Unfinished Business: Climate-driven Contract Suspensions as Firm Liquidity Shocks

Luigi Dante Gaviano*

August 2025

Abstract

Climate change has disruptive effects on economic activity. This paper focuses on one channel through which this disruption materialises: contract suspension. If adverse weather conditions force the interruption of work, contractual payments to the firms assigned to the project are delayed. The delay imposes several costs on these firms. They experience lower liquidity due to deferred payments, coupled with uncertainty regarding the duration of the suspension. Concurrently, firms cannot fully reduce their labour and capital costs, since they may have to suddenly resume work when the suspension is lifted. I quantify these costs using data on Italian construction firms and public infrastructure projects, obtained from a new database on the universe of public procurement contracts in the country. The suspension channel is isolated with a staggered DiD design matching similar firms. Suspensions lead to extensive financial damages, with sales dropping on average by 30%, employment by 15.3%, and total assets by 18.5% in the years after a firm's first suspension. This contraction in firm operations arises both from the adverse liquidity effects of weather suspensions, and from their knock-on effects on firms' other contracts, which are also hit by delays.

*University of Cambridge, Trinity College (ldg26@cam.ac.uk). I am grateful to Elisa Faraglia, Tiago Cavalcanti, Constantine Yannelis, Aksel Mjøs, Ambrogio Cesa-Bianchi, Ludovica Gazzè (discussant) and participants of the Bank of England Macro Brown Bag Seminar, 2025 EAERE Annual Conference, 2025 FMA European Conference, Spring 2025 Midwest Macro Meeting, Fourth SaNE Conference, 2025 ifo Institute Climate Adaptation Workshop, various University of Cambridge seminars, and 2024 Austrian Economic Association Meeting for their invaluable comments. I am indebted to Maria Rosa Simonetti for her assistance with ANAC data.

Keywords: climate change, firm dynamics, liquidity risk, construction sector

JEL codes: D22, G32, L25, L74, Q54

1 Introduction

Weather shocks are rising in frequency and intensity due to climate change. In addition to physical damages to infrastructure and to communities, adverse weather can also disrupt economic activity. In particularly weather-exposed sectors, such as construction, weather shocks may force the interruption of a contract when advancing work becomes technically impossible or unsafe for workers. As public contracts are suspended due to adverse climate events, payments from the public sector to the firms assigned to the project (affected firms) can be halted, thereby imposing heavy liquidity costs on affected firms. My goal in this paper is to quantify firms' costs arising from weather-driven contract interruptions.

To address this question, I use data on Italian public contracts in construction (public works) due to several reasons. Firstly, construction is a large sector, the gross value added of which amounted to 18.7% of Italian GDP in 2018 (*European Construction Sector Observatory*, June 2018). It is also directly exposed to weather events, since most of its activities take place outdoors, meaning that weather-driven disruptions in this sector should leave visible traces (Schuldt et al., 2021). Furthermore, the construction sector is characterised by great network centrality, given its large raw material and labour requirements and its role in facilitating connectivity between sectors, e.g. through transport infrastructure (Ive and Gruneberg, 2000). Public investment in infrastructure is also sizeable and matters for aggregate fluctuations (Basso et al., 2024).

Secondly, I focus on public, rather than private, contracts for reasons of data granularity and comparability. In order to ensure that Italian public funds are transparently accounted for, the country's anti-corruption agency (*ANAC*) has launched a *National Database of Public Contracts* (BDNCP), which covers the universe of procurement contracts for the entire public sector (ranging from schools to hospitals and ministries). Importantly, the data reported in the BDNCP is publicly available, contract-level, and harmonised, ensuring comparability over time, across regions and between types of public entities.

Thirdly, the Italian setting is particularly suited to my research question for two further reasons. On the meteorological front, the country is regularly hit by a range of weather and climate-related shocks, including floods, landslides and heat waves (Mysiak et al., 2018), which suggests that such shocks should be a tangible concern in weather-exposed sectors like construction. Indeed, BDNCP data suggests that around 8% of Italian public infrastructure projects by contract value - amounting to €28.8bn - have been suspended due to weather since 1999. On the legal front, the implications of a weather-related suspension for affected firms are very clearly codified in Italian law - namely, the interruption of any cash flows from the public entity to the firm for the duration of the suspension. This demand shortfall is the key channel through which this type of contract disruption should negatively affect firms, and thus precisely identifying it is crucial.

By merging the BDNCP data on public contracts with yearly firm balance sheet and income statement data from *Orbis Historical*, I can measure the economic repercussions of contract suspensions on firms, chiefly on their size and financial viability.

I achieve this through a matched difference-in-differences (DiD) research design, in which I compare constructions firms that have received a contract suspension for the first time with firms of a similar size, sector and in a similar location that have not yet received a suspension, but that will in the future (the *not yet treated*). I argue that this pairing of treated and control groups allows me to net out any unobserved endogenous determinants of weather suspensions that might be correlated with firm outcomes, like a firm's technology level or political connections.

In my preferred specification, I find that weather suspensions reduce firm sales on average by around 30%, total assets - a proxy of firm size - by 18.5%, and the number of employees by 15%. These substantial drops in the operations of affected firms persist over 3-4 years.

As some of a firm's contracts are suspended, labour and machinery can only be reallocated imperfectly, and liquidity constraints may bind. I show that this gives rise to delays across the firm's unsuspended contracts, as the production disruption spills over to the rest of the firm's operations. This explains the vast negative effects of suspensions on firms.

Overall, I interpret these results as being suggestive of limited insurance coverage of Italian construction firms. It seems that this specific risk to their operations - that of a cash flow freeze arising from weather-driven contract interruption, which need not also be associated with physical damage to the firm's assets - is materialising in my sample, with severe repercussions for the financial viability of affected firms.

These results have great relevance beyond the firms narrowly affected by contract interruptions. The suspension channel of climate change is likely to transmit along the production network, as the subcontractors and suppliers of treated firms will experience sharp drops for the demands of their own products and services. Furthermore, in the absence of improved insurance package design and take-up, or of other risk-sharing agreements between construction firms and their clients, the probability of contract interruption is likely to lead to costlier, longer construction projects being drawn up in the first place. In the case of public procurement, this amounts to onerous demands on taxpayers' money.

The remainder of this paper is structured as follows. Section 2 reviews the related strands of the literature on weather shocks, public procurement, and their impact on firm dynamics. Section 3 examines the notion of public contract suspensions, and clarifies what the associated costs for firms are. Section 4 introduces the two data sources used in this project. Section 5 outlines my matched DiD approach. Section 6 presents my results, arranged into liquidity effects of suspensions on firms and into knock-on effects on these firms' other contracts. In section 7 I draw out the broader implications of my results, and in section 8 I conclude. Finally, in the appendix (section 9) I provide further details on how my dataset is constructed, and outline further findings.

2 Literature

This project speaks most directly to three strands of the literature. The first is concerned with estimating the total impact of weather and climate shocks on firms, typically without seeking to identify the specific economic or financial channels through which this impact materialises (Grover and Kahn, 2024). Instead, headquarters' location is frequently used as a proxy for a firm's climate exposure. For instance, using data on over 3 million listed and non-listed firms in 24 countries, Cevik and Miryugin (2022) report that firms operating in countries that are more vulnerable to climate change exhibit higher leverage, interest burden, and lower profitability and TFP. Ponticelli et al. (2023) find that temperature shocks raise energy costs and lower the productivity of small US manufacturing plants, though no significant negative effects on employment or positive effects on firm exit are detected. Elliott et al. (2019) report large but short-lived negative effects of typhoons on Chinese manufacturing plants, particularly on liquidity, turnover and profits. Yu and Shi (2024) analyse how high temperatures cause delays and cancellations in the US airline industry, with substantial costs materialising for affected firms. Caggese et al. (2024) rely on Italian firm-level financials and weather data, and find negative effects of extreme temperatures on firm sales, material and labour costs, and productivity. Their general equilibrium model then allows them to aggregate these productivity losses in order to evaluate economic damages under different projected climate change scenarios.

The effects of hydrogeological events have also been examined: Clò et al. (2024) focus on episodes of floods and landslides in Italy since 2010, and they find a negative effect on survival probability, revenues and employment. These effects are particularly strong among smaller firms, and in the services and construction sectors. More broadly, Fatica et al. (2022) consider the impact of floods on European manufacturing firms. They report negative effects on firm total assets, sales and employment up to 7-8 years after a flood event. Conteduca and Panon (2024) study the effect of natural disasters on firm markups in Italy and France. Furthermore, some have used local surveys in the aftermath of extreme events like floods in order to gauge the extent of reported damages and firm responses (Endendijk et al., 2024; Sakai and Yao, 2023; Sultana et al., 2018).

