Audit publicity and tax compliance: a natural experiment*

Audit publicity and tax compliance

Pietro Battiston[†], Denvil Duncan[‡], Simona Gamba[§], Alessandro Santoro[¶]

Abstract We use confidential data on Value Added Tax payments at the sector level, in two large Italian cities, to estimate the effect of audit publicity on tax compliance of local sellers. By employing a Difference-in-Differences identification strategy, we find that such publicity has a positive effect on fiscal declarations made shortly after. The results suggest that increasing awareness on future audits via the media can be an important instrument in the hands of tax authorities.

Keywords: Tax evasion, Audits, Media coverage, Social norms.

JEL classification: H26 K34 K40

*We thank the Italian Revenue Agency for providing us with the data. Authors bear the exclusive responsiblity for the use and interpretation of these data as well as for results and comments provided in the paper.

[†]Sant'Anna School of Advanced Studies, 56127 Pisa, Italy. me@pietrobattiston.it

[‡]School of Public and Environmental Affairs, Indiana University, 47405 Bloomington, USA. duncande@indiana.edu

[§]University of Verona, 37129 Verona, Italy. gamba.simona@gmail.com

[¶]University of Milan-Bicocca, 20126 Milano, Italy. alessandro.santoro@unimib.it

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1111/sjoe.12330

I Introduction

Tax evasion is a worldwide phenomenon with significant budgetary, efficiency and equity implications. For example, it is estimated that closing the tax-gap would provide resources corresponding to approximately 60% of the UK 2013 budget deficit, 155% of the US 2006 budget deficit and 180% of the Italian 2015 budget deficit.[1] There is also evidence that tax evasion affects allocative efficiency by influencing market prices (Kopczuk et al., 2016) and the elasticity of labor supply to tax rate changes (Doerrenberg and Duncan, 2014). While precise measures of tax evasion are not available for all countries, it is commonly accepted that tax evasion is widespread and that it is a major problem especially in developing countries. Given the implications of this phenomenon, a vast academic literature has focused on understanding its determinants (see Hashimzade et al., 2013 for a comprehensive review).

While many contributions highlight the importance of tax rates, audit probability and fines (Allingham and Sandmo, 1972; Yitzhaki, 1974; Rincke and Traxler, 2011) in influencing the decision to evade, it is generally understood that tax evasion is also sensitive to other factors, such as the information taxpayers receive on the activities of tax authorities (Kasper et al., 2015). Information can reach taxpayers through three main channels: administrator-to-taxpayer communications, taxpayer-to-taxpayer communications, and media reports. These communication channels can provide information on audit frequencies and audit targets, which may affect taxpayers' perceived audit probability. Additionally, communication regarding an auditing event can influence an individual's percep-

[1]The tax gap is defined as the difference between the amount of tax that should, in theory, be collected by the revenue agency, and what is actually collected. For UK: the tax gap figure is taken from Table 1.1 in the HMRC's document *Measuring Tax Gaps Tables 2015* while the deficit figure is from Table T4.35 in the Office for Budget Responsibility's *November 2015 Economic and Fiscal Outlook: Charts & Tables*. For US: the tax gap figure is from the Internal Revenue Service while the deficit figure is from Table 1.1 of the *Historical Tables* produced by the Office of Management and Budget. We refer to the 2006 budget deficit because the most recent tax gap estimates are for that year. For Italy: the tax gap figure is taken from *Scenari Economici n. 25, Dicembre 2015* of the Centro Studi Confindustria, while the deficit figure is from Table 4 of the December 2015 "Bollettino Statistico" ("*Statistical Bulletin*") published by Bank of Italy.

tion of the proportion of evaders in the population, which in turn can affect the individual's perception of the social norms governing tax evasion. Although administrator-to-taxpayer and taxpayer-to-taxpayer communication have been shown to affect tax evasion (Slemrod et al., 2001; Alm et al., 2009; Kleven et al., 2011; Konrad et al., 2017), there is very little information on the extent to which media coverage influences the decision to evade. The impact of public disclosure of information on tax compliance has gained the attention of several governments around the world: for example, in Ireland, a list of tax defaulters was reported in national and local newspapers. According to the tax agency, this measure "aims to raise the profile of compliance and provide a continuous deterrent to other potential tax evaders" (Bø et al., 2015). However, so far, the literature has focused on the disclosure of information concerning the outcome of individual tax audits, rather than audit campaigns themselves.

The present paper contributes to this stream of literature by identifying whether the publicity of an audit affects the propensity to evade. Our identification strategy is based on an evaluation of audit blitzes which recently took place in Italy. Blitzes are defined as a set of unexpected tax verification activities taking place within a short period of time, in a small area, and on some predefined business sectors. Importantly, in recent years, some blitzes in Italy were carried out in private, while others received significant media coverage (we will refer to them as "public blitzes"). In this paper, we exploit this difference in order to identify the effect that publicity has on compliance, as detailed below.

Our dataset is provided by *Agenzia delle Entrate* (the Italian Revenue Agency) and includes data for two blitzes which took place in the Italian cities of Milan and Genoa, covering 18 business-to-consumer (B2C henceforth) sectors (the sample selection is discussed in Section II). We focus on these blitzes for several reasons. First, they both took place in January 2012. Second, they had similar characteristics, including being unannounced and focusing on a similar set of industries. Third, they differed greatly in their media coverage, with the Milan blitz being extensively covered by news outlets due to an explicit decision of the Italian Revenue Agency. Finally, Genoa and Milan are part of the same Italian macro region (North-West), are comparable in socio-economic terms and, according to province-level figures provided to us by the Revenue Agency, are similar in terms of estimated tax gap (as detailed in Section II). Hence, even though the treatment assignment was not random, the two groups of taxpayers are comparable at the baseline.

We identify the effect of media coverage through a Difference-in-Differences strategy which compares the behavior of taxpayers in the two cities using a non-parametric approach.[2] The treatment group is composed of sectors in Milan, while the control group is composed of the corresponding sectors in Genoa and the treatment is defined as exposure to the news of the local blitz. We find evidence of a positive and significant effect of publicity on VAT payments for the month of the blitz: those in Milan increased relative to the ones in Genoa. We estimate the aggregate effect of publicity (on declarations for the month of the blitz and in the 18 sectors considered) to be more than 7 million euros. The findings are robust to a range of alternative specifications and suggest that the public blitz increased compliance in the short run. Clearly, businesses in Genoa might have also reacted to the news of the Milan blitz in the same direction; our estimates should be viewed as lower bounds if this is true.

Our findings are consistent with other empirical evidence on the publicity of audit strategies. For example, there is evidence that Germany experienced an increase in voluntary disclosure of evaded taxes after publicizing the purchase of CDs containing a list of potential tax evaders.[3] This has important implications for both academics and policymakers interested in understanding the determinants of tax evasion.

The remainder of the paper proceeds as follows. Section II provides institutional details that support our decision to focus on Italy and, in particular, on the cities of Milan and

[2] We adopt a non-parametric approach for reasons related to the structure of our dataset and the limited support for normality assumptions. See Section IV for more details.

[3] See `http://edition.cnn.com/2008/WORLD/europe/02/19/tax.evasion` and Langenmayr (2017); Bethmann and Kvasnicka (2016).

Genoa. The data are described in Section III, our identification strategy in Section IV, and results in Section V. Finally, we discuss the policy implications of our results in Section VI.

II Tax evasion and blitzes in Italy

Italy provides a suitable context for testing the effect of media-publicity on tax compliance as tax evasion is widespread, the government regularly conducts blitzes, and some recent blitzes differed substantially in their media coverage. This section describes each of these features in detail.

