sales in two separate returns: a *payment-linked* return that triggers tax liability and *credit-linked* return that provides tax credits to their clients. Performing a cross-check of these reports, it is immediately apparent that there are substantial discrepancies favoring the taxpayer, i.e. where credit-linked sales exceed payment-linked sales. Our first finding is that about 12 to 14 percent of these discrepancies arise through the exact channel that the policy targets: non-filing of a payment-return while filing a credit-linked return, suggesting that it could have large effects (Figure 1). Taxpayers are classified as FCGs if they have non-negligible discrepancies in the pre-treatment period, between January and June 2021. Although these firms are about 0.3 percent of all taxpayers, they account for 60 percent of excess credits in the economy and their clients are 3.5 percent of all firms.[5]

We find that FCGs stop providing credits unbacked by payments in response to the policy. Compared to non-FCG firms with otherwise similar pre-policy filing behavior, FCGs become 3 percentage points more likely to cease filing credit returns (a 50 percent increase relative to baseline non-filing rates). They are no more likely to start making payments, the other avenue through which taxpayers could have become compliant. Instead, targeted FCGs seem to cease operations and risk de-registration: they double the rate at which they stop all filing.

A counter-balancing concern is that the policy is causing firms to shut down that otherwise might have contributed to tax revenue in the long-run. For example, a firm that was cash-constrained and unable to make their tax payments may nevertheless have wanted to pass on credits expecting to make their tax payments when they eventually had the liquidity. The policy would make it harder for such firms to continue operating, thereby depressing future revenue growth. Our results on the clients of FCGs suggest that this is unlikely to be the case. These downstream firms increase net tax payments by 5 percent over the next eight months and show no decline in sales. They do not replace lost FCG credits even eight months after the policy change. If FCGs were otherwise legitimate firms shutting down activity, we would expect that lower input purchases among clients would also translate to some decline in sales, which we do not find. Our findings suggest that the blocking policy curtailed fake credits without adversely affecting legitimate economic activity.

Back-of-the-envelope estimates suggest that the policy generated annual revenue gains in the range of INR 1.6 billion to 5.4 billion[6]—which is 2 to 5 percent of all suspected evasion or 1% of net GST collections in Delhi—at negligible additional administrative or compliance cost. The current literature reveals that effectiveness of enforcement tools is proportional to the resource costs involved. For instance, notifying the clients of fake firms—identified after an intensive and costly "antighost" initiative—increased tax payments by 13.7% among treated firms in Ecuador (Carrillo et al., 2022). In Pakistan, use of sophisticated software that verifies

[5]Our estimates compare favorably with the national estimates published by the tax department in December 2023, which found 29,273 fake firms (Lok Sabha Questions, 2024), about 0.23% of all GST taxpayers (Goods and Services Tax Network, 2023a)

[6]1 USD equaled INR 74 in 2021.

authenticity of each invoice increased net tax collections by 10% (Shah, 2026). Our estimates, although smaller than more costly interventions, indicate promising potential of low-cost interventions built on existing digital architecture to address evasion at scale.

Our works shows the promise for a low-cost intervention to target fake credits in contrast to more costly resource-intensive audits (Carrillo et al., 2023; Waseem, 2023) and algorithm-based risk scoring (Mittal et al., 2018; Shah, 2026). These tools have not realized full potential in developing countries because of reasons ranging from institutional frictions hindering the integration of sophisticated algorithms in enforcement (Barwahwala et al., 2024) to limited state capacity in recovering known evasion (Okunogbe and Tourek, 2024; Best et al., 2021). In this study, we show that enforcement tools in the pre-filing stage can be effective in curbing evasion without relying on state's capacity to recover lost revenue post-filing. These findings can be useful in other contexts such as property taxes where delinquency remains prevalent despite being easily observable by tax authorities (Castro and Scartascini, 2015; Brockmeyer et al., 2021; Weigel, 2020).

Our findings also highlight a technology that replaces rather than augments the work of a tax official. Many interventions increase the productivity of state officials and aid in detecting tax liability rather than collecting taxes owed. High-tech tools have been successfully deployed to improve property tax collection in several contexts (Dzansi et al., 2022; Casaburi and Troiano, 2016; Meiselman, 2018). E-filing, e-invoicing and electronic billing machines have been introduced with varying degrees of success (Bellon et al., 2022; Okunogbe and Pouliquen, 2022). The policy we study speaks to the advantages of automation over discretion (Bachas et al., 2025). By showing that the policy was effective in targeting FCGs, we also argue that well-designed automated interventions may not hurt compliant taxpayers - a concern documented in the current literature (Almunia et al., 2024; Mascagni et al., 2023; Brockmeyer et al., 2024).

2 Institutional Context

2.1 Tax Filing and Payment in the GST

The Indian Goods and Services Tax (GST) is a nationwide value-added tax: firms remit taxes on their sales net of credits received for taxable inputs. Taxpayers report their total output and total eligible input credits on the "GSTR-3B" form, which we refer to as the *payment-linked* form because they cannot file it without making a payment. They also report invoice-level sales in the "GSTR-1" form, which we refer to as the *credit-linked* form since it generates credits for their clients. Credits available to buyers are limited to what a seller has reported on this form, which automatically populates a "GSTR-2" form for the client.

In theory, total output reported on the payment-linked form must either match or exceed

the sales reported on the credit-linked form, since the latter excludes sales to non tax-registered entities like final consumers. But in practice, there are discrepancies. Sellers can (fraudulently) provide credits without making associated payments by reporting higher sales in their credit-linked form than in their payment-linked form. We term such credits as "fake credit". Similar behavior in other countries might be to issue tax invoices without making associated payments.

2.2 Policy Change: Automated Blocking of Credit Provision

To curb the provision of fake credits, the GST administration implemented a nationwide reform in August 2021: firms were automatically blocked from filing their credit-linked return if they had any pending payment-linked return for the previous two months (i.e. affects payments starting in July 2021). Any firm affected by this policy could respond in three potential ways: (1) file all pending payment-linked returns, (2) stop filing any credit-linked returns, or (3) file payment-linked returns but report different sale amounts for payment and credit provision. Firms who expect to remain active and make payments eventually but faced temporary cash-flow issues or made errors would be expected to respond in the first way. Others, who would not have reason to remain registered if they cannot evade, could respond in either of the latter two ways.

3 Data

We use administrative data from the universe of Goods and Services Tax (GST) taxpayers registered in Delhi, obtained from the Delhi Department of Trade and Taxes. It includes firm-level information on registration status and two key tax forms returns described in Section 2. These returns allow us to construct our main measures of tax compliance: whether a taxpayer issued credits and whether they made a payment associated with those sales. We identify discrepancies between tax payments on sales to registered taxpayers and credits issued to those taxpayers, and leverage these discrepancies to classify firms as potential Fake Credit Generators (FCGs), as we describe in the next subsection.

The returns also allow us to build firm-to-firm trade linkages, enabling us to trace how the effects of the policy propagate through supply chains. We supplement these returns with firm characteristics, including sector (e.g., retail, manufacturing) and age, collected at the time of registration. For these client firms of potential FCGs, we examine the change in our key outcome variable of input credit used as a share of gross tax payments.

Our main sample consists of firms that file GST returns on a monthly basis and were registered for all return periods from January 2021 to April 2022, covering seven months before and eight months after the August 2021 introduction of the blocking policy. Although we analyze the outcomes of an expanded sample of firms that were registered for at least one period between January-June 2021 in Appendix Section B and discuss their contribution

to the aggregate revenue implications of the policy, we focus on continuously registered firms as our primary sample of interest. This restriction enables us to isolate the effect of the policy specifically on firms' filing behavior, and ensures that the results are not affected by tax authority-driven composition changes in registered taxpayers—if the tax authority were more likely to deregister firms that would have stopped filing their credit-linked form in response to the policy, we would overestimate the effect of the policy with the expanded sample. In our context, nearly all deregistrations are initiated by the tax authority and not the taxpayer. A taxpayer that intends to deregister would instead simply stop filing both returns.[7] Therefore, a key outcome of interest in our setting is whether a firm ceased return filing, which we can directly and reliably observe.

We also focus on taxpayers who file returns on a monthly frequency. This restriction leaves out taxpayers with annual taxable revenue less than Rs. 1.5 crore who may choose to file their returns quarterly. As the size restriction suggests, these taxpayers, although numerous, are less consequential to overall revenue. Moreover, it is harder to identify FCGs among quarterly filers because the discrepancies in their filings are more likely to reflect aggregation or reconciliation errors over a three-month horizon rather than deliberate evasion.[8]

3.1 Identifying Fake Credit Generators from Tax Returns

How large is the problem of fake credits, and can they be identified from returns data alone? Fake credits arise when firms report input credits that are not backed by corresponding tax payments. This can occur if no underlying transaction took place or if the transaction occurred between different parties than those declared.[9] We cannot detect the latter form of evasion solely through tax returns. We focus instead on cases where reported credits exceed reported payments—whether because the transaction was fictitious or the taxpayer simply neglected to report and make a payment associated with it, which leads to a loss of tax revenue. Under the GST, this occurs when a seller reports sales in the credit-linked form but fails to report or remit payment on those sales in the payment-linked form. Therefore, any positive discrepancy between sales reported in the credit-linked form that is not reported in the payment-linked form is evasion.

We identify firms likely to be generating fraudulent credits by comparing their credit provision and payment behavior before the policy change. For each firm, we calculate total tax payments in the payment-linked form and total credit issued in the credit-linked form between January and June 2021. Around 6.5% of firms report higher sales in the credit-linked form as

[7] If a firm does not file both the tax forms for 6 months, then it is likely to be deregistered from the portal by the department.

[8] In Appendix Section B, we separately analyze the effect of the policy on quarterly filers and include their small contribution in our final aggregate revenue estimates.

[9] Misreporting transacting partners can reduce tax liability because tax credits may be more valuable to some firms (eg. exporters) than to others. Such misreporting is facilitated by "invoice mill" firms as illustrated in Waseem (2023).

compared to the payment-linked form. Many of these discrepancies are small and could be attributed to mistakes or temporary liquidity issues. Therefore, we classify firms where the discrepancy between the two forms is higher than 10% as likely-fraudulent, yielding 5,644 firms who are potentially intentionally generating fake credits.[10] Additionally, there are 668 firms who have no tax payments between January - June 2021, but still issue tax credits to their buyers. These firms were responsible for 12-14 percent of excess credits issued before August 2021. We classify the subset of these firms with total credit provisions above the 75th percentile as likely generating fraudulent invoices.[11] We refer to firms that fall in either of these categories as Fake Credit Generators (FCGs). Our conclusions are robust to the specific choice of these thresholds.

Appendix Table A.1 compares the characteristics of firms we identify as FCGs to all other firms. FCGs issue more credits while reporting lower sales (captured by gross payments) relative to other firms. They are more likely to be in the services sector where sales are potentially harder to verify relative to manufacturing potentially due to factors such as absence of physical inventory, lower fixed capital, among others. FCGs are less likely to file their payment-linked returns. Importantly, firms that we characterize as FCGs are more likely to be affected by the policy change than other firms. Appendix Figure A.2 shows that FCGs were 2-3 times more likely to have a pending payment-linked form, even when they had filed the credit-linked form.

4 Empirical Strategy

We estimate the causal effect of the blocking policy on firm behavior using an event-study framework. Our main treatment group consists of firms classified as Fake Credit Generators (FCGs), following the procedure described in Section 3.1. All other firms serve as the control group.

To ensure comparability between treated and control firms, we implement Coarsened Exact Matching (CEM) on a set of pre-treatment firm characteristics: firm age, firm size, nature of business, whether the firm reports inter-state sales, and the number of payment returns filed prior to the policy (Blackwell et al., 2009). The matched control group is used in our preferred specification in Table 1. For transparency, we also report our key results using the full unmatched control group, and show that the conclusions remain robust.

[10] Panel A of Appendix Figure A.1 plots the distribution of excess credit provided in the credit-linked form that is not reported in the payment-linked form (scaled by the reported sales in the payment-linked form).

[11] Panel B of Appendix Figure A.1 plots the distribution of excess credit among firms with no tax payments.

Formally, we estimate the following two-way fixed effects event-study specification:

$$y_{it} = \gamma_i + \lambda_t + \sum_{\substack{k=-7 \\ k\neq -2}}^{8} \beta_k \times Treated_i \times \mathbb{1}(t = k) + \varepsilon_{it}, \tag{1}$$

where y_{it} denotes the outcome for firm i in month t, γ_i are firm fixed effects, and λ_t are filing-period fixed effects. The coefficients of interest, β_k, trace the dynamic differences between treated and control firms relative to June 2021 ($t = -2$), which serves as the omitted category.[12,13] Standard errors are clustered at the firm level throughout.

To identify the effect of the policy change on downstream firms, we first use the invoice-level data to trace input tax credits back to their source. Specifically, we identify all firms that received credits between January and June 2021 from suppliers classified as Fake Credit Generators (FCGs). We then define treated downstream firms as those who obtained a non-negligible share (at least 5 percent) of their total pre-period input tax credits from FCGs. Firms that did not receive any credits from FCGs form the control group, while upstream FCGs themselves are excluded from this group to avoid contamination.[14] We then estimate Equation 1, where treated downstream firms are compared to all control firms.