Others have also analysed the firm-level impact of weather and climate events on the extensive margin (firm entry, continuation and exit). In this vein, Cascarano et al. (2023) rely on administrative data on the quasi-universe of Italian firms to quantify the effect of protracted high temperature episodes on firm demography. They document that these weather shocks persistently reduce the entry rates of new firms, and raise the exit rates of incumbent ones, thereby reducing business dynamism. Turning to financial performance, Pankratz et al. (2023) examine the universe of firms in S&P's Compustat Global database (17,000 firms in 93 countries), and find that investors do not correctly anticipate the adverse effects of heat shocks: firms exhibit more negative revenue and operating income surprises and announcement returns as firm heat exposure rises, although the exposure to temperature shocks could in principle be directly observed and taken into account by market participants.

Custódio et al. (2022) isolate the *supply-side* effects of weather shocks on firms, i.e. those that involve changes to labour supply, productivity and costs of production. Custódio et al. (2022) focus on the variation in local temperatures to which suppliers of the same firm are exposed, which allows them to control for firm-specific demand, since they observe supplier-client pair sales in the production network. They find that both higher temperatures and extreme weather events lower firm sales through this supply-side channel. Furthermore, these drops in sales are stronger in labour-intensive industries and in financially constrained firms, which suggests that lower labour productivity and more limited financial flexibility might be important mechanisms through which supply-side costs of weather shocks materialise.

A second strand of the literature addresses the adaptive responses and mitigation strategies put in place by firms affected by climate events (Grover and Kahn, 2024). This literature typically emphasises that the effects of natural disasters on firm outcomes are shaped by the associated supply chain disruptions and production process turbulence (Carvalho et al., 2021). Acharya et al. (2023) use US establishment-level data to examine how climate disasters affect employment, as mediated by firm climate risk management. They show that multi-location firms respond to heat shocks by reallocating employment away from locations hit by shocks into non-shocked ones. Conversely, single-location firms merely cut employment. This geographic shift in employment aims to contain losses in labour productivity due to higher temperatures, and is costly. Indeed, Acharya et al. (2023) show that this geographic reorganisation of labour is larger among bigger, lower leveraged firms with greater ESG awareness among their investors. A similar mechanism is also at play in Castro-Vincenzi (2024), who focuses on car production plants and flood episodes around the globe. Castro-Vincenzi (2024) finds that multi-plant firms reallocate production to plants spared by floods and that produce similar car models to affected plants. This is consistent with Ponticelli et al. (2023), who conclude that large production plants are better able to adapt to temperature shocks than smaller ones, and that there is evidence of some labour reallocation from the latter to the former, in turn raising the concentration of manufacturing activity among larger firms.

Turning to firm adaptation in a developing country that is particularly vulnerable to hydrogeological events, Balboni et al. (2023) use granular data on firm-to-firm transactions, firm locations, and data on geo-located commercial truck journeys in flood affected areas to establish a large range of firm adaptation methods to floods in Pakistan. Affected firms tend to relocate to less vulnerable areas. They also adapt their supply chains to reduce their indirect flood exposure, diversifying their suppliers and turning more to suppliers based in less flood-prone areas and reachable via less flood-prone roads. Echoes of these results can also be found in Pankratz and Schiller (2024) and Castro-Vincenzi et al. (2024). Pankratz and Schiller (2024) leverage data on the global supply chains of large companies to show that firms are more likely to terminate their relationships with a supplier if the supplier's realised exposure to heat is greater than that expected ex ante. They also provide evidence of firms switching to substitute suppliers with a lower climate exposure. Castro-Vincenzi et al. (2024) employ firm-to-firm data from a large Indian state to examine the effect of climate risk on transactions between affected suppliers and downstream buyers. Castro-Vincenzi et al. (2024) show that firms diversify the locations from which they purchase inputs as part of a risk mitigation strategy. Moreover, in their event study design evidence is provided for short-lived but sizeable drops in the sales of flood-affected suppliers, and in the purchases of their downstream buyers.

Using World Management Survey data, Keiller and Van Reenen (2023) show that good management practices attenuate the negative effects of natural disasters on firms. In particular, well-managed firms have a lower post-disaster exit rate, and exhibit smaller drops in employment and output. They argue that this good management premium arises due to a better perception of firm-specific climate-related risks, and due to already-implemented risk-reducing measures in well-managed firms. Relatedly, Downey et al. (2023) establish that although heavy precipitation is costly for construction firms, the earlier rainfall is correctly forecasted, the greater is firms' ability to reallocate labour, thereby minimising the reduction in worker productivity, and in turn cushioning any profit losses arising from (anticipated) weather disturbances.

Taken together, these results also point to the importance of adaptation in mitigating the adverse firm-level effects of climate change. Interestingly, Addoum et al. (2020) find that US publicly listed firms exhibit no statistically significant effects of temperature shocks on their sales, productivity and profits, whether at the firm or establishment level, and whether in heat-exposed or non-heat-exposed sectors. This is suggestive of how high the upper bound of the adaptation potential of firms is, particularly if they are sufficiently large and financially unconstrained.

Thirdly and finally, my specific empirical setting - public procurement - raises the question of how public contracts affect firm dynamics in the first place. Using specific quasi-randomising features of Brazilian electronic procurement actions, including their randomised end time and the full distribution of bid amounts and bidding times, Ferraz et al. (2015) measure the benefits of securing public contracts. These lead to positive and persistent effects on firm employment growth, lasting at least two years, and which are stronger for younger firms (controlling for firm size). Over the medium-term, close auction winners participate in more auctions, and are more likely to win again, compared to close losers, and they also tend to diversify their customer base both geographically and in terms of product type.

There are also sizeable financial benefits from winning a public auction. Using administrative data on Portuguese public procurement contracts, coupled with balance sheet and credit registry information for procurement firms, Gabriel (2024) shows that being awarded a public contract increases firm credit while lowering the associated cost of credit. Winning a procurement contract therefore increases the net worth of a firm by boosting its future cash flows, thereby alleviating borrowing constraints and improving the firm's liquidity position. Ultimately, this channel has a positive effect on investment and employment by the firm.

Cappelletti et al. (2024) focus on the positive effect of public contracts on firm survival. Using data on Italian public procurement, they find that securing a public contract increases the survival probability of the firm over the next 2-3 year period. Concurrently, they show that winning firms substitute private with public customers (earnings substitution), rather than growing their customer base. However, the benefits of winning a public contract do materialise in terms of a rise in creditworthiness and in the volume of credit granted through uncollateralised loans, which they link to the superior survival prospects of winning firms. This loosening of firm credit constraints occurring thanks to public contracts is consistent with the broader literature: using Spanish procurement data and a model to quantify macroeconomic effects, di Giovanni et al. (2023) show that winning a public contract loosens borrowing constraints that are both earnings-based (on impact) and asset-based, i.e. that require capital as collateral, later on. Naturally, to the extent that weather suspensions disrupt public contracts, one would expect them to partially vitiate these credit and liquidity benefits for firms. This is exactly the mechanism of interest in this project.

3 Implications of weather suspensions for firms

The engineering and operations' research is replete with revealing examples of how weather conditions may force the interruption of a construction project (Schuldt et al., 2021). Subzero temperatures create difficulties for machinery usage, e.g. as frost penetration in the soil makes it harder to excavate and then refill terrain. Strong winds may create hazardous environments for workers, particularly due to flying debris, while making it dangerous to rely on tower cranes for lifting in construction sites. Even light precipitation can disrupt steps in the construction process, chiefly by negatively affecting the placement of materials. For instance, concrete may be placed only when the water-cement ratio is in a specific range of values, and is thus vulnerable to excess or insufficient humidity. Asphalt paving also usually has to be interrupted whenever a risk of rain, snow or hail is detected. Heavy rainfall also requires time being spent shielding sensitive materials or areas from water, especially when digging, while raising the risk of the construction site being flooded and of some haul routes becoming inaccessible.

In all of these scenarios, construction work would presumably have to be halted. Specifically in the setting of Italian procurement contracts, these weather suspensions are legally defined as *force majeure* contract interruptions due to "adverse climatic conditions". The implications for contracting firms are explicitly detailed in legislation, having undergone almost no change over 1999-2023.[1] Crucially, weather suspensions need not arise only when there is

[1] This section is directly based on the relevant Italian laws and decrees, namely: DPR 21 December 1999,

physical damage to the firm or to the actual construction site, nor necessarily when it is physically impossible to advance with the building work. Instead, the phenomenon is much broader: any circumstance in which work cannot be carried out "to a satisfactory standard" gives rise to a suspension, e.g. if workers' health might be endangered due to being exposed to severe heat, or if extreme temperatures mean that some building materials have decomposed and are thus no longer safely usable for construction.