Tax Evasion in Italy Italy is known to have one of the highest tax evasion rates among OECD countries (Buehn and Schneider, 2016), of which VAT evasion represents a significant share. Indeed, Italy is estimated to have the fifth highest VAT gap among European Union countries (CASE, 2016). Although evasion is a nation-wide issue, the propensity to evade is well known to be heterogeneous across regions and sectors (Marino and Zizza, 2012; Pisani, 2014). According to estimates provided by the Italian Revenue Agency, the regional propensity to evade tends to be lower in the North of the country, and higher in the Center and especially in the South.[4] Unpublished data provided by the Revenue Agency show that, at the *province* level, the tax gap estimated for Milan is close to that estimated for Genoa (18.8% vs. 19.4% in the pre-blitz year 2011, while the estimate for all Italy is 28.8%). While estimates are not available at the *city* level, both cities represent the vast majority of commercial activities in their own province; hence, the similarity suggests that our results are unlikely to be driven by pre-intervention differences in evasion between the two cities.

[4]For example, in the period 2007 to 2010, the estimated VAT gap was below the national average of 26.04% in six of the eight Northern regions and above 32% in six of the eight Southern regions (D'Agosto et al., 2014).

As for heterogeneity across sectors, the available evidence indicates that B2C sectors are more prone to tax evasion. Indeed, they are less exposed than business-to-business sectors to the so-called "VAT paper trail" (Pomeranz, 2015).

Blitzes The Italian government has implemented many policies to address tax evasion. These include campaigns aimed at improving tax morale and public consciousness, the use of presumptive taxes, the increase in penalties for evaders and in the frequency of tax verification activities, including blitzes.

Blitzes in Italy are usually conducted by the *Agenzia delle Entrate* in collaboration with a specialized finance police force (*Guardia di Finanza*) and sometimes with the support of inspectors from the Ministry of Labor as well as local policemen. During blitzes, agents show up unexpectedly to check for the correct issuance of receipts,[5] the integrity of cash registers, the regular updating of account books, the congruity of declarations previously made concerning several aspects of the shop (e.g., number of rooms and electrical appliances), and the presence of workers not on the books. Blitzes not only lead to economies of scale in organizing audits, but they also provide the revenue agency with a comprehensive snapshot of fiscal compliance for a given geographic area or economic sector at a given point in time. Still, because they usually target only a few dozen to a couple hundred businesses, the direct effect that a single blitz can have on the total amount of tax evasion is negligible.

Although complete data are not available, blitzes are not an uncommon instrument among Revenue Agencies: in Italy, at least 1,800 businesses, located in almost all regions, were inspected during blitzes which took place in the first half of 2012 alone (Italian Government, 2013).

[5]This is done either by checking clients exiting a shop, or by agents in plain clothes inside the shop. Furthermore, once agents show up, their presence naturally enforces the release of receipts. This allows a comparison of the amount of registered sales with the amount of registered sales in previous days, by analyzing the cash registers.

Publicity The Italian blitzes of the last few years varied sharply in the amount of media attention they received. Blitzes are usually private, in the sense that they do not receive much media coverage, and only shop owners/employees who are affected by the blitz are aware that one is taking, or has taken, place. On the other hand, in recent years, two blitzes received extensive media coverage so that every shop seller or business owner in the city of the blitz probably ended up being aware that a blitz was occurring or had occurred. The first one took place in Cortina d'Ampezzo in December 2011, and the other one in Milan in January 2012. The public nature of these blitzes was a choice of the Revenue Agency, presumably based on directions from the Italian Government. During winter 2011-2012, the recently installed Monti government was facing a major public finance crisis, and it had just passed a Budget Law including a number of tax increases and expenditure cuts. These policies were not welcomed by a vast part of the population, and several voices raised concerns that lower tax rates could be afforded, had Italy succeeded in reducing tax evasion. Thus, the decision to "go public" was probably motivated by the need to show that the fight against tax evasion was an organic component of consolidation efforts.[6]

Analyzing the effect of the blitz in Cortina d'Ampezzo would be difficult for at least three reasons. First, Cortina, a very famous winter holiday resort, has a unique economic context in which luxury goods and services represent an exceptionally large component of economic activities; comparing it to other towns, even in the same geographic area, would make little sense. Second, Cortina is a small town, and thus time series of aggregated tax payments are noisier (see Section III for more details on the structure of our data). The blitz in Cortina and its media coverage mostly focused on controls on individual possessions (which may not be relevant for the fiscal behavior of sellers), rather than on shop audits. For these reasons, our analysis focuses on the public blitz which took place in Milan.

[6]It should be noted that the two experiments of "public blitzes" were unprecedented and unreplicated: nowadays, blitzes are rarely, if ever, discussed in the media. This in turn may be a consequence of the hot public debate which revolved around the alleged spectacularization of the Cortina and Milan blitzes, as well as a consequence of the end of the Monti government.

The Milan blitz, which started at 8:30 p.m. on Saturday, January 28, 2012, instead focused on restaurants, night clubs and discotheques, and continued the following morning when more restaurants, cafés and shops in the city center were subject to audits. Overall, the Milan blitz covered approximately 350 businesses. Agents mainly verified compliance of sales reports (including the regular release of receipts), and national and local TV stations were informed and allowed to broadcast these activities live.

Our research question asks whether the publicity of the blitz had an effect on compliance. In order to answer this question, we require data on public and private blitzes that occurred in a comparable geographic area, on the same sectors and at approximately the same time. As already mentioned, in 2012 a number of private blitzes were conducted in a number of Italian cities (such as Genoa, Turin, Bari and Cagliari, as well as in smaller towns). However, only the blitz in Genoa took place in the same month as the public blitz in Milan. Moreover, most of these cities or towns are not comparable to Milan in several respects including size, wealth, geographic location, and pre-existing propensity to evade; as the link between these observable variables and the reaction of taxpayers to a blitz is far from obvious, it is difficult to define an objective measure of similarity in this respect.[7] Therefore, we focus on the Genoa blitz held on January 6, 2012, that covered approximately 150 businesses (including ice cream parlors, bars, discotheques, restaurants and clothing shops), lasted until late in the night and, like the Milan blitz, was concentrated in less than 24 hours. As already illustrated, the provinces of Genoa and Milan display a similar estimated tax gap. Also, these two cities are similar in socio-economic terms: together with Turin, they represent the three largest cities in the North West of Italy and are often referred to as the "industrial triangle", due to the important role they share in the history of Italian manufacturing. Today, they are among the richest cities in Italy, with a per-capita GDP of € 21,896.98 (Genoa) and € 29,803.77 (Milan) in 2014.

[7] See Appendix A for an exercise in adoption of a "synthetic control" approach, and caveats.

Table 1: Media coverage of different blitzes

Blitz	Google News	Newspapers I	Newspapers II
Milan	326	35	19
Genoa	40	9	6

Note: Reported is the number of articles on the blitz conducted in each city by data source. The Google News results are as of 26th of November 2015, selecting the category News, the date of the blitz and using "blitz + evasione+city" as search criterion. Newspapers I refers to "*Eco della Stampa*", Newspapers II to the web archives of *La Repubblica* and *Il Corriere della Sera*.

Evidence of Publicity We perform several checks in order to substantiate the difference in media coverage between the blitz in Genoa and the one in Milan. These include archival searches on the following sources:

- the Google News search engine (Italian version);
- the news database "*Eco della Stampa*", for a time window of 30 days before and after each blitz;
- the web archives of the two most important Italian newspapers (*La Repubblica* and *Il Corriere della Sera*), for a time window of 30 days before and after each blitz.

The results, presented in Table 1, confirm that the Milan blitz had much wider media coverage than the Genoa blitz. A *Google News* search for "*blitz*", "*evasione*" (evasion), and "*Milano*", with a date range of the 28th to 29th of January 2012, yields 203 entries reporting 326 articles on the topic. On the other hand, the search of "*blitz*", "*evasione*", and "*Genoa*", with a date range of the 6th to 7th of January 2012, yields 5 entries and 40 articles. This disparity is confirmed by a search in the newspaper archives: while the Milan blitz received between 19 and 35 mentions, Genoa is only mentioned 6 to 9 times.[8] Finally, a search on Twitter for the words "*blitz*" and "*evasione*", limited to tweets written

[8] Results from a different search also confirm the surge in news related to tax evasion (only) in Milan *after* the blitz (see Appendix B.1. The Milan blitz was even covered in a video published on the YouTube channel of the Italian newspaper "Il Fatto Quotidiano": as of November 2015, it had scored 7240 views.