Identification relies on the assumption that, absent the blocking policy, outcomes for treated and untreated firms would have followed parallel trends, and that firms did not alter behavior in anticipation of the reform. The event-study estimates presented in Section 5 show stable pre-trends and no evidence of anticipatory responses, providing support for these assumptions. The main threat to identification is mean reversion: firms with large discrepancies between reported credits and payments in the pre-policy period may mechanically reduce excess credits over time, independent of the blocking policy. This could generate an apparent post-reform decline in discrepancies even without the policy. We provide evidence against this possibility in Section 5, showing that the effects are driven by the reform itself.

[12]Because the policy was implemented within two months of the filing deadline for July 2021, we omit $t = -2$ rather than $t = -1$ to avoid potential contamination from anticipatory responses

[13]When estimating the difference-in-differences version of this specification for filing outcomes, we truncate the sample at February 2022 to avoid overstating average treatment effects due to potential delayed filing in March and April return periods. The data for March and April was obtained from the tax administration in June 2022, so it is likely that some firms filed their March and April returns after the data was collected.

[14]We also exclude from the control group client firms of our expanded samples of FCGs (including deregistered firms and quarterly filers, described in Appendix Section B), so that none of the firms in the control group receive input credits from any FCGs.

5 Impact of Automated Checks on Excess Credits and Revenue

5.1 Impact on Tax Credits Generated

The first key finding is that FCGs stop providing input tax credits unbacked by tax payments in response to the policy. If the firms we classify as FCGs were making real sales but facing temporary liquidity constraints, they would have responded by filing the pending payment-linked form rather than ceasing all filing. Instead, they stop filing even their credit-linked returns, suggesting that previously issued credits were unlikely to be accompanied by payments.

Panel A of Table 1 shows that FCGs become 3 percentage points less likely to file their credit-linked form after the policy, corresponding to about a 50 percent increase in the non-filing rate (Column 1). The event-study estimates in Figure 2 confirm that the control group and FCGs exhibit parallel pre-policy trends in filing of the credit-linked form, while the drop in filing among FCGs is immediate and persistent once the policy is implemented.[15] In contrast, there is no increase—and even a slight decrease—in the likelihood of filing the payment-linked form among FCGs relative to control after the policy (Column 2).[16] FCGs are 1.5 percentage points more likely to cease filing all forms (Column 3), representing nearly a 100% increase in the probability of not filing either form. This cessation of all filing suggests that the average treated firm accepts that they will be eventually de-registered–they persistently do not file either form, as shown in Appendix Table A.2, and they do not resort to reporting zero sales in their payment-linked form to avoid deregistration, as shown in Appendix Table A.3.

Two further pieces of evidence suggest that these filing responses are a result of the policy and that they are likely to involve fraudulent credits.

First, we divide the FCGs into 2 categories. *Filer FCGs* consists of firms who have filed all their payment-linked form in the period from January-June 2021, but had substantial discrepancies between the credit and payment-linked forms. Even though these firms do not have a pending payment form, they are likely to respond to the policy if their perception regarding department's enforcement against fake credits increases. In contrast, firms that did not file at least one payment-linked form in this period are defined as *Non-Filer FCGs*.[17] We find that the effects of the policy are significantly more pronounced among non-filer FCGs, the group more directly affected by the policy because they had pending payment-linked return(s). Panel B of Table 1 shows non-filer FCGs are around 11 percentage points less likely to file their

[15] Appendix Figure A.3 plots the share of firms filing the credit-linked form in each month and shows that FCGs and control firms exhibit strikingly similar levels and pre-trends in filing behavior.

[16] Appendix Figure A.4 shows the corresponding event-study results for the effect of the policy on filing of the payment-linked form. The larger point estimates in the later return periods are likely driven by late-filing among firms we classify as FCGs, since the data was obtained from the tax administration in June 2022.

[17] It is notable that we identify 600 FCGs that were non-filers in the pre-period and this number is a comparable number to 800 ghost firms identified through substantial tax authority effort in Ecuador (Carrillo et al., 2023).

credit-linked return.[18] These firms are no more likely to start making payments (Column 2) and are 10 percentage points more likely to cease filing all forms (Column 3).[19] Taken together, these findings suggest that the policy was effective in reducing excess credits issued by FCGs whose prior behavior was now explicitly blocked by the policy. Non-filing FCGs accounted for about 4 percent of all excess credits in the economy in the first two quarters of 2021 though they account for only 0.1% of all firms in our data. Figure 1 shows that this share falls to almost zero after the policy change.

Second, we show that the magnitude of decline in filing of credit-linked form increases with the magnitude of pre-period discrepancies. Smaller pre-period discrepancies may be more likely to be caused by errors or temporary lapses while larger ones are more likely to be indicate fraudulent behavior. Appendix Figure A.6 shows that taxpayers whose pre-policy credits were at least 50 percent in excess of their payments show a 6 percentage point decrease in filing of the credit-linked form, as compared to a 3 percentage point decrease among firms with a discrepancy of at least 10 percent. Reassuringly, the effect of the policy on firms with positive but negligible pre-period discrepancies—likely caused by accounting errors—is close to 0, as shown in Appendix Table A.5.

These effects on the targeted firms still leave open the possibility that the firms generating excess credits in this manner had provided credits for real transactions but failed to make a payment on those transactions as opposed to generating fictitious transactions. For example, these credits may be coming from firms that were in the process of shutting down and unable to make their tax payments but wished to provide their suppliers with credits. In this case, their cessation of return filing would have occurred regardless of the blocking policy. We provide two pieces of evidence that the excess credits are unlikely to be backed by real economic activity. First, we examine the behavior of FCGs who meet the definition described in Section 3.1 based solely on activity in the first quarter of 2021. If our results are driven by firms that were going to shut down in due course or due to mean reversion induced by our approach of classifying firms as FCGs based on pre-period data, then the effects should be apparent in the second quarter, before the policy was implemented. Appendix Figure A.7 shows the results from classifying firms as FCGs using data from January-March 2021 and estimating Equation (1) using data from April 2021-April 2022: we observe that for all the FCGs (Panel A) and pre-period non-filing FCGs (Panel B), the drop in filing of credit-linked form is significant only after the policy change. The second piece of evidence comes from the fact that clients of these FCGs do not replace the lost credits. If the missing transactions were real transfers of goods and services, we would expect to find that clients either replace these lost input purchases with purchases from other firms or decrease their sales if they are unable to replace suppliers. We investigate this possibility by examining the behavior of the clients of FCGs in the next

[18] Appendix Figure A.5 presents the results from event-study, estimated using Equation 1.

[19] Appendix Table A.4 shows that the policy-induced effects on filing FCGs are comparable to the effects among all FCGs, as they form the majority of this sample.

section.

We also provide descriptive evidence on the impact on firm deregistration. Month-wise deregistration data is only available to us after the implementation of the policy. Panel A of Appendix Figure A.8 shows that deregistration rates increased by around 3 percentage points in the post period. Once we restrict the FCG sample to non-filers in Panel B, deregistration rates increase by more than 30 percentage points compared to the control group.

Our main sample consists of firms that were registered throughout the study period to avoid conflating our results with endogenous tax-authority driven deregistrations. As a sensitivity check and to fully account for the effect of the policy on total revenue, we include firms that were registered in least one month between January and June 2021. Filing outcomes can be accurately measured for all firms regardless of registration status since it does not affect whether we observe a return filed. Appendix Table B.1 shows that among this expanded sample of taxpayers, FCGs became 6 percentage points less likely to provide any credits, 1.7 percentage points less likely to make any payment and 5.4 percentage points more likely to stop filing all returns.

5.2 Impact on Credits and Payments by Beneficiaries of Fake Credits

Do clients of FCGs increase their net tax payments when the likely-fraudulent credit they received from FCGs declines, or do they find another source of credit such that their tax liability remains unchanged?

We begin by characterizing the downstream firms who receive a non-negligible share (at least 5%) of their input tax credits from firms we classify as FCGs and compare them to a control group of firms who receive none of their credit from FCGs; FCGs, themselves, are excluded from this analysis. Appendix Table A.6 shows that network firms tend to be larger than control firms (in terms of gross payments or sales), and like FCGs themselves, they are more likely to be in the services sector. Next, we estimate Equation 1, defining FCGs' clients as the treatment group. Figure 3 shows the impact of the policy on net tax payments and credits used by FCGs' clients. Compared to the control group, network firms permanently decrease the share of their output tax obligations paid using credits. The shift in the mode of payment results in net tax payments increasing by 1.4 percentage points, implying a 5 percent increase in revenue (see Columns 1 and 2 of Table 2). Columns 3 of Table 2 shows that the estimated effect on gross payments is noisy; we cannot rule out that the policy has no effect gross payments and the point estimate of 1.7 percent is not economically large when compared to the decline in credit payments.[20]

The results become starker once we limit our attention to the network of pre-period non-filing FCGs (Columns 4-6, Table 2). The network firms reduce their share of credit used by 1.7

[20] Appendix Table A.7 shows that less than 0.4% additional firms start reporting 0 sales after the policy change. We also don't observe that network firms are more likely to be de-registered following the policy. Appendix Figure A.9 shows that, if anything, network firms are less likely to deregister than control firms.

percentage points, representing a 9 percent increase in net tax payments. The decline in gross tax payments continues to be small and insignificant.[21, 22]

Overall, these findings indicate that the purchases from FCGs were likely to be fraudulent as the clients of these firms do not change their sales and increase net tax payments, as they reduce their input procurement from the FCGs. If the inputs purchases were real, then the clients should have either reduced their output or found another input supplier once the FCGs were impacted by the policy change; our results show that the network firms do not replace the lost credits even 8 months after the policy change.

5.3 Revenue Implications

We calculate the implied revenue increase in New Delhi, as a result of the policy. We do this first by calculating the implied total decrease in fraudulent credits, which translates to a net revenue gain for the tax authority. Our results showed that the policy reduced the likelihood that the average FCG generates any credits in a given month by 2.9 percentage points (Column 1, Table 1). The average monthly credit generated by all FCGs combined between Jan-June 2021 is approximately INR 4.5 billion. Multiplying these two figures and then multiplying by 12 translates into an annual decline of fraudulent credit of INR 1.6 billion. The 95% confidence intervals allow us to rule out effects smaller than INR 1.3 billion and larger than INR 1.9 billion. To interpret the magnitude, it is helpful to compare the expected annual decline of fraudulent credit to the tax authority's estimate of the total fraudulent credit in the economy, which in FY 2023-24 was estimated to be INR 121 billion annually in New Delhi.[23] Therefore, the blocking policy reduced total evasion by approximately 1.5%. If we include the reduction in credit generated by firms that are deregistered during our sample period, the revenue effect is much larger – approximately INR 5.4 billion (95% CI: 4.7 billion, 6 billion) or 4.5 percent of estimated fraudulent credit.[24,25]

An alternative approach to evaluating the aggregate revenue implications of the policy would be to estimate the increase in net tax payments (gross payments net of ITC) by client firms after accounting for any changes in their gross payments. We report in Column 2 of

[21]For completeness, we include the results from the event study specification for the effect of the policy on the credit used and net tax payments among the buyers of non-filing FCGs in Appendix Figure A.10.

[22]Appendix B describes results for alternative samples of taxpayers. The results remain similar when we include the clients of firms who were registered for at least 1 period between January-June 2021 in Appendix Table B.2. Also, there is no impact on the net tax payments of the clients of FCGs who file their returns every quarter, instead of every month, since these firms are smaller and are less significant as sources of fraudulent credit for downstream firms.

[23]The suspected amount of evasion in the last quarter of 2023 was INR 30.3 billion Lok Sabha Questions (2024). We multiply this amount by 4 to arrive at the annual estimate.

[24]Monthly pre-policy credits generated by FCGs that were registered in at least one period between January-June 2021 were INR 7.2 billion. We multiply this number with the estimated decline in filing of credit-linked form reported in Appendix Table B.1 and annualize it to get an estimate of the revenue effect.

[25]Repeating this procedure using our results for the sample of firms who elect to file their returns at a quarterly frequency (Table B.3) increases our revenue estimates by only INR 0.1 billion.

Table 2 that client firms reduce the proportion of credits used to make tax payments by 1.4 p.p. because of the policy. The monthly gross payments in the pre-period were INR 20.6 billion which reduce by 1.7% in the post period (see Table 2 Column 3) leading to post-policy gross payments of INR 20.24 billion. After considering the reduction in ITC usage, we estimate an annual increase in revenue of INR 3.4 Billion or 0.8 percent of net GST collections in Delhi.[26] The 95% confidence intervals allow us to rule out effects smaller than INR 2.7 billion and larger than INR 3.9 billion.

When we include client firms of deregistered FCGs, the aggregate effect is INR 4 billion (95% CI: 3.5 Billion, 4.7 Billion).[27] The decline in credits should correspond directly to the increase in net payments since the client firms must replace credits with cash payments. Reassuringly, we are unable to reject the null hypothesis that the implied decrease in fraudulent credit is different from the implied increase in revenue for the tax authority, suggesting that the estimated revenue increase is likely due to the effect of the policy on fraudulent credit in the economy.

6 Conclusion

The increasing digitization of tax administration in developing countries provides opportunities to raise the efficiency of tax collection in new, low-cost ways. We find untapped potential in automated cross-checks using the digital architecture of VAT return filing.