Furthermore, the law requires firms to already price in average weather conditions when submitting bids for the contract, by factoring these into their cost and completion time calculations. Thus, contract suspensions are granted only when the adverse weather event is unexpected given local seasonal patterns, and could not be foreseen when drafting the public contract.

The beginning and end of a weather suspension are envisaged as being automatic, based on the objective occurrence and continuation of the underlying climatic event, and are thus not subject to bargaining between the contract parties. The ultimate duration of a suspension is thus uncertain, being tied directly to the continuation of the underlying weather shocks, which is ex ante unobserved. In other words, the value of the deferred contractual payments due to the interruption is subject to uncertainty, which is itself a cost for firms.

More practically for contracting firms, legislation explicitly forbids any work on the interrupted project during a contract suspension. At the same time, any payments from the public entity to the firm are similarly forbidden, and no indemnity or compensation for the firm arising from the suspension is contemplated. Thus, the way in which public contracts are regulated in Italy places all contractual risk squarely on the firm, and none on the public sector.

In addition to eliminating any cash flows arising from suspended contracts, the law also

n. 554, art. 133, DM 19 April 2000, n. 145, art. 24, DPR 5 October 2010, n. 207, art. 158-160, DL 18 April 2016, n. 50, art. 107, DL 31 March 2023, n. 36, art. 121.

makes clear that the firm continues to be bound by her contractual obligations during a suspension. In practice, this means that the firm can reduce her fixed costs, e.g. by renting out her machinery, placing her workers under short-term work schemes or firing them, only to a very limited extent. The firm must remain able to promptly resume work once the suspension period comes to an end. Furthermore, the firm is obliged to keep the construction site open and secure, while preserving its ability to function, which entails significant costs. The firm's ability to reallocate of factors of production to an unsuspended construction site is thus sharply curtailed.

To summarise, the *suspension channel* through which climate shocks disrupt the economic activity of firms operates via the following mechanisms:

- **Liquidity costs** due to the lack of payments from suspended contracts;
- **Uncertainty** over these liquidity costs due to ex ante unclear suspension duration;
- **Fixed costs** persistence, as contractual obligations continue despite the suspension.

All three mechanisms will hinder firms in their ability to reallocate workers and machinery to other construction projects, both logistically, as contract work may need to be resumed suddenly, and financially, as this reallocation may require cash. It is thus plausible that unsuspended contracts too might end up being delayed and disrupted - this is a form of knock-on effect of weather suspensions I will examine later. Overall, these mechanisms will materialise as a negative demand shock for affected firms, with major implications for their economic position and financial viability.

4 Data

In this project I merge two granular datasets: the *National Database of Public Contracts* (BDNCP), in order to have contract-year data on the universe of Italian public procurement contracts, and *Orbis Historical*, to obtain firm-year balance sheet and income statement data for a large and representative sample of firms (both publicly listed and private). Each is now described in turn.

4.1 Data on public contracts

Italy's anti-corruption agency (*ANAC*) collects data on public procurement contracts in its *National Database of Public Contracts* (BDNCP). The entire public sector, across all levels (local, provincial, regional and national), is obliged to supply data on all of their procurement activities to *ANAC*, who then clean it, harmonise it and make it publicly available as part of the BDNCP. This is thus a rich database both in terms of coverage and of the details of the public contracts covered, including data on:

- Procuring public entities;
- Procurement auctions and bids;
- Winning firms and sub-contractors;
- Contract description;
- Contract start and end;
- Expected and realised costs;
- Payments to the firm by *tranche*;

- Contract modifications, incl. early terminations;
- **Contract suspensions**.

The database splits public procurement contracts into three categories: procurement of supplies, public works and services (figure 1). In this paper I will focus on public works only, which consist of all construction, maintenance and repair work of public infrastructure. Figure 1 shows that although this contract type underwent a sharp decline in 2011-2016, largely as a result of cuts in public investment in the aftermath of the euro area sovereign debt crisis, spending on ongoing contracts has largely stabilised since then to around €10-15bn a year, which remains a sizeable amount.

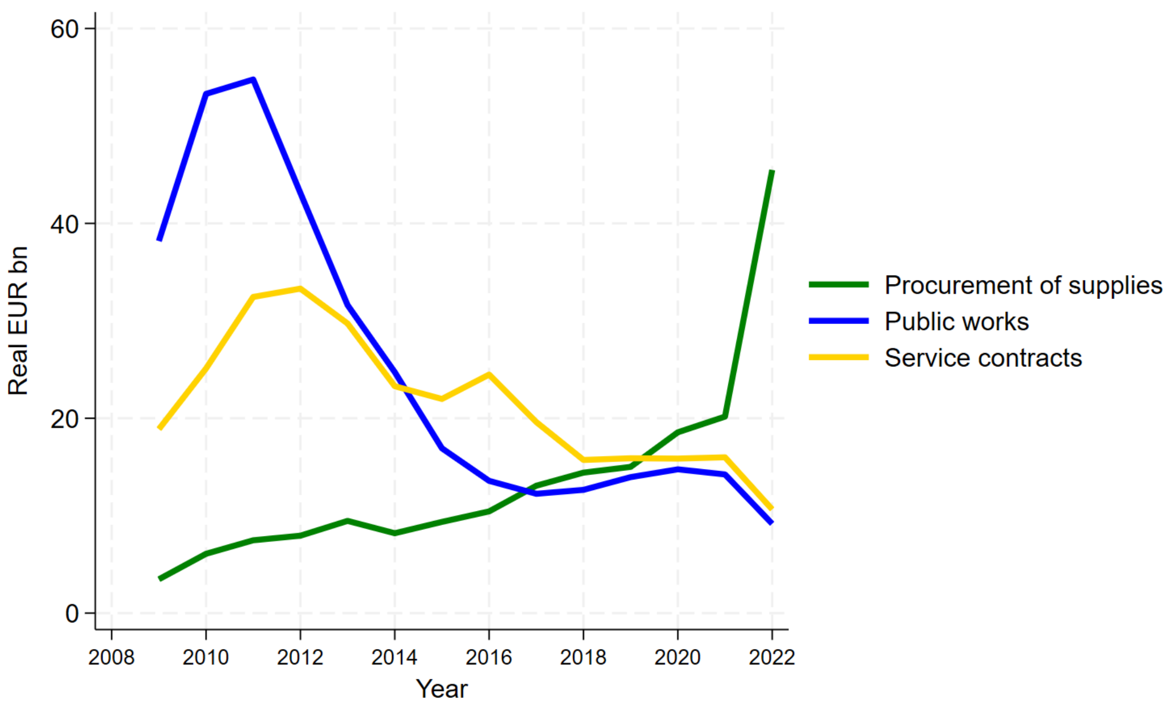

Figure 1: Italian public procurement in the long run
Source: BDNCP

The BDNCP dataset of greatest interest here is that on suspensions, i.e. public contract interruptions.[2] Seven suspension types appear in the dataset, with varying quantitative

[2] A small but growing literature focuses on closely related datasets on Italian public procurement, exploring

importance. Adding up the value of public contracts suspended by type over 1999-2023, with €30bn weather suspensions (*adverse climatic conditions*) emerge as the second most substantial suspension motivation (table 4.1).

Suspension types	**Value of contracts**
Variations to the contract	€36.4bn
Adverse climatic conditions	**€28.8bn**
Interference of technical nature	€20.5bn
Force majeure cause	€18.3bn
Interference of administrative nature	€16.4bn
Public interest or necessity	€11.7bn
Intervention of the judiciary	€0.878bn

Table 4.1: Total value of contracts suspended by suspension type, 1999-2023

The six non-weather suspension types are highly varied in impact and nature. While some are clearly brought about by misdeeds of public sector officials and/or of the firm herself (*Intervention of the judiciary*), others are less neatly classifiable, and might potentially include weather-related disruptions (*Force majeure cause*, *Public interest or necessity*). In the absence of a granular description of the reason for the suspension (only this typology is reported), in this paper I exclusively use data on *adverse climatic conditions* (weather) suspensions. Around one third of suspension periods are due to weather (23k out of 68k); at the same time, these tend to be relatively short in duration compared to non-weather ones (table 4.2). Nevertheless, the median weather suspension period lasts for 45 days, with the top quartile of the distribution lasting well in excess of three months, which is potentially suggestive of large liquidity costs to affected firms.