Figure 1: Discussion of audit blitzes on Twitter

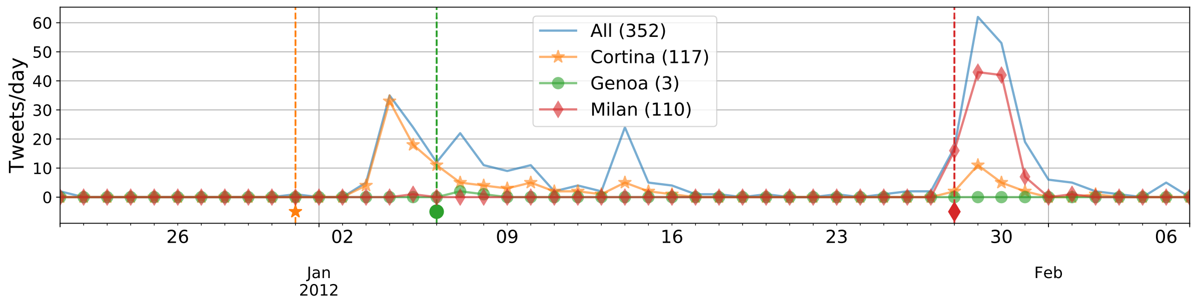

Note: daily count of tweets written in Italian and containing both the words "evasione" and "blitz", in the period of interest. For each of Cortina, Genoa and Milan, the subsample of tweets mentioning the city is also represented, together with a vertical line denoting the date of the local blitz.

in Italian, shows a clear peak after the Milan blitz (and after the previously mentioned Cortina blitz), and nothing comparable for the Genoa blitz (Figure 1).[9]

Auditing Effort A possible concern for identification is that the blitzes in Milan and Genoa differed in auditing effort. We provide some evidence that, on the contrary, auditing efforts were very similar between the two blitzes. In particular, we calculate two "blitz ratio" measures using data on firms from the 18 sectors represented in our sample (see Table 2).[10] The numerator is the number of firms involved in the blitz, while the denominator is equal to the average number of firms making VAT declarations per period (month or quarter – see note to Table 2) for tax year 2011 in Measure 1, and to the number of declarations for January 2012 in Measure 2. Because firms differ in their filing frequencies – monthly or quarterly – each of these measures is defined separately for monthly declarations only, quarterly declarations only, and for all declarations (monthly and quarterly). In general, we observe that the ratios just defined are roughly the same in Milan and Genoa, suggesting that auditing efforts do not significantly differ between the two cities. It should be noted that both Panel A and Panel B provide partial measures, since they take into account either

[9] `https://twitter.com/search?l=it&q=evasione%20blitz%20since%3A2011-12-21until%3A2012-02-07`

[10] Ideally, audit effort should be calculated using data on firms in all sectors. However, we only observe the number of VAT paying entities that belong to the 18 B2C sectors included in our analysis (see Section III).

monthly or quarterly declarations. Figures from Panel C, where both firms paying VAT monthly and firms paying VAT quarterly are considered, are more reasonable estimates of the share of blitzed firms, although the reported 4% is still an upper bound (a firm does not issue any VAT declaration for a given month/quarter if it has no VAT due).

Table 2: Blitz ratios

	Measure 1 (year 2011)			Measure 2 (Jan/Q1 2012)		
	Audits	N. Obs	Ratio	Audits	N. Obs	Ratio
	Panel A: monthly declarations					
Milan	350	2059.7	0.17	350	2031	0.17
Genoa	150	672.8	0.22	150	644	0.23
	Panel B: quarterly declarations					
Milan	350	4614	0.08	350	6575	0.05
Genoa	150	2419.6	0.06	150	3541	0.04
	Panel C: all declarations					
Milan	350	6673.7	0.05	350	8606	0.04
Genoa	150	3092.4	0.05	150	4185	0.04

Note: Data from the 18 sectors considered in our study. Measure 1: average of monthly/quarterly/all declarations for 2011. Measure 2: declarations for January 2012/first quarter of 2012/January 2012 and first quarter of 2012.

Effects of a blitz There are multiple channels through which a blitz can be expected to affect tax compliance:

1. *while agents are present in the shop*, unreported sales virtually disappear. This effect however lasts only a few hours, and affects only a relatively small number of shops (less than 4% in the two blitzes analyzed, as shown in the previous paragraph);

2. shortly after an audit, audited shops might change their behavior (the sign of this change is not obvious: audits might signal an increase in investment of the local Revenue Agency in repressive activities, but the perceived probability of being audited twice in a short time span might be particularly low);

3. word-of-mouth about the audits independent from media coverage can reach owners of business that are not being audited;

4. media coverage: owners of other activities can get to know about the blitz from official publicity (or from word-of-mouth which originates from it).

Except for the last, we can expect that these effects are present in Milan and Genoa to a similar extent. Hence, the comparison of the two cities allows us to identify specifically the effect of official publicity. It is important to observe that "publicity" included more than just producing official press releases: agents coordinating the operation released interviews, and passers-by could see not just agents, but also a large number of journalists, at work.

Our setting and our approach are complementary to those of Rincke and Traxler (2011). In their experiment, they exploit microdata on the enforcement of TV licenses in Austrian households. Enforcement is directed at individual households and is *not* publicly observed. Thus, their focus is on the impact of interpersonal communication, i.e., word-of-mouth *independent* from publicity; they do not measure the effect of publicity per se, which is the main target of our study. Another difference is that they are dealing with individuals rather than with firms. On the contrary, Pomeranz (2015) identifies an increase in compliance for firms trading with those that receive the announcement of an audit. In such a case the enforcement spillover is driven by a network effect because increasing audit probability at the end of the production chain increases compliance along the chain. This is again similar, in terms of spreading mechanism, to word-of-mouth (*independent* from media coverage), and differs strikingly from a publicity effect. Thus, our paper is related to the literature on spillover effects of tax audits, but it considers a channel that of media coverage which has not been investigated so far.

Media reports can play an important role in shaping tax morale (in a country in which a former Prime Minister claimed that evading taxes was "morally justified"), by stressing the negative consequences of non-compliance on the public budget or by highlighting the burden that non-compliance places on the shoulders of compliant taxpayers. Both these arguments are rooted in the public discussion on tax evasion in Italy, and both were mentioned in articles and declarations concerning the Milan blitz.

One may wonder whether the Milan blitz might have been relevant for Genoa taxpayers also. Interestingly, since these two cities belong to different regional directorates, and since budget and audit targets are allocated ex ante among these directorates,[11] there is no reason for a rational taxpayer in Genoa to change their compliance after the blitz in Milan. However, the media coverage and discussion about the blitz in Milan may have had overall effects that also affected Genoa taxpayers by enhancing tax morale or the salience of the event of an audit. In such a case, VAT payments by Genoa taxpayers could actually have been increased by the Milan blitz, and our Difference-in-Differences estimation would represent a lower bound of the effect of the public blitz.

[11]This is stated in Article 5, Point 6, of the Administration Act (*Regolamento di Amministrazione*), available at `http://www.agenziaentrate.gov.it/wps/file/Nsilib/Nsi/Agenzia/Chi+siamo/Statuto+e+appositi+regolamenti/Il+regolamento+di+amministrazione/T1_reg_amministrazione_1+gennaio+2016_II_pubblicato.pdf`

III Data

Our empirical analysis is based on a confidential database obtained from the *Agenzia delle Entrate*. The database includes a panel of monthly IVA (the Italian VAT) payments for 18 B2C sectors from January 2009 to November 2013, in Milan and Genoa.[12] As a general rule, Italian firms are required to pay IVA to the Italian Revenue Agency monthly.[13] Each IVA payment consists of the (declared) difference between IVA collected on sales and IVA paid on purchases.[14] Such payments are reported on Form "F24", which is submitted electronically. The deadline for presenting monthly IVA payments is the 16th of the following month; importantly, in our data, each month m corresponds to the amounts *due* for that month, and hence refers to earnings and costs in month $m-1$. As an exception, declarations for the month of December, the last month of the fiscal year, include VAT payments based on both actual November sales *and* projected December sales.