We find that simply requiring firms to make a payment to enable credit provision through the online portal reduces the provision of credit without payment by 1.5 to 4.5 percent. Firms that were providing these credits stop all activity and are likely then deregistered. Their clients, who were previously benefiting from these credits, permanently increase net tax payments by 5 percent without decreasing sales. Reassuringly, the balancing concern that affected taxpayers are making minor mistakes or only temporarily non-compliant does not hold.

[26] Gross collections and refunds in 2021-22 were INR 462.5 billion and INR 27 billion, respectively (Goods and Services Tax Network, 2021).

[27] The pre-period monthly gross payments for these firms is INR 25 Billion; we scale this number using the effects in Table B.2 using the same procedure as for the network of continuously registered FCGs to obtain the aggregate increase in net tax payments.

References

Almunia, Miguel, Jonas Hjort, Justine Knebelmann, and Lin Tian, "Strategic or confused firms? evidence from "missing" transactions in Uganda," *Review of Economics and Statistics*, 2024, *106* (1), 256–265.

Bachas, Pierre, Anne Brockmeyer, Alipio Ferreira, and Bassirou Sarr, "Algorithms and bureaucrats: evidence from tax audit selection in Senegal," Technical Report, The World Bank 2025.

Barwahwala, Taha, Aprajit Mahajan, Shekhar Mittal, and Ofir Reich, "Is Model Accuracy Enough? A Field Evaluation Of A Machine Learning Model To Catch Bogus Firms," Technical Report, National Bureau of Economic Research 2024.

Bellon, Matthieu, Era Dabla-Norris, Salma Khalid, and Frederico Lima, "Digitalization to improve tax compliance: evidence from VAT e-Invoicing in Peru," *Journal of Public Economics*, 2022, *210*, 104661.

Best, Michael, Jawad Shah, and Mazhar Waseem, "Detection Without Deterrence: Long-Run Effects of Tax Audit on Firm Behavior," 2021.

Blackwell, Matthew, Stefano Iacus, Gary King, and Giuseppe Porro, "cem: Coarsened exact matching in Stata," *The Stata Journal*, 2009, *9* (4), 524–546.

Brockmeyer, Anne, Alejandro Estefan, Karina Ramírez Arras, and Juan Carlos Suárez Serrato, "Taxing property in developing countries: Theory and evidence from mexico," Technical Report, National Bureau of Economic Research 2021.

_ , Giulia Mascagni, Vedanth Nair, Mazhar Waseem, and Miguel Almunia, "Does the Value-Added Tax Add Value? Lessons Using Administrative Data from a Diverse Set of Countries," *Journal of Economic Perspectives*, 2024, *38* (1), 107–132.

Carrillo, Paul, Dave Donaldson, Dina Pomeranz, and Monica Singhal, "Ghosting the Tax Authority: Fake Firms and Tax Fraud in Ecuador," 2022.

_ , _ , _ , and _ , "Ghosting the tax authority: fake firms and tax fraud in Ecuador," *American Economic Review: Insights*, 2023, *5* (4), 427–444.

_ , Dina Pomeranz, and Monica Singhal, "Dodging the taxman: Firm misreporting and limits to tax enforcement," *American Economic Journal: Applied Economics*, 2017, *9* (2), 144–164.

Casaburi, Lorenzo and Ugo Troiano, "Ghost-house busters: The electoral response to a large anti–tax evasion program," *The Quarterly Journal of Economics*, 2016, *131* (1), 273–314.

Castro, Lucio and Carlos Scartascini, "Tax compliance and enforcement in the pampas evidence from a field experiment," *Journal of Economic Behavior & Organization*, 2015, *116*, 65–82.

Dzansi, James, Anders Jensen, David Lagakos, and Henry Telli, "Technology and tax capacity: evidence from local governments in Ghana," Technical Report, National Bureau of Economic Research 2022.

Goods and Services Tax Network, "Table: State wise Tax Collection," https://www.gst.gov.in/download/gststatistics 2021.

_ , "Table: Return (GSTR-3B)," https://www.gst.gov.in/download/gststatistics 2023.

_ , "Table: State wise Tax Collection," https://www.gst.gov.in/download/gststatistics 2023.

Jensen, Anders, "Employment structure and the rise of the modern tax system," *American Economic Review*, 2022, *112* (1), 213–234.

Kleven, Henrik Jacobsen, Claus Thustrup Kreiner, and Emmanuel Saez, "Why can modern governments tax so much? An agency model of firms as fiscal intermediaries," *Economica*, 2016, *83* (330), 219–246.

_ , Martin B Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark," *Econometrica*, 2011, *79* (3), 651–692.

Kumler, Todd, Eric Verhoogen, and Judith Frías, "Enlisting employees in improving payroll tax compliance: Evidence from mexico," *Review of Economics and Statistics*, 2020, *102* (5), 881–896.

Lok Sabha Questions, "Unstarred Question No. 435: Unearthing of Fake Input Tax Credit," Answered by the Minister of State in the Ministry of Finance, https://sansad.in/getFile/loksabhaquestions/annex/1715/AU435.pdf?source=pqals 2024.

Mascagni, Giulia, Roel Dom, Fabrizio Santoro, and Denis Mukama, "The VAT in practice: equity, enforcement, and complexity," *International Tax and Public Finance*, 2023, *30* (2), 525–563.

Meiselman, Ben S, "Ghostbusting in Detroit: Evidence on nonfilers from a controlled field experiment," *Journal of Public Economics*, 2018, *158*, 180–193.

Mittal, Shekhar, Ofir Reich, and Aprajit Mahajan, "Who is bogus? using one-sided labels to identify fraudulent firms from tax returns," in "Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies" 2018, pp. 1–11.

Okunogbe, Oyebola and Gabriel Tourek, "How can lower-income countries collect more taxes? The role of technology, tax agents, and politics," *Journal of Economic Perspectives*, 2024, *38* (1), 81–106.

_ **and Victor Pouliquen**, "Technology, taxation, and corruption: evidence from the introduction of electronic tax filing," *American Economic Journal: Economic Policy*, 2022, *14* (1), 341–72.

Pomeranz, Dina, "No taxation without information: Deterrence and self-enforcement in the value added tax," *American Economic Review*, 2015, *105* (8), 2539–2569.

Shah, Jawad, "Using computerized information to enforce VAT: Evidence from Pakistan," *Journal of Development Economics*, 2026, *178*, 103595.

Waseem, Mazhar, "Overclaimed refunds, undeclared sales, and invoice mills: Nature and extent of noncompliance in a value-added tax," *Journal of Public Economics*, 2023, *218*, 104783.

Weigel, Jonathan L, "The participation dividend of taxation: How citizens in Congo engage more with the state when it tries to tax them," *The Quarterly Journal of Economics*, 2020, *135* (4), 1849–1903.

Table 1: Effect on Filing Behavior of Fake Credit Generators

	(1)	(2)	(3)
	Filed Credit-Linked Form	Filed Payment-Linked Form	Filed Neither Form
Panel A: Effect of Policy on All Fake Credit Generators			
Fake Credit Generator × Post	-0.029	-0.007	0.015
	(0.003)	(0.002)	(0.002)
Outcome Mean	0.945	0.964	0.016
Observations	2,858,030	2,858,030	2,858,030
Panel B: Effect of Policy on Non-Filer Fake Credit Generators			
Fake Credit Generator × Post	-0.109	-0.002	0.097
	(0.017)	(0.014)	(0.017)
Outcome Mean	0.811	0.627	0.164
Observations	57,372	57,372	57,372

Notes: This table presents the causal effect of the policy on filing outcomes of Fake Credit Generators. Panel A presents the effects of the policy on all firms we classify as Fake Credit Generators, relative to a control group of all other firms. Panel B presents the effects of the policy on Fake Credit Generators that don't file at least one payment-linked return between January-June 2021, relative to a control group of all other firms that don't file at least one payment-linked return between January-June 2021. Column 1 presents the effect on filing of the credit-linked form in a given month. Column 2 presents the effect on filing of the payment-linked form in a given month. Column 3 presents the effect on filing neither form in a given month (where the outcome is equal to 1 if a firm files neither form, and 0 otherwise). In each case, the causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. The outcome mean is the average of the outcome for the treatment group before the policy was implemented. Standard errors are clustered at the firm level.

Table 2: Effect on Network Firms

	(1)	(2)	(3)	(4)	(5)	(6)
	All Network Firms			Non-Filer Network Firms		
	Net Tax Payments (Share)	Credit Used (Share)	Log Gross Payments	Net Tax Payments (Share)	Credit Used (Share)	Log Gross Payments
Network Firm × Post	0.014	-0.014	-0.017	0.017	-0.017	-0.022
	(0.002)	(0.002)	(0.012)	(0.005)	(0.005)	(0.028)
Outcome Mean	0.259	0.741	11.862	0.181	0.819	11.799
Observations	1,957,302	1,957,302	1,957,302	1,875,425	1,875,425	1,875,425

Notes: This table presents the causal effect of the policy on the network (buyers) of firms we classify as Fake Credit Generators. We define the network of Fake Credit Generators as those firms who receive at least 5% of their total input tax credit between January-June 2021 from Fake Credit Generators. We compare these network firms to all firms who receive none of their input tax credits between January-June 2021 from Fake Credit Generators. We exclude all Fake Credit Generators themselves from this analysis. Columns 1-3 focus on the network of all firms we classify as Fake Credit Generators. Columns 4-6 focus on the network of the firms we classify as Fake Credit Generators that additionally do not file at least one payment return between January-June 2021. The outcome in Columns 1 and 4 is defined as the ratio of net tax payments made by a firm using cash divided by the firm's total tax payments. The outcome in Columns 2 and 5 is defined as the ratio of tax payments made by a firm using input tax credit divided by the firm's total tax payments. The outcome in Columns 3 and 6 is defined as the log of the firm's gross tax payments in a given time period. We estimate the causal effects using the two-way fixed effects design described in Section 4. The outcome mean is the average of the outcome for the network firms before the policy was implemented. Standard errors are clustered at the firm-level.

Figure 1: Trends in Excess Credit Attributable to Non-Filer Fake Credit Generators

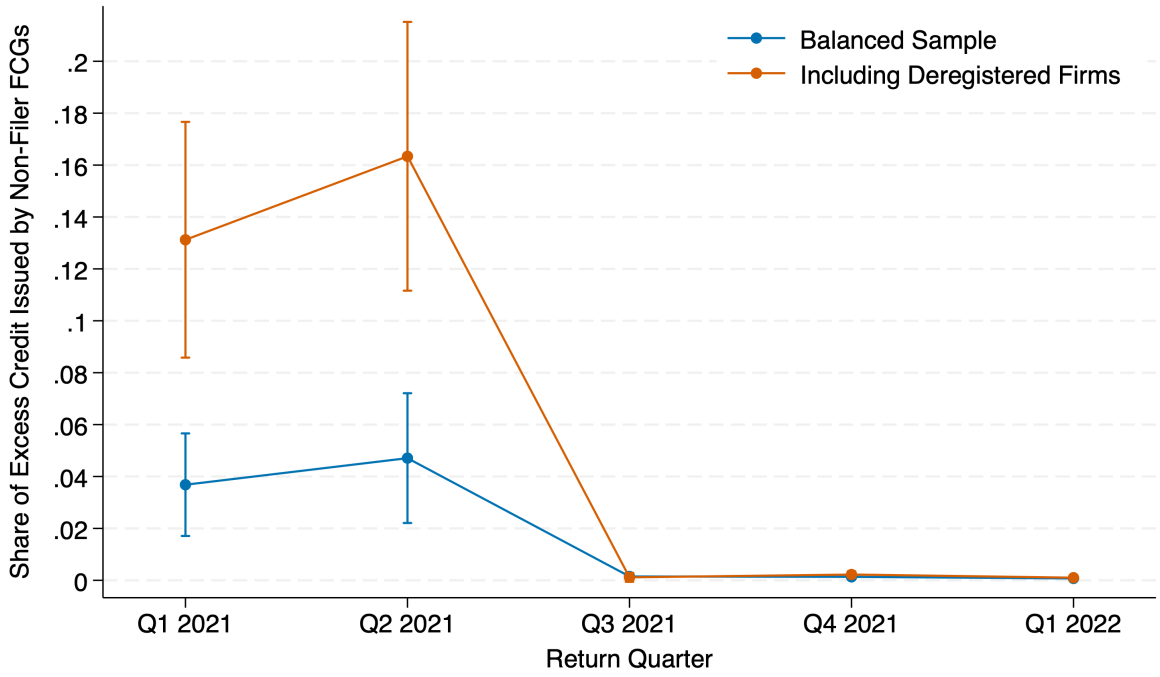

Notes. This figure plots the proportion of Excess Credit in a given quarter that is attributable to Non-Filer Fake Credit Generators (FCGs). Excess Credit for a firm is defined as the total credit issued by the firm to their suppliers in that quarter subtracted by the total tax payments by the firm in that quarter. Non-Filer FCGs are the firms that we classify as Fake Credit Generators who additionally don't file at least one payment-linked form between January-June 2021. In this figure, we show the ratio of the sum of Excess Credit attributable to Non-Filer Fake Credit Generators in a quarter divided by the total amount of Excess Credit in the economy in that quarter, separately for non-filer FCGs that are registered in every period between January 2021 to April 2022 ("Balanced Sample") and for non-filer FCGs that are required to be registered in every period between January 2021 to September 2021.