To observe the geographic spread of public contract interruptions, in figures 2 and 3 I plot the share of public works' contracts suspended over 1999-2023, overall and then due to weather only (respectively). Interestingly, the provinces with the highest weather suspension

contract rather than firm outcomes, e.g. the efficiency and discretion in the selection of suppliers (Baltrunaite et al., 2021), and the impact of procurement management quality on the speed of contract awarding and project completion (Baltrunaite et al., 2023).

Table 4.2: Distribution of suspension durations in days

	N	25th P.	Median	75th P.
Weather suspensions	23,058	20	45	100
Non-weather suspensions	45,227	42	98	200
Total	68,285	31	76	165

incidence (9-15% of public works contracts) are well-distributed over the whole peninsula, likely reflecting the prevalence of different types of weather events in different regions, e.g. frost in the alpine provinces (*Trentino Alto Adige*), floods and landslides in the central ones (*Toscana, Marche*), and heat waves in the southern ones (*Calabria*).

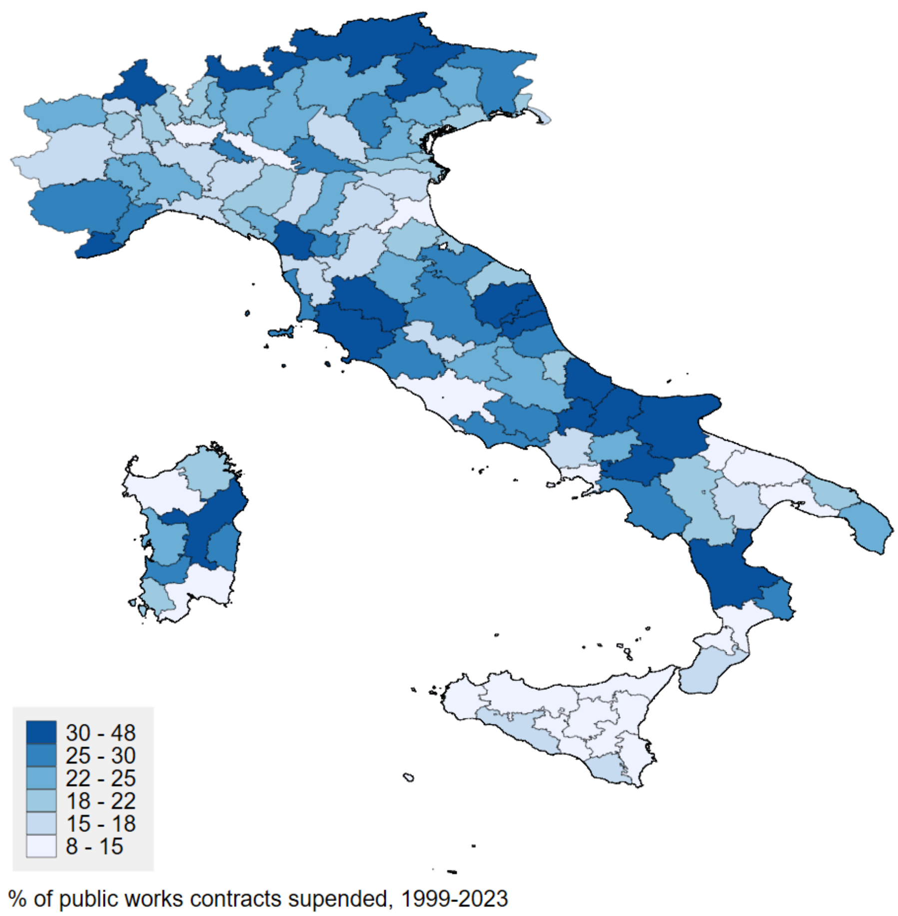

Figure 2: All suspensions as a share of public contracts, by province
Source: BDNCP

To examine the correlation between meteorological conditions and weather suspensions, I

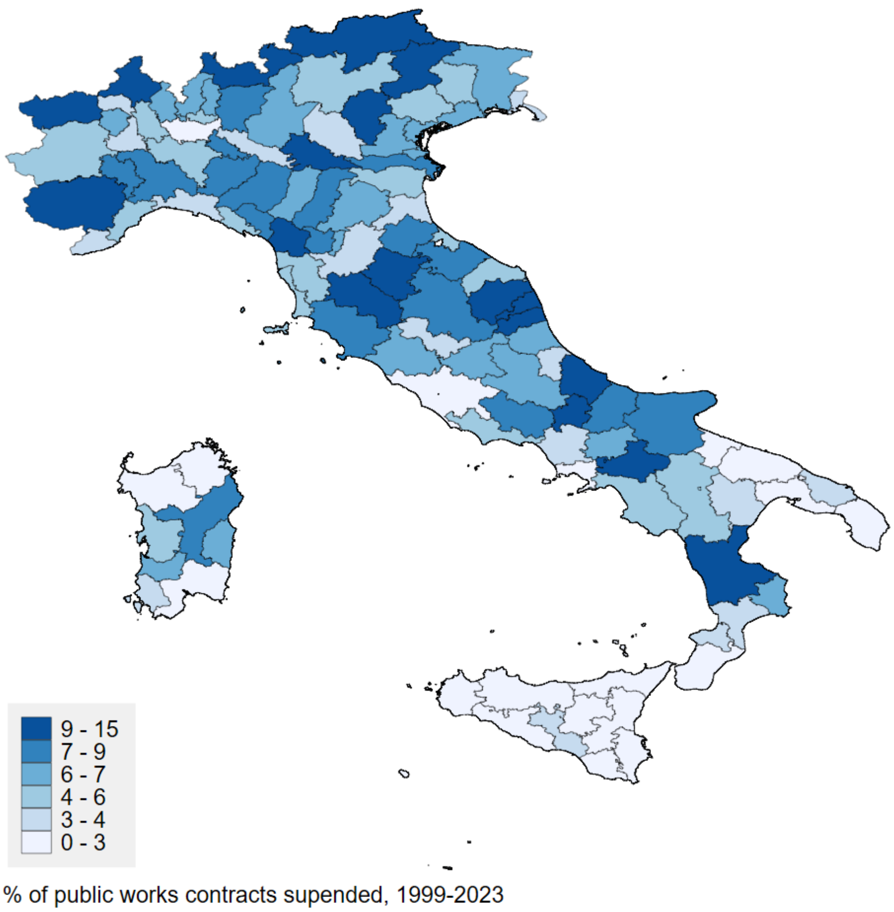

Figure 3: Weather suspensions as a share of public contracts, by province
Source: BDNCP

assemble weather data and build a city-level panel spanning 2001-2022 at monthly frequency, covering the 3,735 Italian cities experiencing contract suspensions.[3] I then estimate the following linear probability model to predict the beginning of a suspension episode:

$$susp_{ct} = \beta_0 + \beta_{\mathbf{1}}\mathbf{weather_{ct}} + \omega_{\mathbf{FE}} + v_{ct} \tag{1}$$

where $susp$ is a dummy equal to one whenever in city c at date t at least one weather suspension begins, $weather$ is a vector of different temperature and precipitation variables, and ω_{FE} refers to city, year and month fixed effects.

Tables 4.3 and 4.4 report the results of this predictive exercise. Firstly, table 4.3 indicates that higher precipitation and lower temperatures are associated with a higher probability of weather suspensions. By discretising the precipitation and temperature variables into four bins, as defined by the quartiles, it becomes clear that this relationship holds over the entire weather distribution, though with a lower intensity (coefficient magnitude) when there is little precipitation or when temperatures are higher. In addition, outliers in both the precipitation and temperature distributions (values above and below the 99th and 1st percentiles, respectively) are highly significant predictors of the beginning of a weather suspension (table 4.4).

This suggests that weather suspensions arise in the aftermath of more extreme weather episodes (such as floods and snowstorms), but also that they are activated by (unexpected) weather variation at lower intensity levels. Consequently, suspensions are likely to hinder firm activity more frequently and in more locations than rare natural catastrophes alone would, and are thus a distinct disruptive phenomenon for the economy.

[3] I obtain grid-level precipitation and weather data from the CRU TS (v4.07) and IMERG GPM (v07B) repositories, and aggregate it to the level of Italian municipalities following Goodman et al. (2019) and Runfola et al. (2020).