For privacy reasons, the data provided to us are aggregated at the sector level for each city-month, so they do not allow an analysis at the firm level. For each city-sector-month, however, we know both the sum of IVA payments and the number of taxpayers. In total, our panel comprises 2124 observations (59 months, 2 cities, 18 sectors).

Firms must remit the amounts declared on the F24 form to the Italian Revenue Agency; therefore, non-compliance in the act of filling out the F24 form represents an act of tax evasion. Notice that, in the eventuality of tax audit controls, the tax police can verify that the payments correspond to the difference between IVA on sales and IVA on costs: each must correspond to the sum of amounts reported on receipts issued and received, respectively. Typically, evading IVA involves selling a good or service without issuing a

[12] A complete list of the 18 sectors is provided in Table B.1 of Appendix B.

[13] Firms whose turnover in the previous year was below some specified thresholds are allowed to pay VAT quarterly.

[14] Importantly, while the VAT rate depends on the type of product and has changed several times over the years, it applies uniformly to the entire Italian territory, i.e., changes affected Genoa and Milan at the same time.

receipt[15] (Fabbri and Hemels, 2013; Battiston and Gamba, 2016) and hence under-reporting sales. Indeed, for B2C activities, this is much easier than fabricating evidence of non-existent purchases of inputs in order to over-report costs (this is the essence of the already mentioned "paper trail"). Tax audits in shops typically focus on non-compliance on the sales side, i.e., by checking or enforcing that receipts are regularly issued. The finance police have limited ability, during the blitz, to ensure that costs are recorded accurately (because, for instance, in B2B transactions, invoices do not typically travel with the goods).

We deflate aggregated IVA payments using city-specific monthly price indexes released by the Italian National Institute of Statistics (ISTAT). We then divide them by the number of IVA-taxpayers in each city-sector-month: in other words, we look at *average* payments[16] rather than at the total amount for the sector. This normalization allows us to take into account differences in size, and to control for changes in the population of IVA-taxpayers within a city over time. Table B.2 in Appendix B presents descriptive statistics at the sector level.

Figure 2 shows total IVA payments over time for Genoa and Milan (see Figure B.1 in Appendix B for per-taxpayer data): the two time series are clearly highly correlated,[17] and strongly shaped by seasonal fiscal deadlines. This brings additional support to the soundness of using Genoa as a counterfactual for Milan, and at the same time highlights the importance of considering seasonal effects.

IV Identification Strategy

We adopt a Difference-in-Differences identification strategy comparing the effect of the public blitz in Milan to the private blitz in Genoa. Let $M_{m,y,s}$ denote deflated average per

[15] Of each receipt released to a client, a copy is kept in the cash register.

[16] The 2124 observations become 2056 once taking into account observations with no payments (see Table B.3, Appendix B).

[17] At a sector level, estimated correlation between Milan and Genoa is positive for 17 sectors out of 18, and significantly so for 16 sectors out of 18: it is, on average, 0.493.

Figure 2: History of reported VAT payments

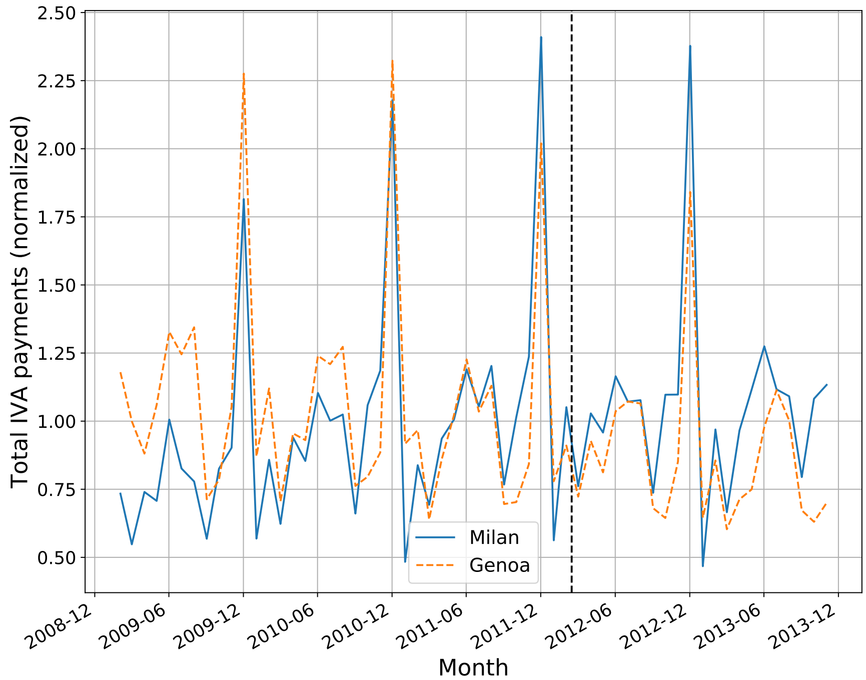

Note: Total VAT payments, across the 18 sectors analyzed. For ease of comparison, both series are normalized so that they average to 1. The dashed line identifies the month of the blitz.

taxpayer VAT payments in Milan in month m of year y, for sector s, and $G_{m,y,s}$ denote the same for Genoa. We are interested in testing the hypothesis that, in the month following the blitz, the difference between $log(M_{m,y,s})$ and $log(G_{m,y,s})$ is larger than in other months.[18] We disregard quarterly VAT payments because of their low temporal resolution.

A crucial feature of our data is the presence of multiple time series, one for each of the 18 sectors. Although the sectors are unified by a theoretical opportunity to evade, they differ in many other respects. Some sell services or goods which are consumed daily (e.g., bakeries), others provide goods for which consumption may be more volatile (e.g., clothing shops); they also differ in the typical size of the firm, in the average value of goods or services offered, and possibly in the seasonality of sales and fiscal deadlines. In addition, there may be a different response to information about blitzes between sectors in which the customer goes to the service provider (e.g., restaurants) and sectors in which the service provider goes to the customer (e.g., plumbers). More generally, there is no support for the assumption that VAT payments are similarly distributed across sectors, and sectors are extremely different in their relative importance (see Table B.2 in Appendix B). These issues prevent us from directly testing our main hypothesis via a pooled test on the entirety of our data, as the aforementioned issues would be only partly solved by including sector- or time-fixed effects in our model specification. Instead, we will run our analysis sector by sector, and then aggregate the results.

Even within sectors, and controlling for seasonal effects, there are no hints that the distribution of VAT payments over time should be normal,[19] and the number of available observations at the sector level (59 per city) makes asymptotic assumptions inappropriate. That is, the assumptions for OLS-based inference are unsatisfied. Hence, we rely on a non-parametric approach, which abstains from any distributional assumption. For each

[18] We work on the logarithm of VAT payments because we expect the effect, if any, to be proportional to the pre-blitz level. Notice that we focus on the *immediate* impact of the blitz, i.e., on the payments regarding the month of the blitz. See Section V for a check of existence of long-term effects.

[19] We checked whether residuals of Equation 3 are normally distributed: normality is rejected for most sectors - see Appendix B, Table B.3.

city and sector, we can decompose the time series of deflated VAT payments per taxpayer in a component explained by year- (Y_y) and month- (Z_m) specific dummies, and a residual. So, for any sector s, we will have for Milan:

$$\log(M_{m,y,s}) = \alpha_M + \sum_{y=2010}^{2013} \beta_{M,y} Y_y + \sum_{m=2}^{12} \gamma_{M,m} Z_m + \epsilon_{M,m,y,s} \tag{1}$$

and for Genoa

$$\log(G_{m,y,s}) = \alpha_G + \sum_{y=2010}^{2013} \beta_{G,y} Y_y + \sum_{m=2}^{12} \gamma_{G,m} Z_m + \epsilon_{G,m,y,s} \tag{2}$$

where year- and month-specific dummies allow us to control for macroeconomic trends and fiscal deadlines (the effect of which is evident in Figure 2), respectively. Any effect of the blitz in Milan in January 2012 will be included in the residual for the corresponding month, and the same can be said about Genoa.