Figure 2: Effect of Policy on Filing of Credit-Linked Return

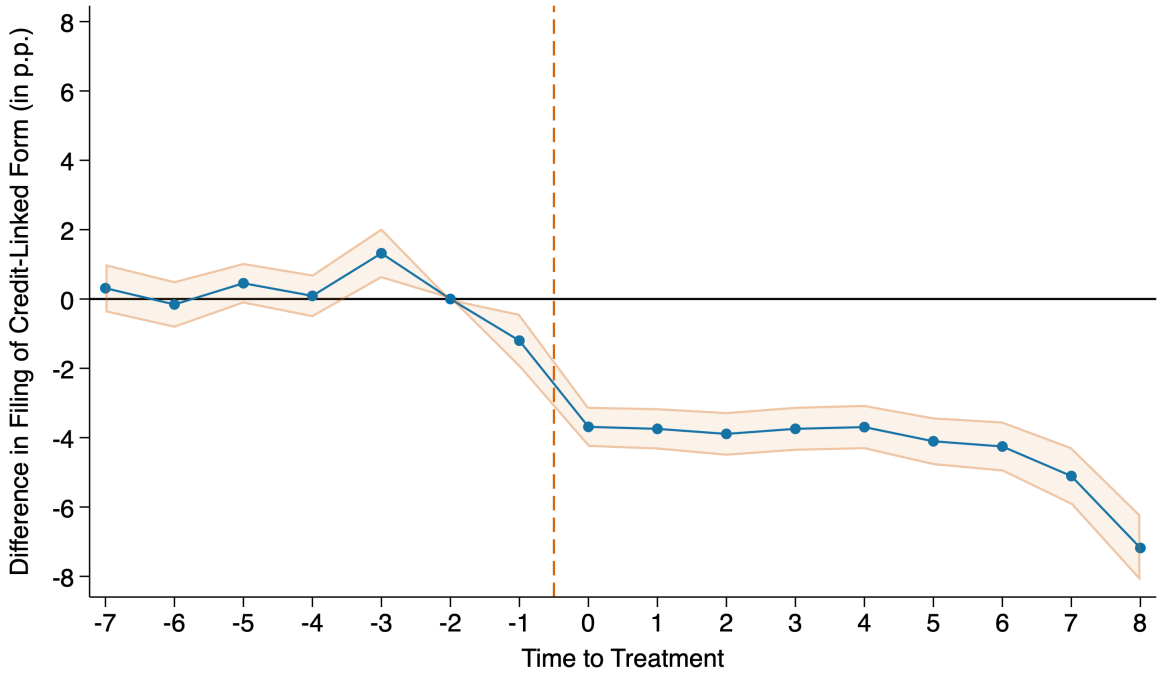

Notes. This figure shows the causal effect of the policy on filing of the credit-linked form by Fake Credit Generators. It plots the event-study coefficients from Equation 1 described in Section 4, along with the 95% confidence intervals. The treatment group includes all firms we classify as Fake Credit Generators. The control group includes all firms we do not classify as Fake Credit Generators. Standard errors are clustered at the firm-level.

Figure 3: Effect of Policy on Tax Payments by Network Firms

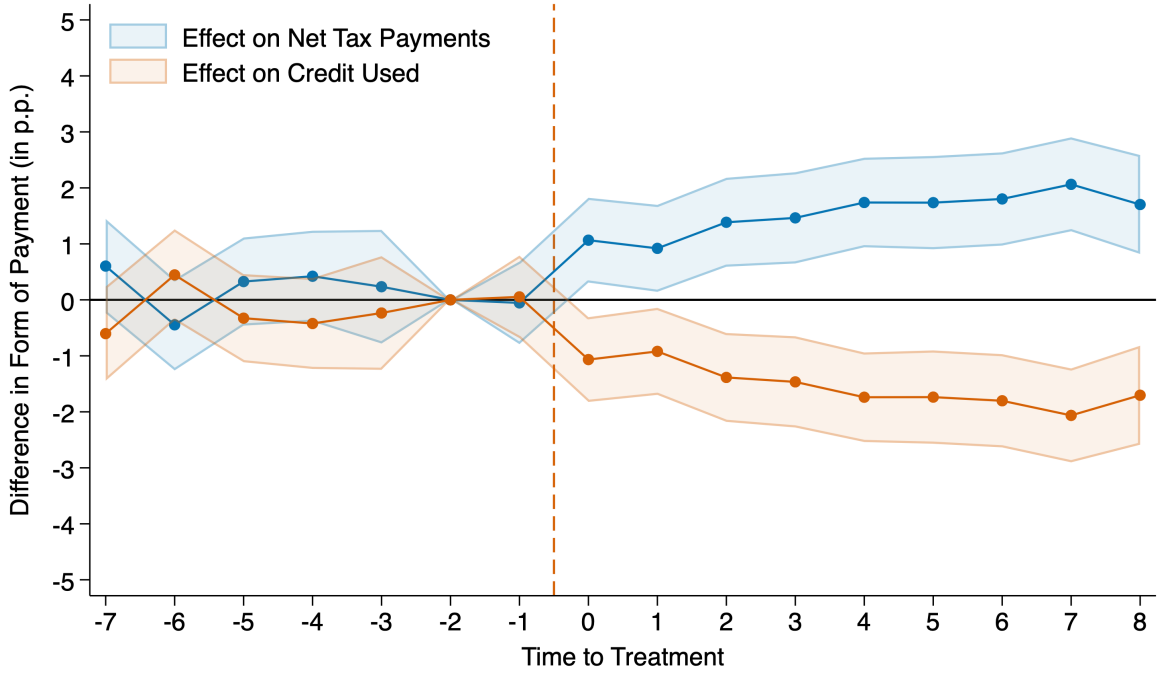

Notes. This figure shows the causal effect of the policy on the tax payment method used by the network of firms we classify as Fake Credit Generators. We define the network of Fake Credit Generators as those firms who receive at least 5% of their total input tax credits between January-June 2021 from Fake Credit Generators. We compare these network firms to all firms who receive none of their input tax credits between January-June 2021 from Fake Credit Generators. We exclude all Fake Credit Generators themselves from this analysis. The outcome is defined as the ratio of tax payments made using cash or input tax credits to the total tax payments in that return period. For a given firm, this ratio must necessarily sum to one in a given return period. We plot the event-study coefficients from Equation 1 described in Section 4 along with the 95% confidence intervals, defining the network of FCGs as the treated firms. Standard errors are clustered at the firm-level.

Appendix

A Appendix Figures and Tables

Table A.1: Summary Statistics

	(1)	(2)	(3)	(4)	(5)
	Control Firms	Fake Credit Generators	Non-Filer Fake Credit Generators	Difference: (2)-(1)	P-Value: (2)-(1)
Number of Credit-Linked Returns Filed	5.718	5.713	5.008	-0.006	0.598
Number of Payment-Linked Returns Filed	5.913	5.765	3.720	-0.148	0.000
Filed All Pre-Period Payment-Linked Returns	0.974	0.897	0.000	-0.077	0.000
Age (in months)	30.629	31.335	26.247	0.705	0.000
Has Inter-State Sales	0.653	0.692	0.623	0.039	0.000
Credit Issued	3,027,093	4,671,994	2,521,705	1,644,901	0.062
Gross Payments	4,115,650	3,261,315	1,514,462	-854,335	0.253
Nature of Business					
Retail	0.328	0.217	0.273	-0.110	0.000
Services	0.176	0.316	0.193	0.139	0.000
Wholesale/Manufacturing	0.240	0.182	0.283	-0.058	0.000
Other	0.256	0.285	0.250	0.029	0.000
Number of Firms	200,868	5,811	600		

Notes: This table reports the summary statistics for all firms in our baseline sample, using data before the policy was implemented between January-June 2021. Column 1 reports the mean value of the corresponding variable for the firms we classify to be in the control group, i.e. firms that are not fake credit generators. Column 2 reports the mean value of the corresponding variable for the firms we classify to be Fake Credit Generators. Column 3 reports the mean value of the corresponding variable for the firms we classify to be Fake Credit Generators that additionally do not file at least one payment-linked return between January-June 2021. Column 4 reports the difference between Column 1 and Column 2. Column 5 provides the p-value from a t-test of the difference between Column 1 and Column 2. "Filed All Pre-Period Payment-Linked Returns " is equal to 1 if a firm filed all payment-linked returns between January-June 2021. "Age" refers to the number of months the firm had been registered for at the start of our sample period (Jan 2021). Because the GST system was introduced in July 2017, the largest value the age variable can take is 42. "Has Inter-State Sales" takes a value of 1 if a firm has any buyer that is not registered in New Delhi, and 0 otherwise. "Credit Issued" is the total rupee value of input tax credits supplied by the firm to their buyers between Jan-June 2021. "Gross Payments" is the total rupee value of tax payments made by the firm between Jan-June 2021.

Table A.2: Effect on Filing of Neither Form for Past 3 Months

	(1)
	Filed Neither Form For 3 Months
Fake Credit Generator × Post	0.011
	(0.003)
Outcome Mean	0.006
Observations	2,449,740

Notes: This table presents the causal effect of the policy on filing neither form in for the past 3 months (where the outcome is equal to 1 at time t if a firm files neither form at t,$t-1$, and $t-2$, and 0 otherwise). To obtain clean pre- and post-policy periods, we exclude the return periods August and September 2021 from this analysis. The causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Table A.3: Effect on Filing of NIL Payment-Linked Form

	(1)
	Filed NIL Payment-Linked Form
Fake Credit Generator × Post	-0.004
	(0.003)
Outcome Mean	0.249
Observations	2,858,030

Notes: This table presents the causal effect of the policy on filing of a NIL payment-linked return by Fake Credit Generators. Specifically, the outcome is equal to 1 in a given return period if a firm filed the payment-linked return and reported no sales, and is 0 otherwise. The causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Table A.4: Effect on Filing Behavior of Filer Fake Credit Generators

	(1)	(2)	(3)
	Filed Credit-Linked Form	Filed Payment-Linked Form	Filed Neither Form
Fake Credit Generator × Post	-0.020	-0.007	0.006
	(0.002)	(0.002)	(0.001)
Outcome Mean	0.959	0.999	0.000
Observations	2,806,552	2,806,552	2,806,552

Notes: This table presents the causal effect of the policy on filing outcomes of Fake Credit Generators. We evaluate the effects of the policy on fake credit generators that file all payment-linked returns between January-June 2021, relative to a control group of all other firms that file all payment-linked returns between January-June 2021. Column 1 presents the effect on filing of the credit-linked form in a given month. Column 2 presents the effect on filing of the payment-linked form in a given month. Column 3 presents the effect on filing of the neither form in a given month. In each case, the causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Table A.5: Effect on Filing of Credit-Linked Form in Placebo Treatment Group

	(1)
	Filed Credit-Linked Form
Placebo FCG × Post	-0.004
	(0.001)
Outcome Mean	0.975
Observations	2,493,134

Notes: This table presents the causal effect of the policy on filing of a credit-linked return by firms with a positive but negligible discrepancy between credit issued and gross payments. Specifically, the treatment status is set to 0 for all firms who did not issue credit in excess of gross payments, and it is set to 1 for all firms who issued between 0-10% more credit than their gross payments. All firms with a discrepancy larger than 10% are dropped from this analysis. The outcome is equal to 1 in a given return period if a firm filed the credit-linked return and is 0 otherwise. The causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Table A.6: Summary Statistics - Network Firms

	(1)	(2)	(3)	(4)	(5)
	Control Firms	Network Firms	Non-Filer Network Firms	Difference: (2)-(1)	P-Value: (2)-(1)
Number of Credit-Linked Returns Filed	5.707	5.785	5.774	0.078	0.000
Number of Payment-Linked Returns Filed	5.902	5.980	5.975	0.078	0.000
Age (in months)	29.741	34.715	34.813	4.974	0.000
Has Inter-State Sales	0.617	0.825	0.876	0.207	0.000
Credit Issued	1,650,879	11,064,002	4,249,417	9,413,123	0.000
Gross Payments	2,108,201	17,727,650	5,259,054	15,619,450	0.000
Ratio of Net Tax Payments to Gross Payments	0.315	0.259	0.194	-0.056	0.000
Nature of Business					
Retail	0.328	0.308	0.371	-0.020	0.000
Services	0.183	0.192	0.099	0.009	0.065
Wholesale/Manufacturing	0.232	0.326	0.238	0.006	0.235
Other	0.257	0.262	0.203	0.005	0.363
Number of Firms	176,291	6,973	929		

Notes: This table reports the summary statistics for the network of fake credit generators, using data before the treatment was implemented between January-June 2021. Column 1 reports the mean value of the corresponding variable for the firms that don't receive any input tax credit from fake credit generators between January-June 2021. Column 2 reports the mean value of the corresponding variable for the firms that receive at least 5% of their input tax credit from Fake Credit Generators. Column 3 reports the mean value of the corresponding variable for the firms that receive at least 5% of their input tax credit from Fake Credit Generators that additionally do not file at least one payment-linked return between January-June 2021. Column 4 reports the difference between Column 1 and Column 2. Column 5 provides the p-value from a t-test of the difference between Column 1 and Column 2. "Filed All Pre-Period Payment-Linked Returns " is equal to 1 if a firm filed all payment-linked returns between January-June 2021. "Age" refers to the number of months the firm had been registered for at the start of our sample period (Jan 2021). Because the GST system was introduced in July 2017, the largest value the age variable can take is 42. "Has Inter-State Sales" takes a value of 1 if a firm has any buyer that is not registered in New Delhi, and 0 otherwise. "Credit Issued" is the total rupee value of input tax credits supplied by the firm to their buyers between Jan-June 2021. "Gross Payments" is the total rupee value of tax payments made by the firm between Jan-June 2021.