Table 4.3: Predicting suspensions using weather bins

	Suspension start	
	(1)	(2)
Precipitation	0.00015***	
	(0.00001)	
Temperature	-0.00526***	
	(0.00037)	
Precipitation Q1		0.00012***
		(0.00003)
Precipitation Q2		0.00014***
		(0.00002)
Precipitation Q3		0.00013***
		(0.00001)
Precipitation Q4		0.00016***
		(0.00001)
Temperature Q1		-0.00695***
		(0.00043)
Temperature Q2		-0.00698***
		(0.00040)
Temperature Q3		-0.00540***
		(0.00038)
Temperature Q4		-0.00464***
		(0.00036)
Constant	0.10491***	0.10857***
	(0.00499)	(0.00520)
Observations	435096	435180
R^2	0.14	0.14
City FEs	Yes	Yes
Year, Month FEs	Yes	Yes

Temperature is the monthly average of
daily maximum temperature in degrees celsius.
Precipitation is the monthly average of
daily mean precipitation in mm per hour.
Q1-4 denote the four bins in which
the temperature and precipitation variables are split,
following the quartiles of the variables by city.
Data sources: CRU TS (v4.07), IMERG GPM (v07B).
* p<0.10, ** p<0.05, *** p<0.01

Table 4.4: Predicting suspensions with weather outliers

	Suspension start		
	(1)	(2)	(3)
Precipitation	0.00015***		0.00015***
	(0.00001)		(0.00001)
Temperature	-0.00516***		-0.00517***
	(0.00036)		(0.00038)
Precipitation p99		0.04060***	0.01360***
		(0.00506)	(0.00526)
Precipitation p1		0.00257	0.01722***
		(0.00313)	(0.00310)
Temperature p99		0.01601***	0.02268***
		(0.00214)	(0.00222)
Temperature p1		0.02223***	0.01025**
		(0.00424)	(0.00435)
Constant	0.10203***	0.05093***	0.10158***
	(0.00484)	(0.00007)	(0.00501)
Observations	435096	435180	435096
R^2	0.14	0.14	0.14
City FEs	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes
Month FEs	Yes	Yes	Yes

Temperature is the monthly average of
daily mean temperature in degrees celsius.
Precipitation is the monthly average of
daily mean precipitation in mm per hour.
p99 (p1) is a dummy denoting the 99th (1st) percentile by city
of the monthly average of daily maximum (minimum)
temperature or precipitation.
Data sources: CRU TS (v4.07), IMERG GPM (v07B).
* p<0.10, ** p<0.05, *** p<0.01

4.2 Data on firm balance sheets and income statements

I obtain income statement and balance sheet data for Italian firms from *Orbis Historical*. This information is derived from the compulsory and periodic data submissions of incorporated firms to their local business registry.

The key advantage of the *Orbis* data is its high coverage of firms across the size distribution - namely, both large and publicly listed ones, as well as small ones. Conversely, this source has two chief limitations: on the one hand, its low frequency (yearly), which is determined by how often firms submit their balance sheets, and on the other, the poor coverage of firm entry and exit dynamics. More specifically, firm market entry need not coincide with entry into the dataset, due to incomplete reporting of firms in their early years. Furthermore, firm exit from the *Orbis* sample is not well documented, given that the dataset does not clearly distinguish between exits that are voluntary (e.g. due to a merger) or involuntary (e.g. due to bankruptcy) (Bajgar et al., 2020).

Consequently, the literature examining firm entry and exit using Italian data typically relies on an additional data source (*InfoCamere*), directly produced by the Italian business registry and not publicly available (Cappelletti et al., 2024; Cascarano et al., 2023). Thus, in this paper I focus on dependent variables that are conditional on the firm remaining in the *Orbis* sample (e.g. sales, employment), rather than considering the impact on market entry and exit. Despite these shortcomings, *Orbis* firm samples have been shown to be nationally representative and thus useful to analyse economy-wide fluctuations, while having well-documented data cleaning routines to remove likely errors and outliers, such as negative values for total assets or inconsistent balance sheets (see section 9 and Kalemli-Ozcan et al. (2023)).

Table 4.5: Sample of firms with public works contracts

	N	25th P.	Median	Mean	75th P.
Sales	376,056	0.5	1.3	9.6	3.2
Employees	328,273	4.0	9.0	45.3	19.0
Total assets	376,324	0.6	1.4	24.8	3.8

Sales, assets in million EUR. 33,049 firms in the sample.

Merging the *Orbis* with the BDNCP data leaves a sample of 33,049 firms with public works contracts, with data on financials and on firm characteristics - this is around 68% of the original BDNCP sample of firms to which public works contracts have been assigned. 25,882 of the firms in this merged dataset are in the construction sector (NACE codes 41-43) - the only one I shall use in the DiD analyses to come. Much of the remainder operates in closely related ones - chiefly industrial and electric machinery, metals and metallic products. Table 4.5 shows that this sample is made up of mainly medium rather than small firms, with nine median employees and over one million euros in median sales and median total assets.

4.3 Event study samples

To examine the impact of suspensions in my event study specification, I focus on the subset of firms in this merged BDNCP-Orbis dataset with contract suspensions. More specifically, I construct two samples:

- **Wide** event study sample: *All firms experiencing at least one contract suspension.* (8,131 firms)

- **Narrow** event study sample: *Contracts suspended for at most one year, matched with firms experiencing a single year of contract suspensions.* (1,997 firms)

Though the former sample is much larger than the latter, it is less restrictive in specifying the timing of the suspension treatment at work. On the other hand, the narrow sample includes only firms that are treated with a contract suspension just once (i.e. for at most a year). This brings some advantages when delineating the long-term effects of suspensions. Since the two elements of this trade-off matter, both samples will be used in my matched DiD design.

Table 4.6: Event study samples' characteristics

	N	25th P.	Median	Mean	75th P.
Wide					
Sales	101,397	0.7	1.8	5.7	4.0
Employees	88,647	6.0	11.0	25.0	22.0
Total assets	101,463	0.9	2.0	9.2	4.9
Narrow					
Sales	21,441	0.4	1.1	4.5	2.5
Employees	18,764	4.0	8.0	21.2	16.0
Total assets	21,451	0.5	1.3	8.0	3.0

Sales, assets in million EUR.

Table 4.6 provides the distribution of firm sales, number of employees and total assets. The two samples are similar, though the wide one does include slightly larger firms. The single contract suspension restriction used in constructing the narrow sample may have driven a selection towards smaller firms with fewer active contracts.

Table 4.7: Suspension exposure of treated firms

	N	25th P.	Median	Mean	75th P.
Wide event study sample					
Value of susp. contract / sales (single firms)	4,500	8.8	20.2	902.2	48.1
Value of susp. contract / sales (consortia)	1,944	20.9	47.8	314.4	116.5
Narrow event study sample					
Value of susp. contract / sales (single firms)	1,219	7.2	18.3	90.1	45.4
Value of susp. contract / sales (consortia)	340	14.0	38.8	147.2	109.7

Sales refer to pre-suspension firm averages. N is the total number of firms.

It is important to note that the firms in these event study samples are highly exposed to the suspensions of the public contracts they hold. Table 4.7 drives this point home by reporting the ratio between the total value of a firm's suspended public contracts in a given year and her pre-suspension average sales. This is approximately the upper bound of the importance of suspensions for a firm given the size of her overall portfolio of projects. This is because if suspensions occurred exactly at the beginning of each contract, and not later, this measure of contract value would match the stock of expected future payments that are deferred due to the suspension.

It is more intuitive to examine the "single firms" rows in table 4.7, i.e. those with firms that undergo suspensions only on the contracts that they hold individually, and not in a consortium with others. This means that all contract payments would have contributed to those firms' sales only, and in turn that numerator and denominator refer to the same entity. For the median firm in my event study samples, 18-20% worth of her sales are taken up by suspended contracts. For about a quarter of the most exposed firms in the two samples (above the 75th percentile), the value of suspended contracts reaches nearly half of their total sales (table 4.7). These are sizeable shares, suggesting that suspension risk is likely to be a major threat to the financial viability of firms.

Another way of examining the salience of contract suspensions for firms is to consider the duration of their suspension spells, i.e. how long periods with contract interruptions last once they have begun. In the year in which a firm experiences her first suspension, the median number of her suspension days is around one month in both samples. Considering the total duration of all suspension spells experienced by a firm, the median figure reaches 98 days, with one quarter of intensely treated firms experiencing above 231 suspension days (nearly 8 months) during their lifetimes.

Table 4.8: Suspension spells for treated firms

	N	25th P.	Median	Mean	75th P.
Narrow event study sample					
Suspension days	1,992	15.0	28.0	39.9	52.0
Wide event study sample					
Suspension days when first treated	7,997	14.0	28.0	40.5	52.0
Total suspension days	7,997	39.0	98.0	226.8	231.0

Total suspension days are summed over a firm's active contracts and lifespan; first treated refers to the first susp. year only. The two coincide in the narrow sample. N is the total number of firms.