In order to analyze the effect of blitz publicity, we take the difference between the two city-specific equations. That is, we estimate the following model through OLS:

$$\delta_{m,y,s} = \log(M_{m,y,s}) - \log(G_{m,y,s}) = \alpha + \sum_{y=2010}^{2013} \beta_y Y_y + \sum_{m=2}^{12} \gamma_m Z_m + \epsilon_{m,y,s}, \tag{3}$$

where $\epsilon_{m,y,s}$ is hence the component of the *difference* in payments between Milan and Genoa that cannot be explained by year- and month-fixed effects. In principle, we could have based our analysis directly on the difference (across sectors, years, months) between $\epsilon_{M,m,y,s}$ and $\epsilon_{G,m,y,s}$. The reason to focus on $\epsilon_{m,y,s}$, i.e., to run the estimate directly on $\delta_{m,y,s}$, is that this allows us to net out any episodic shock affecting both cities *before* the estimation, hence reducing the noise. That is, while the validity of our analysis requires no specific assumption on the correlation of residuals between Milan and Genoa, the decision to focus on the cross-city difference is made precisely in order to exploit any such correlation.

The residual component $\epsilon_{m,y,s}$ may be shaped by a multitude of unobservable factors which affect the two cities in different ways. This makes it challenging to distinguish any medium- or long-term effect of the blitz from confounding factors (e.g., difference of the business cycle between the two cities). We instead exploit the discontinuity represented by the blitz by focusing our attention on its *immediate* effect, i.e., the effect on VAT declarations immediately following the blitz. Our identification strategy thus relies on the assumption that, with the exception of the blitz, the probability of such factors being exceptionally strong precisely in the month of the blitz is very low.

After estimating Equation 3 for each sector s separately, we pool the residuals $\epsilon_{m,y,s}$ and run a non-parametric Mann-Whitney (MW) test[20] on the null hypothesis that values in the set

$$\mathcal{B} = \{\epsilon_{m,y,s} | (m, y) = (2, 2012)\}$$

follow the same distribution of the values in the set

$$\mathcal{C} = \{\epsilon_{m,y,s} | (m, y) \neq (2, 2012)\}$$

against the alternative hypothesis that values in $\mathcal{B}$ are larger. This one-sided test is run both for a period from January 2009 to February 2012 (specification "$\mathbb{PRE}$"), and for the entire sample period (January 2009 to November 2013: specification "$\mathbb{ALL}$"), hence including the months after February 2012 in the control sample. The $\mathbb{PRE}$ specification is more reliable in the presence of any medium- or long-term trend, while the $\mathbb{ALL}$ specification can achieve higher precision by better exploiting available data. In what follows, results refer to the $\mathbb{ALL}$ specification unless stated otherwise (see Section V for more evidence supporting this choice). Results for the $\mathbb{PRE}$ specification do not differ significantly.

[20]See Appendix C for details.

Comparison with standard Difference-in-Differences approaches As with any Difference-in-Differences approach, our model is parsimonious in terms of data, since by design it does not exploit explanatory variables which are supposedly constant either across time (e.g., average income, number of shops in the city) or across cities (e.g., fiscal deadlines). However, standard Difference-in-Differences approaches are based on the assumption of i.i.d. observations, which may result in severely understating the standard deviations of the estimators (Bertrand et al., 2004). In our context, this problem could, in principle, arise as a consequence of correlation between cities and/or correlation across sectors.

As for the first part, we believe that we are not affected by this problem because we are not assuming independence between cities. We would be doing so if Equation 1 (or an analogous model) was estimated using pooled VAT declarations for Genoa *and* for Milan (including a dummy variable for the city among the regressors) rather than estimated using their difference, as we are doing. We account for non-independence by collapsing data across the space dimension, which is analogous to the suggestion of Bertrand et al. (2004) to collapse data into the time dimension (our approach is feasible, of course, only because the data are restricted to two cities). As for the second part, i.e., independence across sectors, we verify that residuals are not correlated.[21]. We find that pairwise Pearson correlations between pairs of different sectors are concentrated around 0, the average correlation is very small (0.001) and not significantly different from 0.

The assumption of identical distribution required by conventional standard errors is clearly not required by our non-parametric method, as previously mentioned.

[21]With this test, we ensure that the correlation across sectors is explained by seasonal and yearly effects

Figure 3: Average residual over time

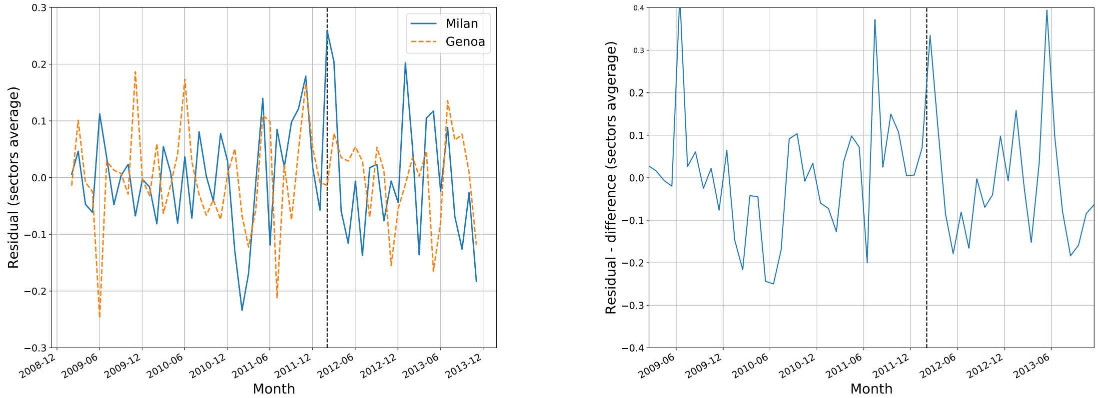

Note: Left: residuals from equations 1 and 2, averaged over sectors. Right: residuals from Equation 3, averaged over sectors. The dashed line identifies the month of the blitz.

V Results

Main Findings

Figure 3 features residuals over time. The right panel shows a spike for February 2012. This is not the only one, as three other spikes are visible in the figure. However, from the left panel, it can be seen that these other spikes are associated with negative peaks for residuals in Genoa, rather than with a particularly high value of residuals in Milan. Instead, the peak recorded in February 2012 (right panel) is due to an exceptionally large positive value of residuals in Milan, hence it is consistent with a positive impact of the blitz.

Turning to the statistical analysis, the one-sided MW test on the null hypothesis of equality between values in $\mathcal{B}$ and values in $\mathcal{C}$ rejects it with a p-value of $p = 0.015$ ($p = 0.019$ in the $\mathbb{PRE}$ specification): values in $\mathcal{B}$ are larger. This means that the increase in tax payments in Milan for the month of the blitz was significantly larger than in Genoa.

In order to estimate the magnitude of the effect, we look at the values of the unexplained component $\epsilon_{m,y,s}$ across sectors:[22] we find that such residual is on average 0.271 (the

[22]The values of ϵ for each sector are shown in Figure B.2, Appendix B.

average across all months being zero by definition) for February 2012. An effect of 0.271 in logarithmic terms (i.e., on $\delta_{m,y,s}$) corresponds to an increase by $e^{0.271} - 1 = 31.1\%$ of the ratio $e^{\delta_{m,y,s}} = \frac{M_{m,y,s}}{G_{m,y,s}}$. This in turn translates to € 4,060 of extra VAT payment per taxpayer,[23] for a total of € 7,690,899 in Milan for the month and sectors considered (corresponding to 23.7% of the average monthly payments in Milan).

This is the most appropriate estimate if we expect that differences in the effect across sectors are random, and we aim at extrapolating an average effect of a generic publicized blitz on the whole population of shops in Milan. If we expect instead that different sectors may have intrinsically different propensities to react to news of the blitz, it is more appropriate, in order to estimate the total effect of publicity, to calculate the absolute effect in each sector, and then sum up the results: this yields a total of € 6,222,934 extra VAT payments. The fact that this estimate is smaller than the previous one means that sectors with higher declared revenues per taxpayer tended to react less. Although the difference is minor, it matches the hypothesis that smaller shops (e.g., shops in which the owner is typically also a seller) have a relatively larger tendency to evade. This measure, however, different from the previous one, is heavily dominated by the results for bigger sectors in terms of VAT payments.