Table A.7: Effect on Any Sales of Network Firms

	(1)
	Any Sales
Network Firm × Post	-0.004
	(0.002)
Outcome Mean	0.859
Observations	2,837,792

Notes: This table presents the causal effect of the policy on the network (buyers) of firms we classify as Fake Credit Generators. We define the network of Fake Credit Generators as those firms who receive at least 5% of their total input tax credit between January-June 2021 from Fake Credit Generators. We compare these network firms to all firms who receive none of their input tax credits between January-June 2021 from Fake Credit Generators. We exclude all Fake Credit Generators themselves from this analysis. The outcome is equal to 1 if the network firm reported non-zero sales in that return period, and 0 otherwise. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Figure A.1: Distribution of Excess Credit, January-June 2021

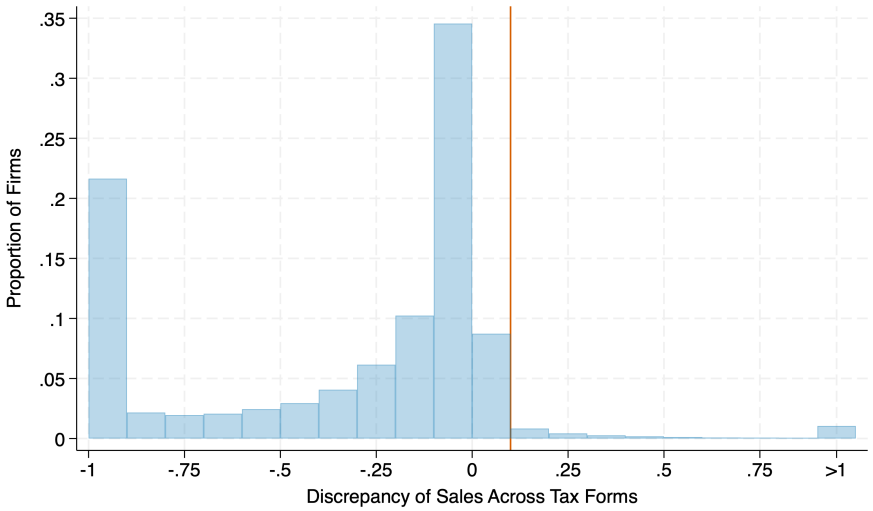

(a) Excess Credit/Gross Payments

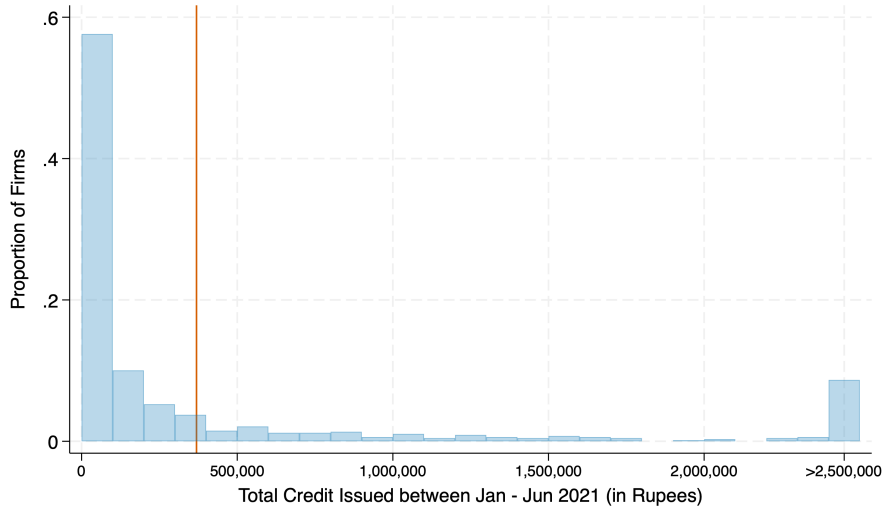

(b) Excess Credit Conditional on No Reported Sales

Notes. This figure plots the distribution of Total Excess Credit issued by firms between January-June 2021. Total Excess Credit for a firm is defined as the total credit issued by the firm to their suppliers between January-June 2021 subtracted by the total gross payments by the firm in this time period. In Panel A, we scale Total Excess Credit by the Total Gross Payments (defined as the sum of all tax payments between January-June 2021) and plot its distribution for firms with positive sales between January-June 2021. All firms that fall to the right of the vertical line (corresponding to a value of 0.1) are classified as Fake Credit Generators. In Panel B, we plot the distribution of Total Excess Credit among firms with no reported sales between January-June 2021. All firms to the right of the vertical line (corresponding to the 75th percentile of this distribution) are classified as Fake Credit Generators

Figure A.2: Exposure to Policy Change

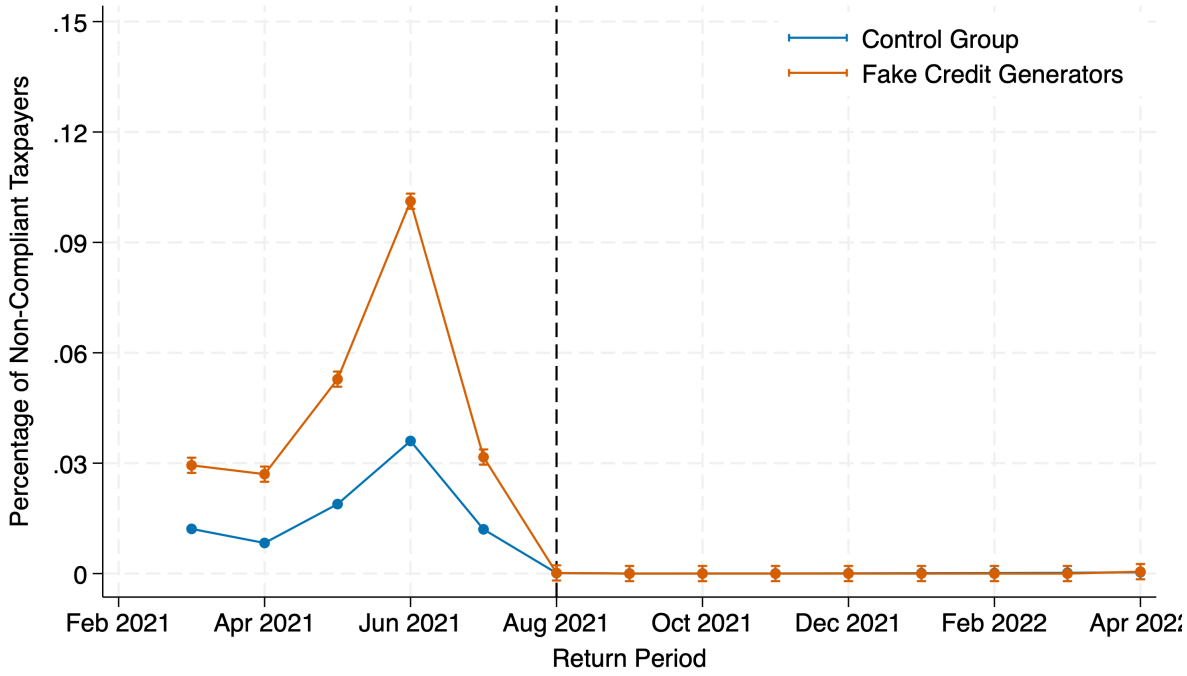

Notes. This figure plots the proportion of taxpayers that are non-compliant with the policy described in Section 2. Specifically, this figure shows the proportion of taxpayers who satisfy the following two conditions: (a) file a credit-linked form in a given return period, and (b) had not filed at least one payment-linked from two or more return periods ago.

Figure A.3: Trends in Filing of Credit-Linked Return

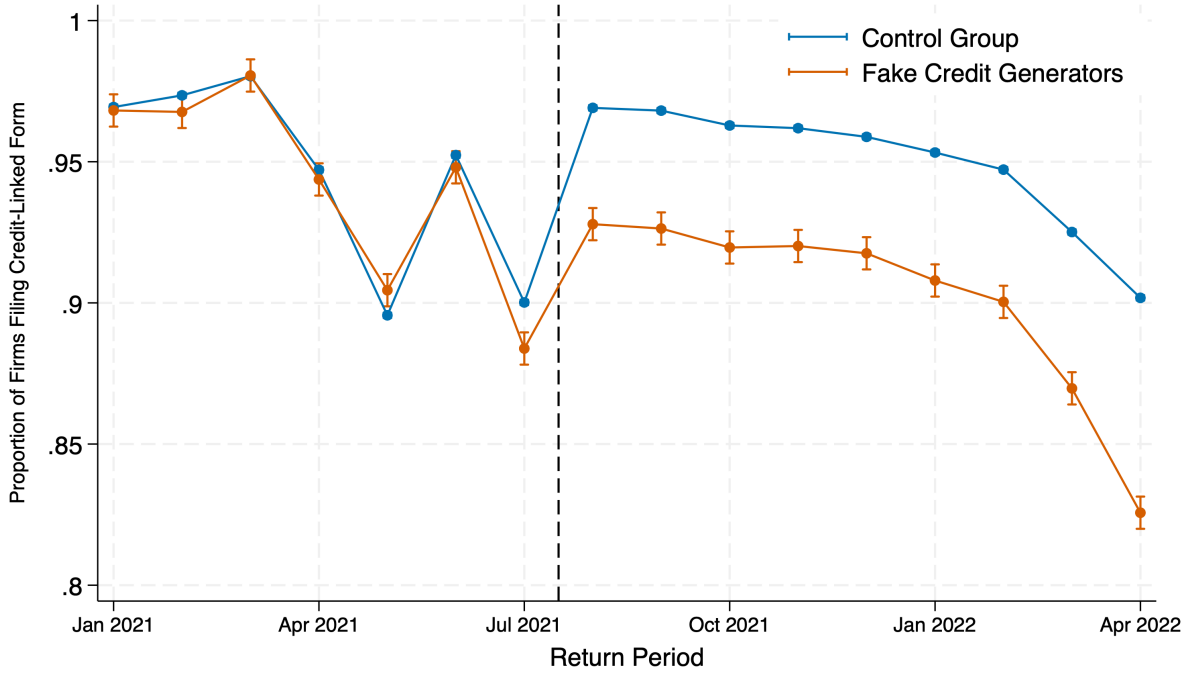

Notes. This figure shows trends in filing of the credit-linked return by firms we classify as Fake Credit Generators, compared to the trends in filing of the credit-linked return by all firms we do not classify as Fake Credit Generators (the control group).

Figure A.4: Effect of Policy on Filing of Payment-Linked Form

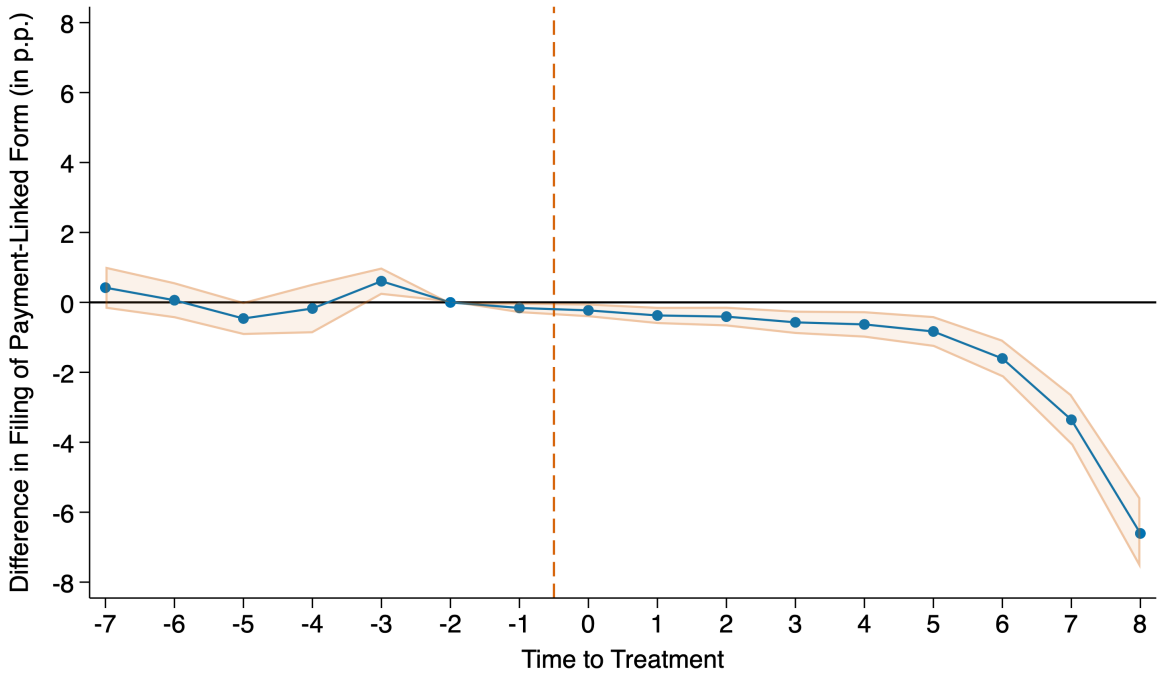

Notes. This figure shows the causal effect of the policy on filing of the payment-linked form by Fake Credit Generators. We plot the event-study coefficients from Equation 1 described in Section 4, along with the 95% confidence intervals. The treatment group consists of all firms we classify as Fake Credit Generators. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). Standard errors are clustered at the firm-level.