Given the severe impact of each suspension - the complete halt of revenues arising from affected contracts - this sort of suspension spell duration is clearly an important disruption to firms' liquidity positions. Quantifying this disruption will require an appropriate empirical design, to which I now turn.

5 Empirical set-up

5.1 Identification challenge

Isolating the *suspension channel* through which adverse weather events disrupt economic activity requires careful empirical design choices, since there are in fact multiple channels at work whenever weather (suspensions) take place:

- **Negative public demand effect**

  due to cash flows to the firm being interrupted. This is the *suspension channel*;

- **Positive local demand effect**

  due to greater reconstruction, maintenance and repair work in localities hit by the weather event;

- **Negative local demand effect**

as existing construction contracts are subject to disruption, cancellation, litigation due to the physical impact of the weather event;

- **Negative supply effect**

  due to weather damages to the factors of production of the firm herself.

Hence, my empirical approach will be guided by the need to separate the suspension channel from these other local demand and supply effects. I will chiefly achieve this by comparing the performance of treated and control firms that are in the same sector (construction) and location, which should ensure that they have access to the same market and to the same physical exposure to the weather event.

5.2 Difference-in-Differences (DiD) methodology

With contract suspension as the staggered treatment, the relevant two-way fixed effect (TWFE) regression to estimate its effect would be the following:

$$log(Y_{it}) = \alpha_0 + \alpha_1 susp_{it} + \boldsymbol{\alpha_2 X_{it-1}} + \gamma_{\mathbf{FE}} + u_{it} \tag{2}$$

where Y is a firm outcome, *susp* is a dummy equal to one whenever firm i in year t has at least one public contract suspension, X is a vector of time-varying firm financial controls, and γ_{FE} is a vector of fixed effects, including interactions of firm, sector, location and year FEs.

Since in my setting treatment effects may be time-varying (as the treatment is staggered, i.e. suspensions hit different firms and places at different times), TWFE is likely to fail to correctly weigh the average treatments effects across units and time, as TWFE typically compares also newly- with previously-treated units. This means that in this setting TWFE cannot produce the causal estimates of interest (Goodman-Bacon, 2021). The recognition of this problem has fuelled a vast literature developing DiD estimators in applied settings (Baker et al., 2022).

The DiD estimator I adopt here - Callaway and Sant'Anna (2021) - is particularly suited to my empirical application, where the weather suspension treatment is binary, staggered, irreversible and arising from a weather shock exogenous to the firm (no anticipation conditional on observables).

I define my suspension treatment in line with common Callaway and Sant'Anna (2021) practice, based on when each firm is *first treated*. This means that the start of the first suspension experienced by a firm marks the beginning of her treatment, and that, once treated, firms may never count as untreated again, or be part of the control group. This set-up means that the duration of suspension episodes is not used to construct the treatment.[4] The DiD control group is defined as the *not yet treated* firms, as opposed to the never treated ones. This is because if there are omitted technological, political or personal reasons for which the never treated firms never receive a suspension - say, their machines' and employees' great ability to work at subzero temperatures, or their close ties to the public administration involved in ruling on whether the contract should be suspended or not - then this would invalidate them as the appropriate control group for the treated. These concerns should however disappear when using the not yet treated as controls, since they too will eventually receive the suspension treatment. This means that they are unlikely to be fundamentally different from another treated firm of a similar size, location and sector.

On the operational side, I rely on the Callaway and Sant'Anna (2021) estimator also because of its several advantages. Firstly, it produces a set of "group-time average treatment effects on the treated", $ATT(g,t)$, which are defined by treatment group (cohort) (g), i.e. time

[4] This seems prudent, since recent research on the endogeneity of project completion times to firm decision-making (Fernandes and Rigato, 2024) might alert us to the fact that suspension duration could in practice be influenced by some unobserved bargaining between the public entity and contracting firm.

when first treated, and time (t). In other words, these raw ATTs are both heterogeneous by treated cohort and dynamic (time-varying). I mostly report these under their event study aggregation, showing how ATTs vary the longer each firm is exposed to the treatment. Secondly, Callaway and Sant'Anna (2021) already implicitly incorporates individual and time fixed effects, and supports (preferably time-invariant) covariates evaluated at pre-treatment levels, which are used in the matching between treated and control groups (Baker et al., 2022; Daw and Hatfield, 2018). Thirdly, this DiD technique automatically clusters standard errors at the level of the panel variable (in this case, the firm level),[5] and can be run with a doubly robust estimator (Sant'Anna and Zhao, 2020).

The key assumption in Callaway and Sant'Anna (2021) needed for identification is that of *Conditional parallel trends based on not yet treated groups.* Borrowing from their notation, this means assuming:

$$E[Y_t(0) - Y_{t-1}(0)|X, G_g = 1] = \tag{3}$$

$$= E[Y_t(0) - Y_{t-1}(0)|X, D_s = 0, G_g = 0]$$

where Y denotes firm outcomes, G_g indicates when a firm's group is treated (periods 0 and 1, in this case), $D_s = 0$ denotes the unobserved counterfactual (had the cohort treated at time zero not been treated), and X is a vector of covariates. In a nutshell, this means that, conditional on the values of the covariates used in the matching process, had they not been treated, treated firms would have followed the same path as the control group (the not yet treated). The common, indirect way of checking whether this is a plausible assumption is by checking for common trends between the two groups of firms before the treatment is administered.

[5]This is helpful, as in my setting the firm level is also the level of the treatment allocation and of the identifying variation used in the DiD.

Table 5.1: Treatment rollout by event study sample

Year	**Number of firms**	
	Narrow	Wide
2008	25	512
2009	86	794
2010	135	842
2011	81	455
2012	200	738
2013	113	532
2014	141	559
2015	138	387
2016	91	370
2017	126	415
2018	140	494
2019	190	678
2020	125	539
2021	186	551
2022	127	259

Firms by first treatment year.

This is more likely to hold the "better" the matching covariates are. The key constraint on their use is defined by the number of firms treated each year, which I report for the two event study samples in table 5.1. This number essentially determines the pool of treated and control firms available for matching in each year. Higher numbers of treated firms allow for a richer matching, based on multiple firm characteristics. Hence, in my baseline DiD results I will focus on the wide sample, the relatively large numbers of which permit a more granular matching of firm characteristics.

The DiD covariates used for matching between treated and control firms, expressed as pre-suspension average levels by firm, are the following:

- **Total assets:** a measure of firm size, which is itself a proxy for a firm's available technology levels and degree of financial constraints;

- **Firm province:**[6] this captures the legally registered location of a firm, and can thus define the local construction market she has access to, together with the local institutions she operates under, e.g. local government offices and courts;[7]

- **Number of public contracts held:** a measure of the strength of a firm's links to local political administrations, which might influence how public contracts are allocated and decided upon once in place;

- **Total value of public contracts held as a share of firm sales:** to capture a firm's overall reliance on the public sector for her cash flows; firms that are similarly reliant might be similar in other (unobserved) respects too, e.g. in strategy and supply chain management.

This rich set of matching variables in my DiD design should allow me isolate the suspension channel, as the impact of the other channels (local construction market access, exposure to direct physical damage) is differenced out when comparing matched treated and control firms. Additionally, I investigate the direct physical damage (negative supply effect) potential threat to identification in section 9.4 of the appendix, and show that it is not a serious concern in my sample.

[6] For the narrow sample, I instead match on a firm's region, i.e. on the higher administrative unit compared to the province, due to the limited sample size.

[7] In this sense, a finer matching by city rather than province would not necessarily be an improvement, given that larger construction firms tend to operate concurrently in a large number of cities. A less disaggregated location variable is thus preferable to capture the reach of their local construction market.

6 Results

6.1 Liquidity effects

Here I put my matched DiD design to work, and first consider dependent variables that concern the liquidity effects of suspensions on firms, and ultimately their financial viability. In table 6.1 I compute the average effect of the suspension treatment on firm sales, number of employees and total assets, comparing this between the two event study samples.

Several finding emerge from table 6.1. Firstly, this research design in successful in comparing treated and control firms that are not statistically different from each other pre-treatment, and that then diverge sharply post-treatment (i.e. the *Pre* row coefficients are all insignificant). This is additional evidence that one might serve as counterfactual for the other after the treatment is administered.

Secondly, considering the wide sample, the impact of weather suspensions on firms is substantial: on average, weather suspensions reduce a firm's sales by around 30%, her number of employees by 15.3%, and her total assets (a proxy for firm size) by 18.5% compared to still-unsuspended firms.