Estimates of prevented VAT evasion presented so far obviously focus only on the 18 sectors under analysis; the total effect could in principle be much larger. Also excluded from the analysis are businesses which, due to their small size, are allowed to file VAT payments once every three or twelve months, rather than every month (the data we used do not include their payments, because of their coarser temporal accuracy).

In general, our approach assumes no externalities, i.e., that the publicity of the blitz in Milan had no effect on compliance decisions in Genoa. However, this does not hinder the

[23]Let R be the observed ratio between payments in Milan and Genoa (averaged over sectors) in the month of the blitz: since the estimated effect is 31.1%, the counterfactual ratio (i.e., estimated in absence of the blitz) is $\hat{R} = \frac{R}{1+31.1\%}$. From such value and the observed payments in Genoa, we can easily calculate the counterfactual value of payments in Milan, and subtract it from the *observed* value, obtaining the given estimate.

significance of our result: if, for reasons discussed in Section II, the Milan blitz increased compliance by Genoa taxpayers, then our estimates are *downward* biased. Moreover, as we mentioned before, the North of Italy is estimated to have lower levels of evasion compared to the Center and to the South. If the repressive effect goes hand in hand with the level of evasion, our estimates will be conservative for what concerns the impact of public blitzes in other Italian cities.

It should be noted that the results presented concern the estimated increase of IVA payments only. If a shop evades IVA by reporting incorrect revenues, evasion of income taxes (due by the owner) is also occurring. The total effect of the blitz in terms of prevented tax evasion could hence be much larger.[24]

Several channels can possibly explain the publicity effect: shop owners could increase reported sales or decrease reported (presumably inflated) costs. Moreover, they may be doing so because of psychological motives (e.g., guilt or social stigma) or strategic decisions (higher perceived probability of auditing). Such distinctions are out of the scope of the present paper. The effect of media coverage on fiscal declarations has *per se* important policy implications.

Additional analyses

Robustness checks In order to check the robustness of our results regarding sector selection, we first rerun the one-sided MW test by removing one sector at a time, looking only at the other 17.[25] The effect always remains significant at the 5% level, as shown in Figure 4. It also remains significant ($p = 0.030$) if we remove sectors known to be involved in the blitzes (restaurants, coffeehouses/bars, discotheques, clothing shops), confirming

[24] If we consider the highest income tax rate (43%, which in Italy is applied on yearly incomes over € 75,000), the amount of prevented income tax evasion can be roughly estimated to lie between 50% and 200% of the prevented VAT evasion, depending on factors such as the markup level and whether input costs are also being incorrectly reported. For instance, if the incorrect reporting concerned exclusively sales, an IVA rate of 21% would correspond to a multiplicative coefficient of $100/21 \times 0.43 \approx 200\%$.

[25] Figure B.3 in Appendix B displays the same two panels of Figure 3 obtained when dropping one sector at a time.

Figure 4: p-value obtained by running the MW test omitting one sector at a time

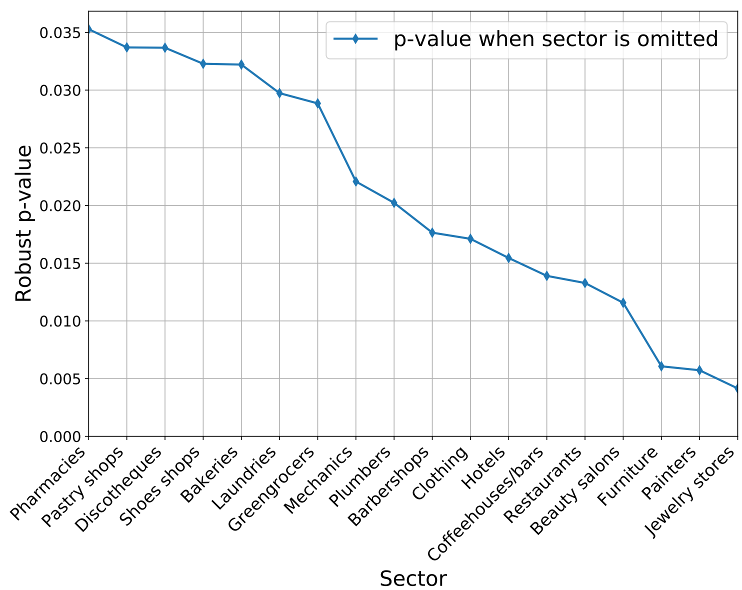

that publicity about the blitz reached all sectors, regardless of their direct involvement. The effect is even more significant ($p = 0.008$) if we drop from the analysis sectors in which the service provider goes to the customer, rather than the opposite. It is worth observing that such sectors (i.e., plumbers and painters, in our sample) pose specific challenges to tax verification authorities and are indeed expected to be less concerned by blitzes.

Second, we replace the pooled MW test with 18 independent MW tests, run on the residuals obtained from each sector separately. Note that each test is run on few observations only,[26] and hence will have low power; however, the (two-tailed) p-values obtained can be aggregated using Fisher's method Fisher, 1925 for meta-analysis.[27] This method guarantees that each sector is attributed the same importance, independently from the sector-specific variance, and results in a p-value of $p = 0.027$. This test, compared to the pooled MW test, should have higher power in the case of larger heterogeneity (in the distribution of residuals) among sectors, and would have much lower power instead if sectors

[26] More precisely, in each of those MW tests, the "treated" set only contains the residual for February 2012, and the resulting p-value is then simply r/T, where r is the rank of such element, and T is the number of observations available for that sector.

[27] Fisher's method must be applied to two-tailed p-values.

were identically distributed. The MW test is a combinatorial test, and performs worse if the sample is split in sub-samples for the analysis.

Third, given the importance of seasonal deadlines, we run placebo tests by assuming that the blitz took place in January of other years covered by the sample (2009, 2010, 2011, 2013) rather than in 2012. In no case can we reject the null hypothesis of equality between values in $\mathcal{B}$ and $\mathcal{C}$ ($p > 0.23$). Figure 5 presents the p-values resulting from placebo tests on *each* month in our sample. The only two months with p-values lower than the month of the blitz feature exceptionally low values of VAT payments in Genoa, as shown in Figure 3. Because January VAT payments are peculiar in that they only include the unanticipated component of December sales (see Section III), we also rerun the analysis while excluding all December payments or all January payments, obtaining p-values of 0.029 and 0.014, respectively.

As a fourth check, we adopt a parametric approach by regressing, through an OLS, the difference $\delta_{m,y,s}$ over sector-, month- and year-fixed effects, with an additional dummy variable "blitz" (again adopting a Difference-in-Differences identification strategy for identifying short-term effects). The coefficient for such variable has a positive sign, and the corresponding one-sided t test is significant ($p = 0.019$). If, rather than a single dummy variable, we add a set of sector-blitz interactions and we aggregate their two-sided p-values using Fisher's method, the result is still significant ($p = 0.019$). See Appendix B.3 for complete estimation results.

Finally, we reproduce our analysis on nominal (non-deflated) payments: differences in results are minimal (the p-value changes from 0.015 to 0.022 in the $\mathbb{ALL}$ specification and from 0.019 to 0.024 in the $\mathbb{PRE}$ specification). The unexplained component is on average 0.270, which results in a 31.0 % increase in the ratio between the two cities, in € $4,049$ of extra VAT payment per taxpayer, and hence € 7,669,639 in total. Adding the ratio of price indexes between the two cities as an additional regressor also results in minor changes (p-values of 0.030 in the $\mathbb{ALL}$ specification and 0.033 in the $\mathbb{PRE}$ specification).