Figure A.5: Effect on Filing of Credit-Linked Return by Non-Filer Fake Credit Generators

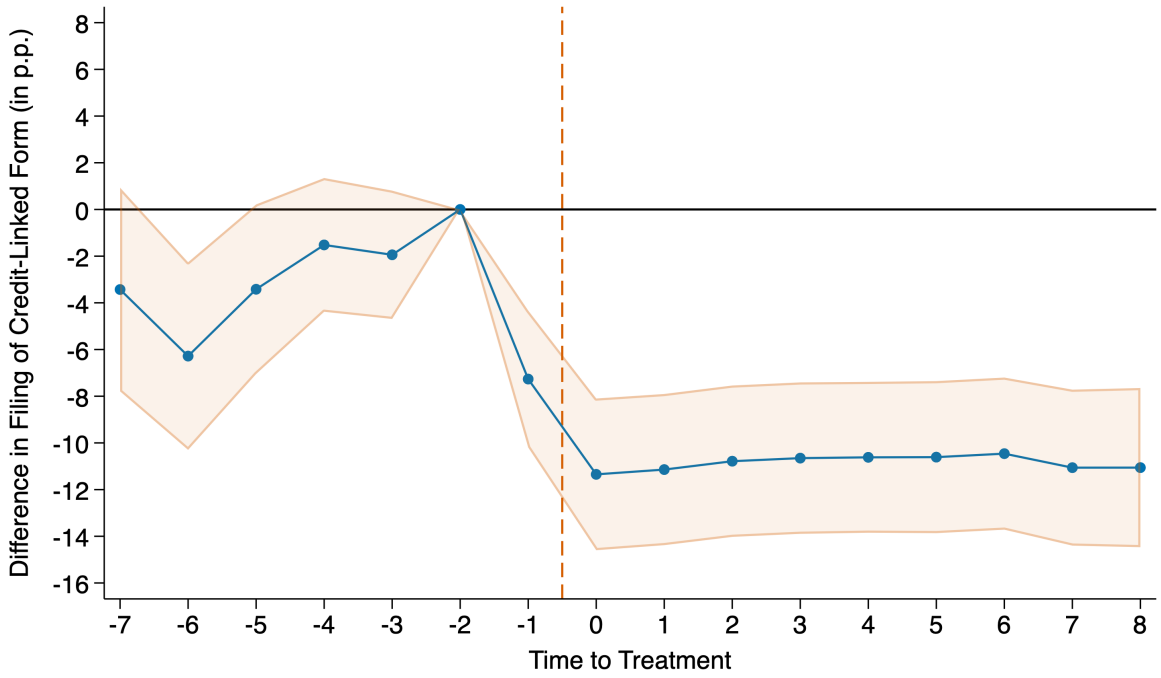

Notes. This figure shows the causal effect of the policy on filing of the credit-linked form by Non-Filer Fake Credit Generators. Non-Filer Fake Credit Generators are the firms that we classify as Fake Credit Generators who additionally don't file at least one payment-linked form between January-June 2021. We plot the event-study coefficients from Equation 1 described in Section 4, along with the 95% confidence intervals. The control group includes all firms we do not classify as Fake Credit Generators that additionally don't file at least one payment-linked form between January-June 2021. Standard errors are clustered at the firm-level.

Figure A.6: Heterogeneity in Effect on Filing of Credit-Linked Form by Pre-Period Discrepancy

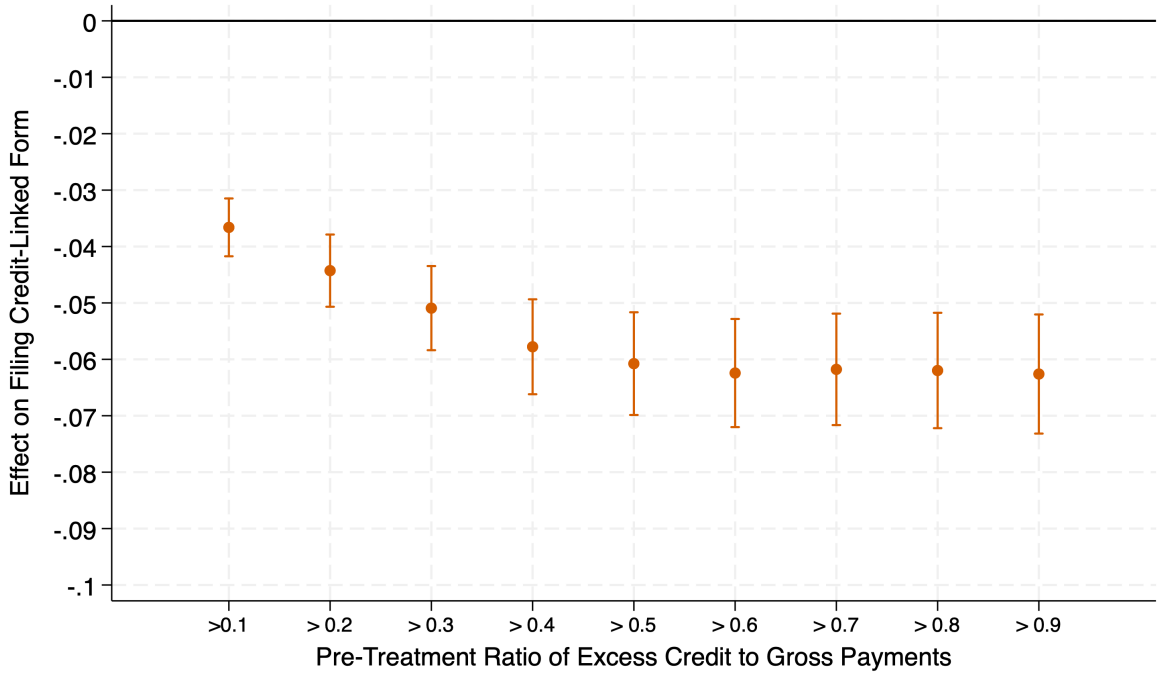

Notes. This figure plots the causal effect of the policy on filing of credit-linked form, separately for different ratios of Excess Credit to Gross Payments between January-June 2021. Total Excess Credit for a firm is defined as the total credit issued by the firm to their suppliers between January-June 2021 subtracted by the total gross payments by the firm in this time period. Total Gross Payments is defined as the sum of all tax payments made by the firm between January-June 2021. This figure excludes all firms with no reported sales. In each case, we compute the Average Treatment Effect on the Treated (ATT) comparing Fake Credit Generators with a pre-treatment ratio of Excess Credit to Total Gross Payments greater than the displayed value on the x-axis, to a control group of all firms with no excess credits between January-June 2021. We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. We also plot the 95% confidence intervals, with the standard errors clustered at the firm-level.

Figure A.7: Effect of Policy on Filing of Credit-Linked Return (Testing for Mean Reversion)

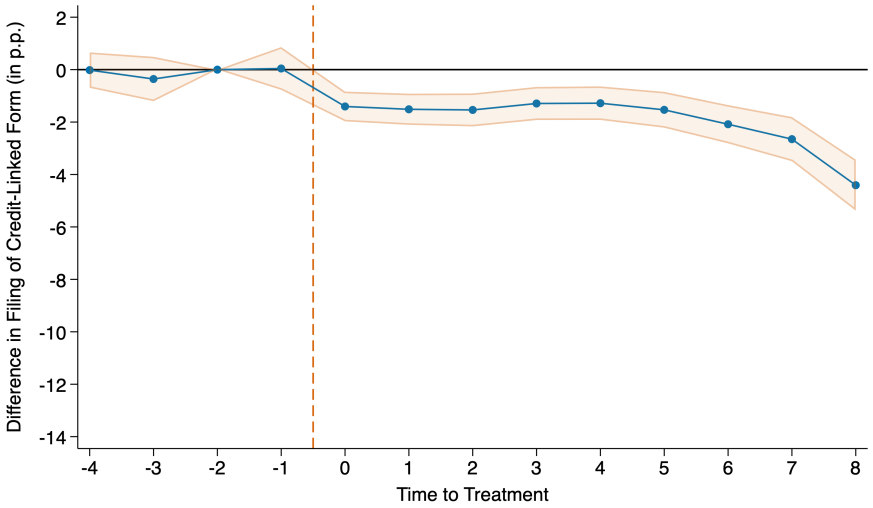

(a) Overall Sample

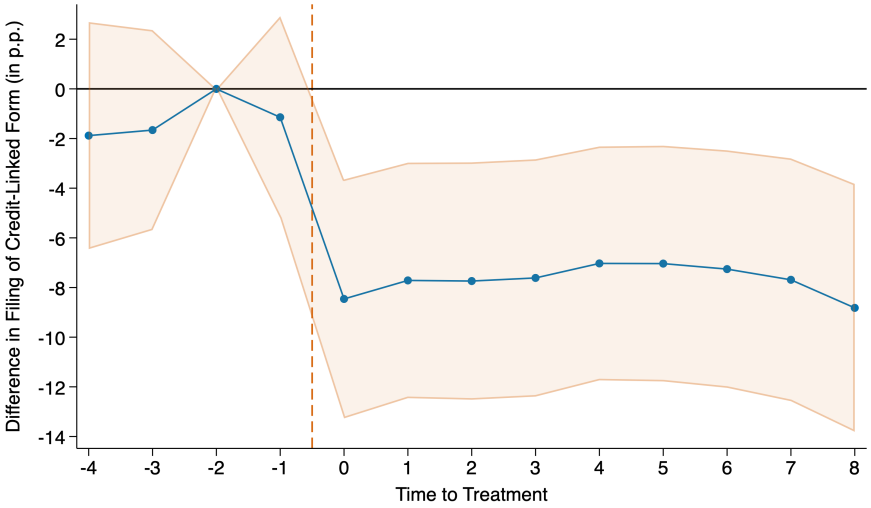

(b) Pre-Period Non-Filers

Notes. This figure shows the causal effect of the policy on filing of the credit-linked form by Fake Credit Generators. Unlike Figures 2 and A.5, this figure only uses data from January-March 2021 to classify firms as Fake Credit Generators. Firms are classified as Fake Credit Generators using the same criteria described in Section 2. The causal effect is estimated with an event-study design using data from April 2021-April 2022. We plot the event-study coefficients along with the 95% confidence intervals. Panel A compares all firms classified as Fake Credit Generators using data from January-March 2021 to all firms not classified as Fake Credit Generators. Panel B additionally restricts the analysis to firms that additionally don't file at least one payment-linked form between January-June 2021. Standard errors are clustered at the firm-level.

Figure A.8: Trends in Registration Status of Firms

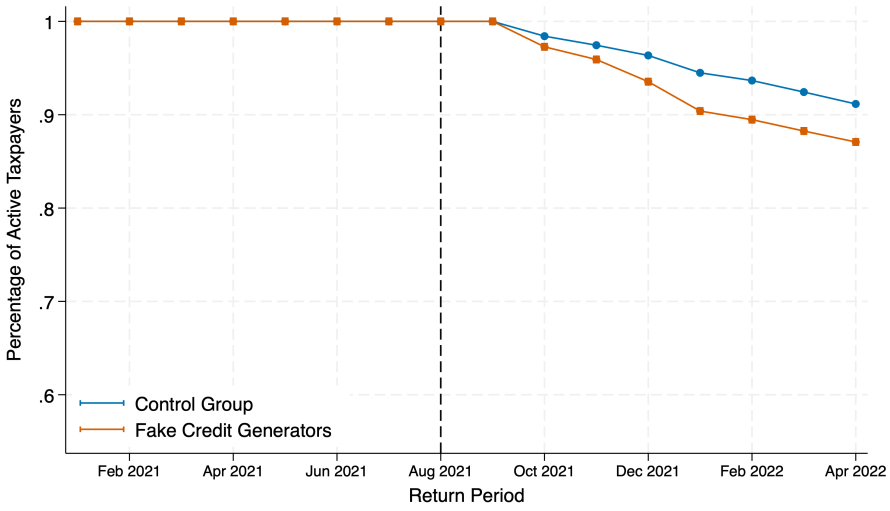

(a) Share of Active Firms

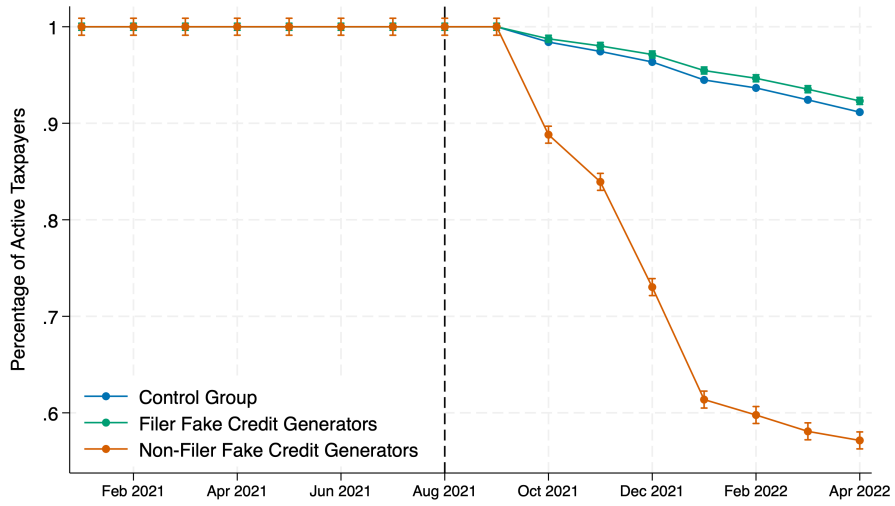

(b) Share of Active Firms by Fake Credit Generator Type

Notes. This figure presents trends in registration status of firms. We begin with a sample of firms registered in every period between January-September 2021 and classify firms as Fake Credit Generators using the procedure described in Section 2. In Panel A, we present the proportion of Fake Credit Generator firms (compared to a control group of all other firms) that are registered in each return period. In Panel B, we separately show the proportion of registered firms for Non-Filer FCGs (that don't file at least one payment-linked return between January-June 2021) and Filer FCGs (that file all payment-linked returns between January-June 2021).