Thirdly, these estimates relate to the wide sample, with firms potentially experiencing several contract suspensions in a row. One might thus expect that they would be able to adapt to suspension risk the more frequently they are exposed to it, e.g. by taking up more comprehensive insurance, borrowing, or adopting more weather-resilient technologies. Instead, in the narrow event study sample we see firms that are hit by a weather suspension only once, and therefore see even larger negative effects, that (mechanically) exclude these adaptation efforts that firms might put in place after having been hit in the past. *Ceteris paribus*, one interpretation of the columns N and W coefficient differences is thus as the scale of firms' climate adaptation and financial mitigation efforts.

Table 6.1: Suspension channel effects

	log(Sales)		log(Employees)		log(Total assets)	
	N	W	N	W	N	W
Pre	-0.032	0.006	0.079	0.001	-0.003	0.016
	(0.061)	(0.012)	(0.070)	(0.007)	(0.044)	(0.010)
Post	-0.484***	-0.298***	-0.200***	-0.153***	-0.355***	-0.185***
	(0.069)	(0.030)	(0.041)	(0.017)	(0.048)	(0.018)
Observations	15,698	73,386	12,593	47,844	16,187	74,452

Treatment: first suspension start. All covariates.
N, W: narrow and wide event study samples, respectively.
Standard errors are clustered at the firm level.
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

In figures 4, 5, 6 and 7, I consider the time-varying nature of these treatment effects by plotting them over event time, i.e. by the number of periods since each firm is first treated. Visually, the insignificant pre-trend finding from table 6.1 is reaffirmed, together with the large economic and financial costs for firms that occur post-suspension. The barely visible, time 0 boost in firm sales is a mechanical effect of the suspension of public contracts, which may trigger an automatic payment to firms for their work done up to the point of the suspension.

Interestingly, these large negative effects are rising in absolute value for up to four years after the first suspension, reaching -40% for firm sales. The impact on earnings (EBITDA) is substantial, but attenuated, since it does not reach -20%, and becomes insignificant after the third year. This points to firms cushioning the impact of this liquidity freeze on their profits, potentially by cutting costs. Indeed, the number of employees drops dramatically and persistently by up to 20% in the third and fourth year after the first suspension, and does not recover. Firm size (total assets) declines by around one quarter (-25%) four years into the first suspension.[8]

[8] It should be noted that these effects are not driven by the fact that treated firms might experience (physical) weather damage to their productive capacity compared to untreated ones - I show this in appendix 9.4.

Overall, these results point to weather-driven suspensions inducing a collapse in cash flows and liquidity, and a persistent down-scaling in firm operations and market presence. This is robust to changing the event study sample from the wide to the narrow one (appendix 9.5), again with the caveat that the effects from this sample will tend to be larger, given the exclusion of adaptation mechanisms when each firm is treated just once. The comparison with the 9.5 figures, which are produced using suspension treatments hitting firms only at period 0, and for at most one year, is helpful, since it indicates that these large and persistent negative liquidity effects of suspensions do not arise simply as suspensions repeatedly hit firms, and thus mechanically produce long-term effects. Instead, even one year-long treatment is enough to produce this impact. This is suggestive evidence of a spillover of weather-suspension-related disruption from the initially affected contracts to other construction projects - I provide evidence for these knock-on effects in section 6.3.

In appendix 9.3, I consider more indirect effects on firms' financial positions, on working capital, investment, leverage and debt maturity. These effects are surprisingly muted. Although there is a noisy but significant drop in working capital, a measure of a firm's short-term liquidity position, investment only slightly drops four years into the first suspension. Debt maturity seems not to be impacted by the treatment, and firm leverage only briefly shoots up in the year of the first suspension, pointing to debts being incurred by firms in the immediate aftermath of a contract interruption, as a substitute for their deferred cash flows. These results are somewhat puzzling, though they do arise from financial variables that are not captured well in my Orbis data - e.g. investment is only observed indirectly as the change in (tangible) fixed assets, and the debt structure is only measured in a binary sense, comparing the short- and long-term debt held by a firm. Thus, delving into the exact financial ramifications of the suspension channel might require richer data on the financial choices available to firms, e.g. loan-level data on their liability position and debt-servicing payments.

Figure 4: log(Sales)

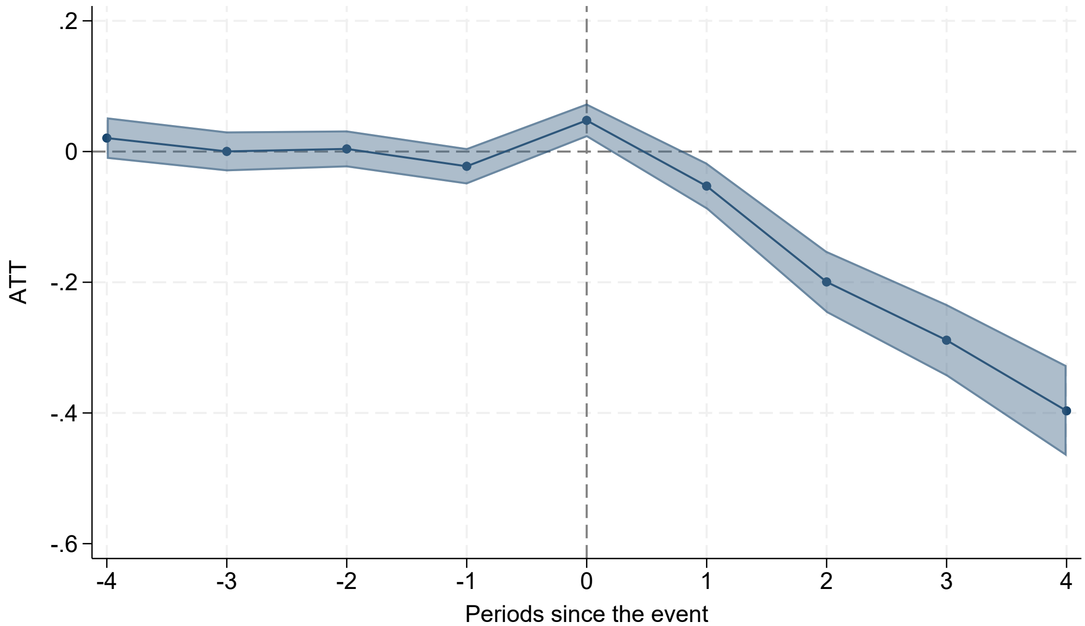

Note: all covariates, wide event study sample

Figure 5: log(EBITDA)

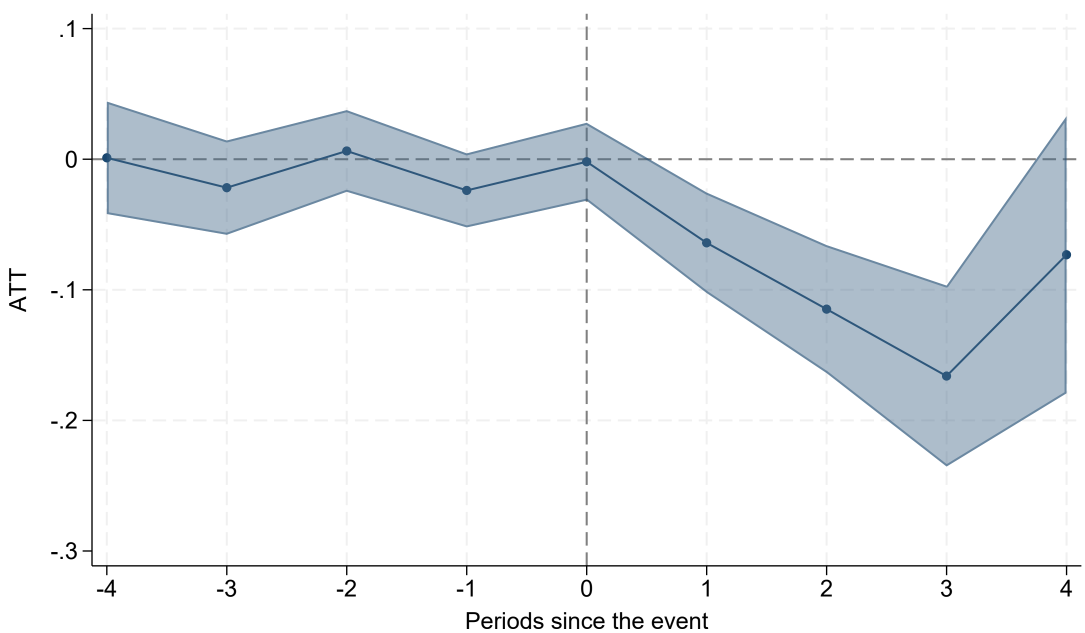

Note: all covariates, wide event study sample

Figure 6: log(Number of Employees)