Timing of blitzes While both blitzes of interest took place in the same month, the one in Genoa took place earlier and, in particular, before the 16th of January, which is the deadline for payments referring to December. If taxpayers waited until the very last minute to determine their VAT returns, rather than doing so as soon as the concerned month was over, then it would be possible that the Genoa blitz affected December 2011 VAT payments due in January 2012, rather than just January payments due in February. We hence use Milan as a counterfactual to check whether or not the Genoa blitz affected local VAT payments referring to December 2011, basically applying our strategy, reversed, to payments due in January 2012, but we do not find any effect ($p = 0.440$). We take this as evidence that the Genoa blitz either only affected January payments, which are due in February, or did not have any effect at all. While we believe that the latter explanation is the most convincing one (due to the Genoa blitz not being publicized), our empirical strategy is compatible with both options. As a further robustness check, we also verify that our main result still holds when removing January 2012 from the sample ($p = 0.029$).

Temporary vs medium term effects The $\mathbb{ALL}$ specification, in which the control sample includes months *after* February 2012, is coherent with the assumption that the blitz only had a temporary effect, or that, if a medium- or long-term effect is present, it cannot be observed due to confounding factors. Although we did verify that results are analogous in the $\mathbb{PRE}$ specification, it is still worth checking if the data in our observation window exhibit some long-term variation. For instance, if a change between the business cycles in Milan and Genoa *had* taken place at some time in our observation window, the results of the $\mathbb{ALL}$ specification could be biased by such a change. In particular, if the business cycle in the period after the blitz negatively affected Genoa compared to Milan, then our results would be biased upwards. Vice-versa, if the business cycle positively affected Genoa compared to Milan, then our results would be biased downwards. This risk is taken into account to a large extent by the presence of year-fixed effects in Equation

Figure 5: Results of placebo tests

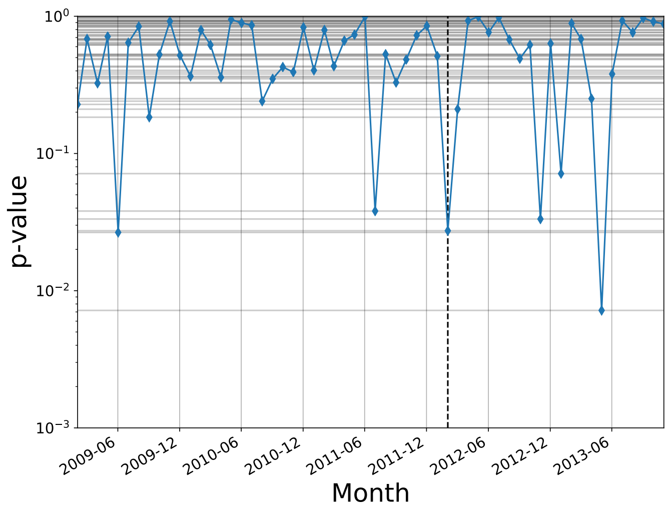

Note: Reported are the p-values of a series of placebo tests. Each of them consists of applying strategy $\mathbb{ALL}$ (see Section IV) assuming that the blitz took place in a different month. The selected month of each placebo is indicated on the x axis. The dashed line identifies the true month of the blitz.

(3); still, in order to completely neutralize it, we first run one-sided MW tests on the hypothesis that after the blitz there was a permanent increase/decrease in the unexplained terms. In both cases, we are unable to reject the null of no difference. Had we found a significant effect, its attribution to the public blitz would have been implausible anyway, for the aforementioned business cycle concerns.

We also check whether or not the publicity effect fades out gradually over the months after the blitz. Namely, we look at payments made in March 2012 (two months after the blitz) and find a positive effect on declarations (€ 3,601,279, roughly half the estimated effect for February), significant at the 10% level ($p = 0.081$).[28] Results are indistinguishable from zero starting with April 2012 (3 months after the blitz: $p = 0.784$). Fiscal data with higher temporal resolution (e.g., receipts released by shops are timestamped) would be needed to estimate more detailed response curves.

[28] This analysis excludes February 2012.

VI Conclusions

Recent empirical studies based on administrative data provide abundant evidence on the specific deterrence effect of audits (Slemrod, 2016). We contribute to the literature by analyzing the role of audit *publicity*, which can be a crucial instrument for sending signals to taxpayers about the willingness to fight tax evasion. In Italy, a fiscal blitz run in Milan in January 2012, achieved an exceptionally broad resonance as a consequence of an explicit choice of the Italian Revenue Agency to direct attention on it, e.g., by releasing information about it to the media. We study the effect of publicity by comparing VAT declarations of Milan taxpayers with those of taxpayers in Genoa, where another blitz was run in the same month without any publicity effort. We exploit confidential data on VAT declarations at the sector level provided by the Italian Revenue Agency; the strong heterogeneity across sectors is accounted for by the non-parametric approach employed. We find a positive, strong and robust short-term effect: VAT compliance is estimated to have increased by more than 7 million euro just in the month after the blitz, and across the 18 B2C sectors considered. Interestingly, the effect remains strongly significant even when discarding sectors directly involved in the blitz. This confirms our hypothesis that media coverage influences tax compliance.

Making a comprehensive assessment and comparison of the benefits and costs related to blitz publicity (e.g., the effort to coordinate audits with media outlets) is out of the scope of this paper, but for reference, we can present a back of the envelope calculation concerning the cost of *running* a blitz (as opposed to the cost of *publicizing* it). Although no public information is available on the number of tax auditors involved in the Milan blitz, we know that approximately 1,200 tax officers were employed and that the blitz lasted approximately 20 hours. Since, according to Revenue Agency estimates, the operational cost of each hour of audit is 55 euros per officer involved,[29] this means that, under the

[29] These figures, which we received privately from the Revenue Agency, are consistent with those provided to the OECD Tax Administration Database:

extreme assumption that each of the officers worked for the entire duration of the blitz, the total cost of the operation was around 1.3 million euros. That is, the upper bound to the cost of the blitz is approximately five times lower than the lower bound to the effect of publicity we provide. These numbers suggest that the net impact of a public blitz is positive.

However, the fact that blitz publicity was a sporadic initiative raises the question of what would happen if blitzes, and tax verification activities in general, routinely benefited from such media coverage. While, on the one hand, publicizing auditing activities more systematically could lead to a strengthening of social norms and habits concerning tax compliance in the long term, on the other hand it is at least plausible that the uniqueness of the publicity of the Milan blitz led to its success in terms of increased compliance. Moreover, it should be taken into account that some of the costs might have a *political* nature, i.e., influencing the popularity of politicians, and hence can be difficult to quantify; there is some evidence, however, that these costs are high. This is suggested by the fact that a change in strategy was publicly advocated for in the direction of strengthening the "mutual trust" between taxpayers and the Revenue Agency after the Milan blitz. In fact, the Italian Revenue Agency has not conducted public blitzes since then.

Appendix A A synthetic control exercise

We run a synthetic control estimation (comparing Milan with cities where *no blitz* took place) which we can interpret as a robustness exercise concerning the effect *of the blitz* (rather than of publicity) in Milan (notice that the effect of the private blitz in Genoa seems to be negligible).

Two problems arise: difficulty finding cities which are "similar to the one representing the case of interest" (Abadie et al., 2015) to be selected as potential donors, and difficulty

www.oecd.org/site/ctpfta/taxadministrationdatabase.htm.

selecting city-level predictors of the phenomenon of interest (VAT evasion) to be used in the determination of the weights attributed to the different donors. To the extent that donors are not similar to the treated unit, the possibility of "interference between units" Abadie et al. (2010) also arises.

Our data were provided by the Italian Revenue Agency per our request to study the Milan blitz; hence, they only include a handful of cities (our request for data reported the exact list of cities and sectors needed). However, we obtained data on VAT declarations for the provinces (NAT-2 subdivisions) in Lombardia, the region of which Milan is the administrative center. We can hence run a synthetic control exercise by comparing the province of Milan[30] with a donor pool including the other 11 provinces.