Figure A.9: Trends in Registration Status of Network Firms

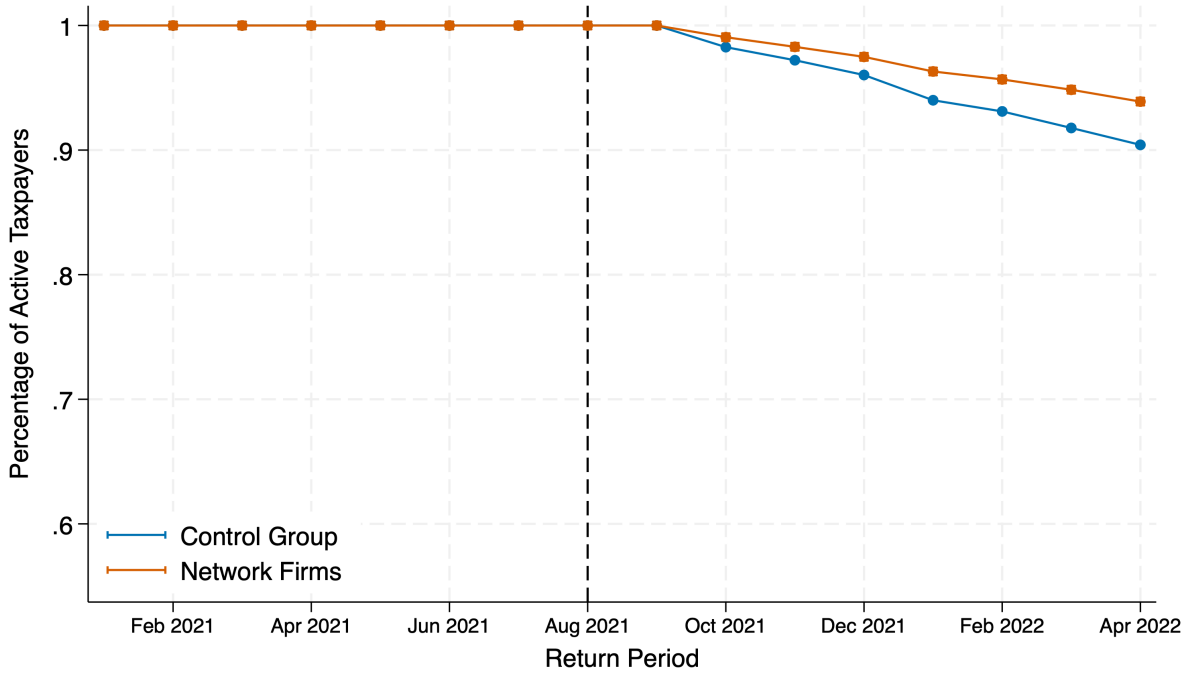

Notes. This figure presents trends in registration status of network firms. We begin with a sample of firms registered in every period between January-September 2021 and classify firms as Fake Credit Generators using the procedure described in Section 2. We define the network of these Fake Credit Generators as those firms who receive at least 5% of their total input tax credit between January-June 2021 from Fake Credit Generators and are registered in all periods between January-September 2021. We compare these network firms to all firms who receive none of their input tax credits between January-June 2021 from Fake Credit Generators and are registered in all periods between January-September 2021.

Figure A.10: Effect of Policy on Tax Payments by Network of Non-Filer FCGs

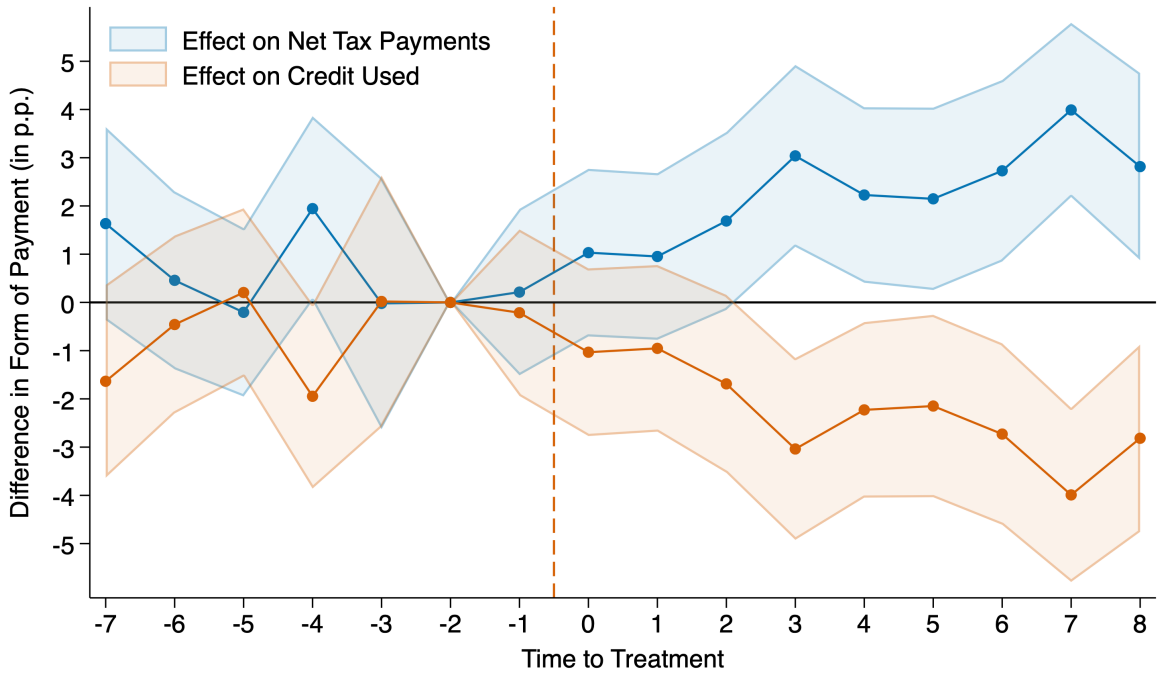

Notes. This figure shows the causal effect of the policy on the form of tax payment method by the network of those firms that we classify as Fake Credit Generators who additionally don't file at least one payment-linked return between January-June 2021. We define the network of these firms as those who receive at least 5% of their total input tax credits between January-June 2021 from Fake Credit Generators who additionally don't file at least one payment-linked return between January-June 2021. We compare these network firms to all firms who receive none of their input tax credits between January-June 2021 from this subset of Fake Credit Generators. We exclude all Fake Credit Generators themselves from this analysis. The outcome is defined as the ratio of tax payments made using cash or input tax credits to the ratio of total tax payments. For a given firm, this ratio must necessarily sum to one in a given return period. We plot the event-study coefficients from Equation 1 (defining the network of non-filer FCGs as the treated firms) described in Section 4, along with the 95% confidence intervals. Standard errors are clustered at the firm-level.

B Alternative Samples

Our main analysis focuses on a balanced sample of monthly filers who are registered throughout the study period of January 2021 - April 2022. In this section, we investigate the effect of the policy on two alternative samples, accounting for firms that deregistered within this study period and firms that elect to file quarterly—rather than monthly—returns.

First, we broaden our definition of FCGs to include firms that were registered for at least one return period between January and June 2021, so that they can be characterized as FCGs based on their discrepancies between payment-linked and credit-link forms based on their pre-policy behavior. Once registered, all firms remain in the sample in each following return period regardless of their contemporaneous registration status. Our outcomes of interest, whether the firm filed a payment-linked or credit-linked form, take a value of zero if the taxpayer did not file the form whether that be because they were no longer registered or if they were registered but non-compliant. Expanding the sample in this manner increases the magnitude of the estimated response to the policy. Table B.1 shows the main results for this expanded sample. Column 1 shows that the FCGs increase their non-filing rate of credit-linked form by 6.2 percentage points, or 31% of the baseline non-filing rate. Figure B.1 shows the event study. We also find that similar to the main analysis, there is a slight decrease in filing of payment-linked form (Column 2). When firms stop filing their returns as a result of the policy, they are more likely to become deregistered and drop out of our main sample: we observe a higher probability of non-filing of either form among FCGs after the blocking policy.Figure 1 shows that while non-filer FCGs account for 12-14 percent of the excess credit in the economy before the policy, this is effectively zeroed out by the policy. We trace the downstream effects of the disruption of this fraudulent credit on firms that are its beneficiaries, as in Section 5.2. Table B.2 shows these results, and Figure B.2 shows the event study. We find similar effects, as in the main sample. We do observe a small negative effect on gross tax payments (Column 3) and account for this potential decline in tax payments when computing the aggregate revenue implications of the policy in Section 5.3.

Next, we examine outcomes among quarterly filers who were excluded from the main analysis sample because we observe them for only two quarters before the policy. Only firms with annual revenue below Rs 1.5 crore can file on a quarterly basis, making this group of firms quite different from the main sample. They may be more likely to make legitimate mistakes in filing both their payment and credit linked forms, as they must reconcile information over three months instead of one. A similar share of quarterly filers are categorized as FCGs following the same definition as for monthly filers based on discrepancies in sales reported in the two forms. These potential FCGs react similarly to monthly filers but are less consequential for their downstream clients.

Column 1 of Table B.3 shows that quarterly filing FCGs also reduce their filing of credit-

linked forms and Column 2 shows that they are unlikely to start making payments. Similar to their counterparts who are required to file at a monthly frequency, quarterly filer FCGs are more likely to deregister as a result of the policy, as there is an increase in their likelihood of filing neither form (Column 3)

However, these quarterly filers are either more likely to be replaced or in other ways less consequential to their client firms. Just as with the monthly filers, we identify the client network of quarterly-filer FCGs as those who purchased intensively from these firms in the quarters before the blocking policy. We then compare their net tax payments and input tax credits used to that of a control group of firms with no purchases from quarterly-filer FCGs. Figure B.3 presents these event-study results. Unlike the network of monthly-filer FCGs, we see no measurable change in the net tax payments or input tax credit usage of these network firms after the September 2021, mirroring the muted decrease in credit filing rates among quarterly-filer FCGs. The revenue effects of the GSTR-1 blocking policy among quarterly filers are therefore smaller.

Table B.1: Effect on Filing Behavior of Fake Credit Generators (Including Deregistered Firms)

	(1)	(2)	(3)
	Filed Credit-Linked Form	Filed Payment-Linked Form	Filed Neither Form
Fake Credit Generator × Post	-0.062	-0.017	0.054
	(0.004)	(0.003)	(0.004)
Outcome Mean	0.800	0.786	0.165
Observations	5,911,922	5,911,922	5,911,922

Notes: This table presents the causal effect of the policy on filing outcomes of Fake Credit Generators. Unlike our baseline sample used in Table 1 where we require all the firms to be registered in every period between January 2021 - April 2022, this table only requires the firms to be registered in at least one period between January-June 2021. If a firm deregisters, they—by definition—do not file any of their forms. Therefore, we impute their outcomes following deregisteration as being equal to 0. We classify firms as Fake Credit Generators using the same procedure as in Table 1, and repeat the same analysis for the aforementioned broader sample. Column 1 presents the effect on filing of the credit-linked form in a given month. Column 2 presents the effect on filing of the payment-linked form in a given month. Column 3 presents the effect on filing of the neither form in a given month. In each case, the causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). We exclude all data after February, 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Table B.2: Effect on Network of FCGs (Including Deregistered FCGs)

	(1)	(2)	(3)
	Net Tax Payments (Share)	Credit Used (Share)	Log Gross Payments
Network Firm × Post	0.015	-0.015	-0.027
	(0.002)	(0.002)	(0.010)
Outcome Mean	0.246	0.754	11.785
Observations	1,910,588	1,910,588	1,910,588

Notes: This table presents the causal effect of the policy on the network (buyers) of firms we classify as Fake Credit Generators. Unlike Table 2 which considers the network of FCGs registered in each period between January 2021-April 2022, this table focuses on the network of the firms we classify as Fake Credit Generators that are registered in at least one period between January 2021-June 2021. In both tables, the network firms themselves are registered in each period between January 2021-April 2022. We define the network of Fake Credit Generators as those firms who receive at least 5% of their total input tax credit between January-June 2021 from Fake Credit Generators. We compare these network firms to all firms who receive none of their input tax credits between January-June 2021 from Fake Credit Generators. We exclude all Fake Credit Generators themselves from this analysis. The outcome in Columns 1 is defined as the ratio of net tax payments made by a firm using cash divided by the firm's total tax payments. The outcome in Columns 2 is defined as the ratio of tax payments made by a firm using input tax credit divided by the firm's total tax payments. The outcome in Columns 3 is defined as the log of the firm's gross tax payments in a given time period. We estimate the causal effects using the two-way fixed effects design described in Section 4. The outcome mean is the average of the outcome for the network firms before the policy was implemented. Standard errors are clustered at the firm-level.