In light of the difficulty of selecting valid predictors for the outcome variable, we follow two different specifications. In one case ("static" specification), we resort to the following province level variables: percentage of working population; per capita income; population density; per-capita number of VAT declarations in each of the sectors of interest, for a total of 21 province-level indicators. The three general indicators are sourced from the Italian National Institute of Statistics (ISTAT) data; sector-level indicators are obtained by dividing the average number of monthly VAT declarations in our database by the province population. In the other case ("dynamic" specification), we use as descriptors the per-taxpayer VAT declarations for each sector (i.e., the province-level equivalent of $log(M_{m,y,s})$ in Section IV), for each month until January 2012, obtaining a more directly related and granular albeit maybe less informationally rich, set of descriptors. In both specifications, the weight attributed to each variable is determined endogenously as a result of the optimization (as done by Abadie and Gardeazabal, 2003), the target of which is the minimization of the distance between VAT declarations (averaged across sectors) of the Milan province, and of the synthetic province, in the pre-treatment period.

[30] It should be noted that Milan and its urban area account for the vast majority of the population, businesses and local media of the Milan province.

Figure A.1: History of reported VAT payments - Milan compared with synthetic control

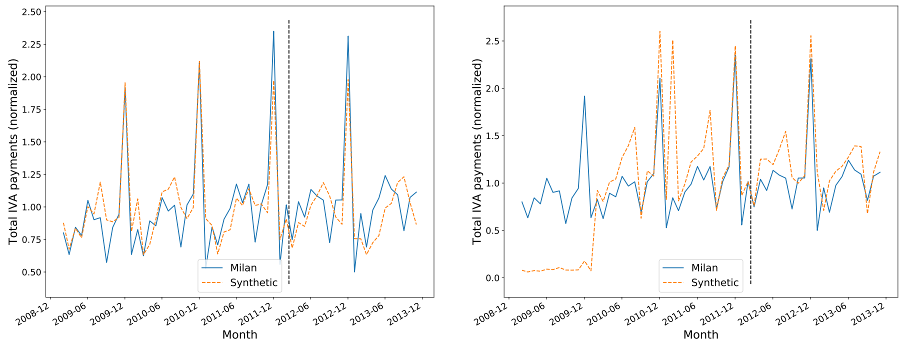

Note: Total VAT payments, across the 18 sectors analyzed. For ease of comparison, both series are normalized so that they average to 1. Left: static specification (Brescia); right: dynamic specification ("Monza e Brianza" and Brescia). Data for "Monza e Brianza" is missing in the fist part of the time span, resulting in the low values which can be observed in the figure. For such periods, the calibration process automatically excluded "Monza e Brianza". The dashed line identifies the month of the blitz.

The loading weights for each province in the donor pool are calculated with the "Synth" package for R. In the static specification, the resulting synthetic province was composed uniquely (weight 0.9999) by the province of Brescia. In the second case, it was composed overwhelmingly (weight 0.9357) by the province of Monza e Brianza, but with a non-negligible weight (0.0643) still attributed to the province of Brescia. See Figure A.1 for the equivalent of Figure 2 comparing Milan with the resulting synthetic control. Running the significance analysis described in Section IV results in a p-value of 0.011 when comparing the province of Milan with the province of Brescia (static specification), and of 0.009 when comparing it with the synthetic control resulting from the dynamic specification, in line with results from our main approach.

References

Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association 105*(490), 493–505.

Abadie, A., A. Diamond, and J. Hainmueller (2015). Comparative politics and the synthetic control method. *American Journal of Political Science 59*(2), 495–510.

Abadie, A. and J. Gardeazabal (2003). The economic costs of conflict: A case study of the Basque Country. *The American Economic Review 93*(1), 113–132.

Allingham, M. and A. Sandmo (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics 1*(3-4), 323–38.

Alm, J., B. R. Jackson, and M. McKee (2009). Getting the word out: Enforcement dissemination and compliance behavior. *Journal of Public Economics 93*(1), 392–402.

Battiston, P. and S. Gamba (2016). The impact of social pressure on tax compliance: A field experiment. *International Review of Law and Economics 46*, 78–85.

Bertrand, M., E. Duflo, and S. Mullainathan (2004). How Much Should We Trust Differences-In-Differences Estimates? *The Quarterly Journal of Economics 119*(1), 249–275.

Bethmann, D. and M. Kvasnicka (2016). International Tax Evasion, State Purchases of Confidential Bank Data and Voluntary Disclosures. *Institute of Economic Research, Korea University, Working Paper Series* (1603).

Bø, E. E., J. B. Slemrod, and T. O. Thoresen (2015). Taxes on the Internet: Deterrence effects of public disclosure. *American Economic Journal: Economic Policy 7*(1), 36–62.

Buehn, A. and F. Schneider (2016). Size and Development of Tax Evasion in 38 OECD Countries: What Do We (Not) Know? *Journal of Economics and Political Economy 3*(1).

CASE (2016). Study and Reports on the VAT Gap in the EU-28 Member States: 2016 Final Report. Technical report, Center for Social and Economic Research, Warsaw. https://ec.europa.eu/taxation_customs/sites/taxation/files/2016-09_vat-gap-report_final.pdf.

D'Agosto, E., M. Marigliani, and S. Pisani (2014). Asymmetries in the territorial VAT gap. Discussion Topics 2, Italian Revenue Agency.

Doerrenberg, P. and D. Duncan (2014). Experimental evidence on the relationship between tax evasion opportunities and labor supply. *European Economic Review 68*, 48–70.

Fabbri, M. and S. Hemels (2013). 'Do You Want a Receipt?' Combating VAT and RST Evasion with Lottery Tickets. *Intertax 41*(8/9), 430–443.

Fisher, R. A. (1925). *Statistical methods for research workers*. Oliver and Boyd, Edinburgh.

Hashimzade, N., G. D. Myles, and T. N. Binh (2013). Applications of Behavioural Economics to Tax Evasion. *Journal of Economic Surveys 27*(5), 941–977.

Italian Government (2013). Rapporto concernente i risultati conseguiti in materia di contrasto dell'evasione fiscale (report concerning results obtained in curbing tax evasion). Nota di Aggiornamento al DEF, Allegato II (Update to the Economic and Financial Planning Document, Attachment II), Senato della Repubblica (Republic Senate).

Kasper, M., C. Kogler, and E. Kirchler (2015). Tax policy and the news: An empirical analysis of taxpayers' perceptions of tax-related media coverage and its impact on tax compliance. *Journal of Behavioral and Experimental Economics 54*, 58–63.

Kleven, H., M. Knudsen, C. Kreiner, S. Pedersen, and E. Saez (2011). Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark. *Econometrica 79*(3), 651–692.

Konrad, K. A., T. Lohse, and S. Qari (2017). Compliance with Endogenous Audit Probabilities. *The Scandinavian Journal of Economics 119*(3), 821–850.

Kopczuk, W., J. Marion, E. Muehlegger, and J. B. Slemrod (2016). Does Tax-Collection Invariance Hold? Evasion and the Pass-through of State Diesel Taxes. *American Economic Journal: Economic Policy 8*(2), 1–36.

Langenmayr, D. (2017). Voluntary disclosure of evaded taxes – Increasing revenue, or increasing incentives to evade? *Journal of Public Economics* (151), 110–125.

Marino, R. and R. Zizza (2012). Personal Income Tax Evasion in Italy: An Estimate by Taxpayer Type. In M. Pickhardt and A. Prinz (Eds.), *Tax Evasion and the Shadow Economy*. Elgar.

Pisani, S. (2014). An approach to assess how the activity of the Italian Revenue Agency affects compliance. Discussion Topics 1, Italian Revenue Agency.

Pomeranz, D. (2015). No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax. *American Economic Review 105*(8), 2539–69.

Rincke, J. and C. Traxler (2011). Enforcement Spillovers. *Review of Economics and Statistics 93*(4), 1224–1234.

Slemrod, J. B. (2016). Tax Compliance and Enforcement. New Research and its Policy Implications. *SSRN Working Paper 2726077*.

Slemrod, J. B., M. Blumenthal, and C. Christian (2001). Taxpayer response to an increased probability of audit: evidence from a controlled experiment in Minnesota. *Journal of Public Economics 79*(3), 455–483.

Yitzhaki, S. (1974). A note on income tax evasion: A theoretical analysis. *Journal of Public Economics 3*(2), 201–202.