Table B.3: Effect on Filing Behavior of Quarterly Fake Credit Generators (Including Deregistered Firms)

	(1)	(2)	(3)
	Filed Credit-Linked Form	Filed Payment-Linked Form	Filed Neither Form
Fake Credit Generator × Post	-0.032	-0.003	0.019
	(0.003)	(0.002)	(0.003)
Outcome Mean	0.977	0.959	0.013
Observations	1,155,285	1,155,285	1,155,285

Notes: This table presents the causal effect of the policy on filing outcomes of Fake Credit Generators who elect to file their forms at a quarterly frequency. Each observation reflects the firms filing outcome in that quarter. Unlike our baseline sample used in Table 1 where we require all the firms to be registered in every period between January 2021 - April 2022, this table only requires the firms to be registered in at least one period between January-June 2021. If a firm deregisters, they—by definition—do not file any of their forms. Therefore, we impute their outcomes following deregisteration as being equal to 0. We classify firms as Fake Credit Generators using the same procedure as in Table 1, and repeat the same analysis for the aforementioned sample of firms that are required to file their forms each quarter. Column 1 presents the effect on filing of the credit-linked form in a given month. Column 2 presents the effect on filing of the payment-linked form in a given month. Column 3 presents the effect on filing of the neither form in a given month. In each case, the causal effect is estimated using the two-way fixed effects approach we describe in Section 4. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). We exclude all data after Q1 2022 to remove any effect of delayed filing on the estimated treatment effect. Standard errors are clustered at the firm level. The outcome mean is the average of the outcome for the treatment group before the policy was implemented.

Figure B.1: Effect of Policy on Filing of Credit-Linked Return (Including Deregistered Firms)

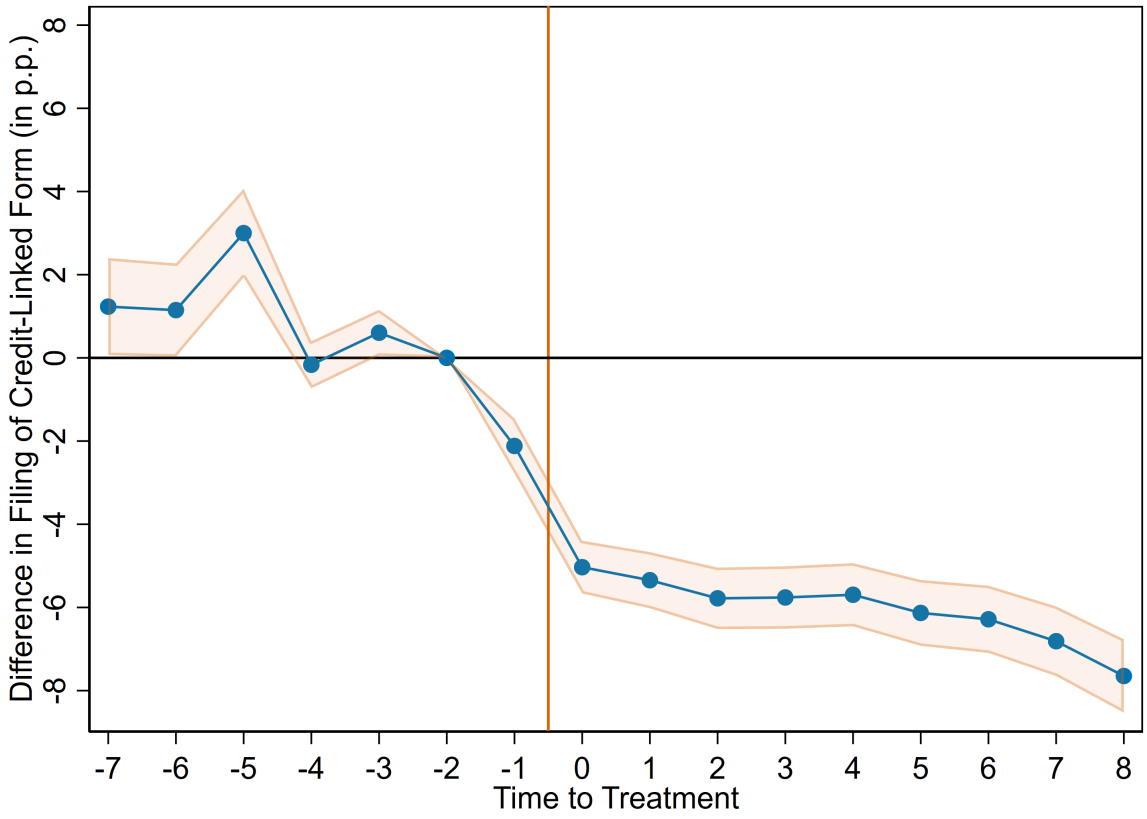

Notes. This figure shows the causal effect of the policy on filing of the credit-linked form by Fake Credit Generators. The sample is restricted to firms that were registered in at least one month between January and June 2021. Once registered, a firm stays in the sample even if they are de-registered. Since de-registered firms do not file the credit-linked form by definition, the outcome is set to 0 for these firms, just as with non-filing among registered firms. We use a matched control group, where we match on pre-treatment firm characteristics (age, size, nature of business, whether the firm reports inter-state sales, and the number of filed pre-period payment returns) using Coarsened Exact Matching (CEM). Standard errors are clustered at the firm-level.

Figure B.2: Effect of Policy on Tax Payments by Network of FCGs (Including De-Registered FCGs)

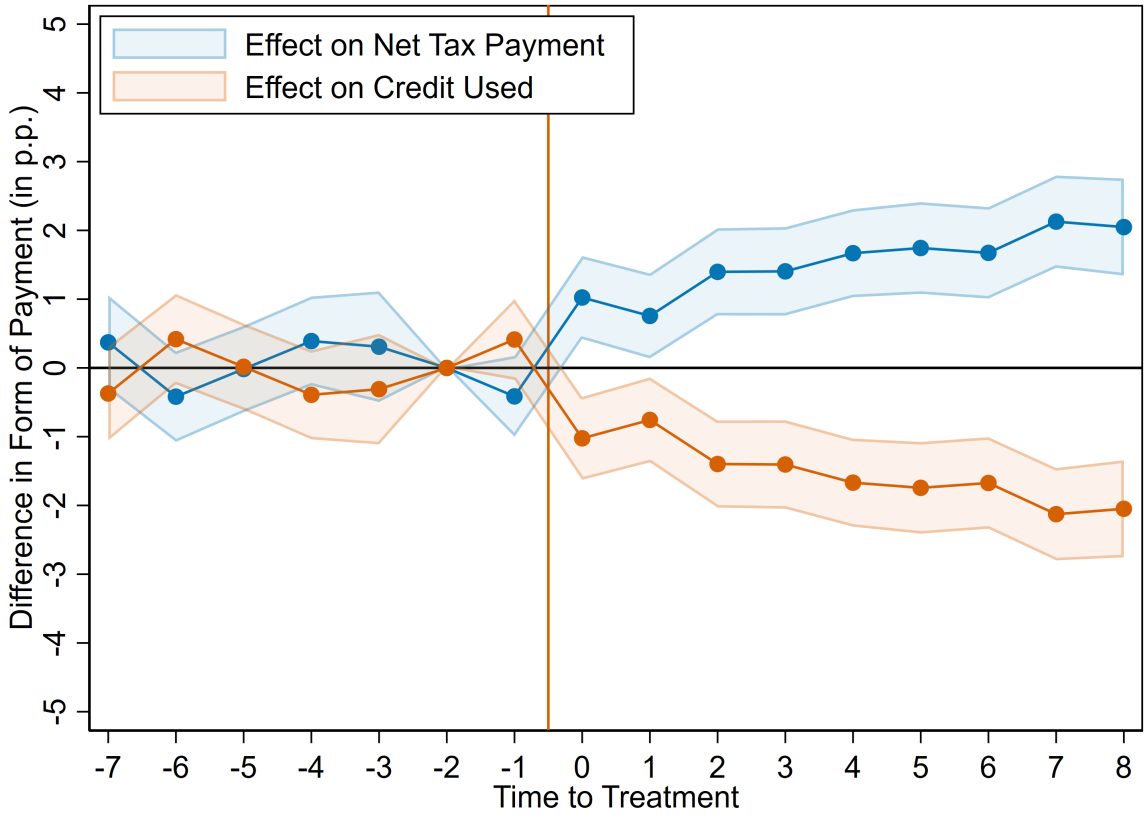

Notes. This figure replicates the results in Figure 3, but relaxes the sample on which FCGs are defined. Unlike Figure 3 which considers the network of FCGs registered in each period between January 2021-April 2022, this figure focuses on the network of the firms we classify as Fake Credit Generators that are registered in at least one period between January 2021-June 2021. In both figures, the network firms themselves are registered in each period between January 2021-April 2022. We define the network of these FCGs as those firms who received at least 5% of their total input tax credit between January-June 2021 from FCGs and were themselves registered in at least one month between January and June 2021. We compare these network firms to all firms who received none of their input tax credits between January-June 2021 from FCGs and were registered in all periods between January 2021 - April 2022. We exclude all FCGs themselves from this analysis. The outcome is defined as the ratio of tax payments made using cash or input tax credit to the total tax payments in that return period. For a given firm, this ratio must necessarily sum to one in a given return period. We plot the event-study coefficients from Equation 1 (defining the network of FCGs as the treated firms) described in Section 4, along with the 95% confidence intervals. Standard errors are clustered at the firm-level.

Figure B.3: Effect on Tax Payments by Network of Quarterly Filer FCGs (Including De-Registered FCGs)

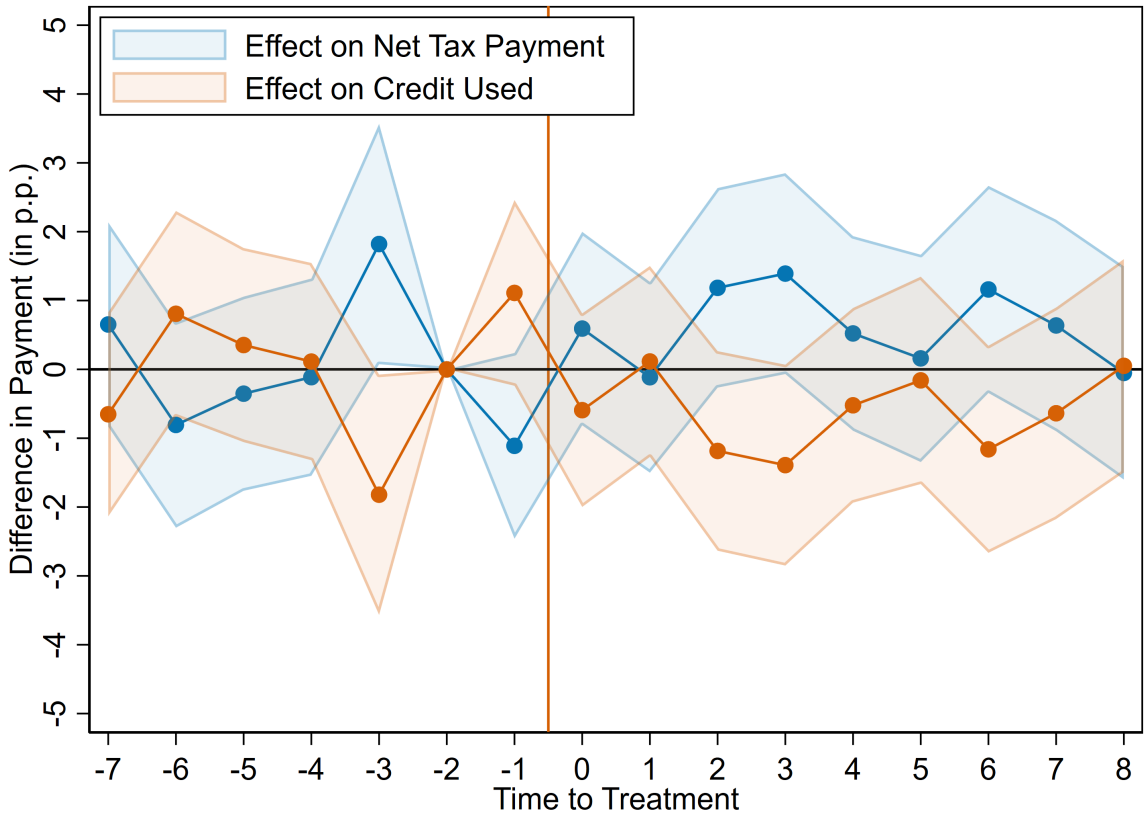

Notes. This figure replicates the results in Figure 3, but for the network of FCGs that elect to file their forms quarterly (specifically, the network of FCGs that are studied in Appendix Table B.3).