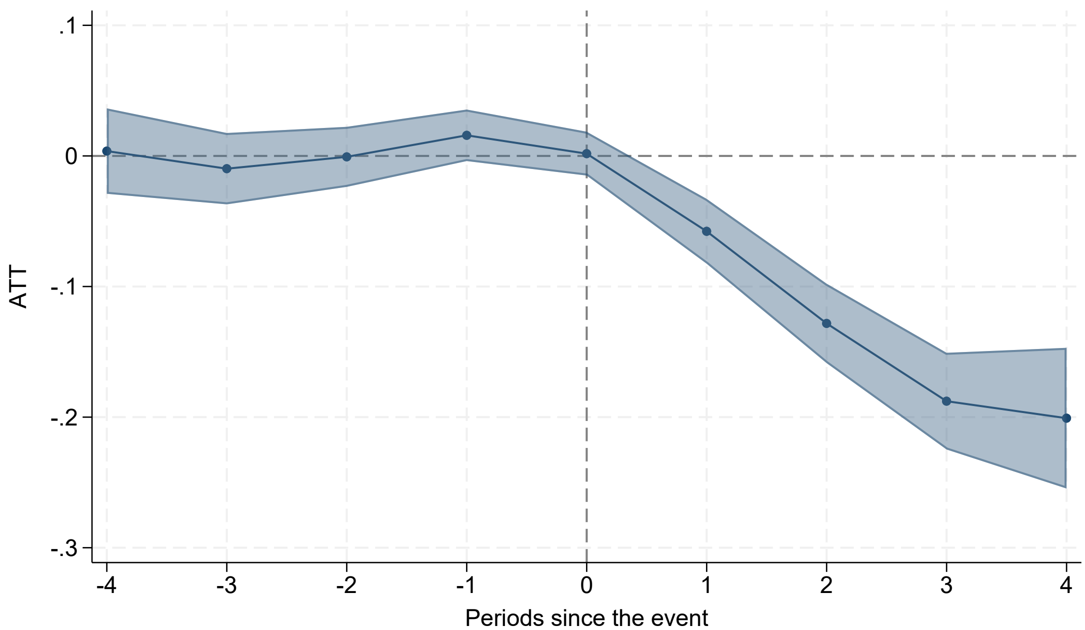

Note: all covariates, wide event study sample

Figure 7: log(Total Assets)

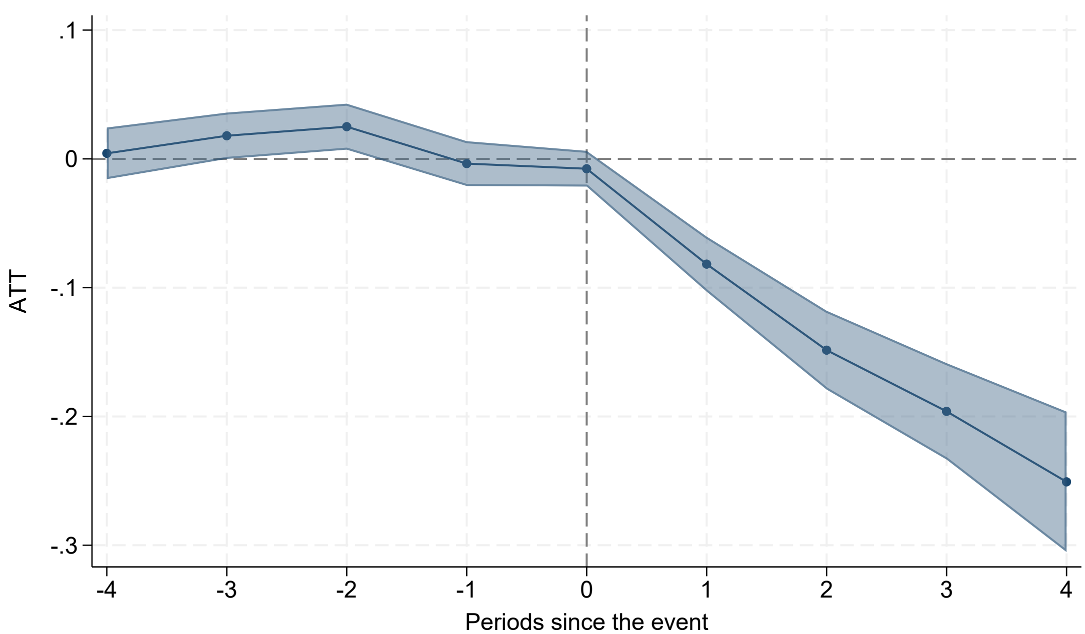

Note: all covariates, wide event study sample

6.2 Robustness to changes in sample composition

In this section I show that the liquidity effects of weather suspensions analysed so far are robust to changes in sample composition. Table 6.2 replicates the baseline effects on firm sales in the wide sample (column *All*), and then proceeds to slice the sample using the percentiles of the distribution of firm age. I thus first remove the firms below the 1st or 5th percentile of age, and then those above the 95th or 99th percentiles (final two columns). Intuitively, this is to test if the nearly 30% drop in average firm sales occurring post-suspension is driven by outliers in the sample. These might be either the very young firms - which perhaps were only set up to run a single construction project, on which they might rely fully - or the very old ones, which might be dependent on some regular, long-term contracts for their income. Though the coefficients along the *Post* row exhibit some variation around the 30% level, these changes remain within the respective confidence intervals, and are not statistically significant.

Table 6.2: Robustness by firm age

	log(Sales)				
	All	$> p1$	$> p5$	$< p95$	$< p99$
Pre	0.006	0.006	0.002	0.004	0.005
	(0.012)	(0.012)	(0.012)	(0.013)	(0.012)
Post	-0.298***	-0.329***	-0.356***	-0.355***	-0.347***
	(0.030)	(0.030)	(0.032)	(0.066)	(0.062)
Observations	73,386	72,832	70,842	58,144	66,074

Treatment: first suspension start.
Percentiles from the firm distribution of age when first treated.
Standard errors are clustered at the firm level.
Wide event study sample. All covariates.
* p<0.10, ** p<0.05, *** p<0.01

A very similar picture emerges when removing the outliers in terms of firm size, as proxied by total assets (table 6.3). Here I again reproduce the baseline estimates, then only keep the firms above the 1st or 5th percentiles, and below the 95th or 99th percentiles. There is no significant difference in the average effect of weather suspensions when either the very small or very large firms are removed from the sample. This is an important finding, since one might expect that larger firms might be more diversified than smaller ones, and that they might have better financial means to mitigate the effects of contract interruptions, such as easier access to bank credit. Nonetheless, it is not the case that the large suspension costs are driven by either group of firms. This is a key result in terms of robustness, which has arisen while maintaining insignificant pre-trends, indicating that the fit of my matched DiD design is also robust to the exclusion of small groups of firms on either end of the age and size distributions.

Table 6.3: Robustness by firm size

	log(Sales)				
	All	> p1	> p5	< p95	< p99
Pre	0.006	0.006	0.002	0.015	0.007
	(0.012)	(0.012)	(0.011)	(0.013)	(0.012)
Post	-0.298***	-0.305***	-0.303***	-0.317***	-0.289***
	(0.030)	(0.030)	(0.031)	(0.054)	(0.030)
Observations	73,386	73,240	71,773	68,687	72,282

Treatment: first suspension start.
Percentiles from the firm distribution of average pre-treatment total assets.
Standard errors are clustered at the firm level.
Wide event study sample. All covariates.
* p<0.10, ** p<0.05, *** p<0.01

Thirdly, in table 6.4 I examine the nature of the contracts suspended, distinguishing firms that receive contract suspensions only on contracts which they hold individually (*Single-firms*) from those whose suspensions involve some contracts held as part of a group of firms (*Consortia*). This is an important distinction, as one might worry that within a consortium,

Table 6.4: Robustness by suspended contract type

	log(Sales)		
	All	Single-firms	Consortia
Pre	0.006	0.003	0.085
	(0.012)	(0.013)	(0.080)
Post	-0.298***	-0.415***	-0.061
	(0.030)	(0.069)	(0.077)
Observations	73,386	43,950	4,641

Treatment: first suspension start.
Single-firms:
all susp. contracts are assigned to a single firm;
Consortia:
at least one susp. contract is assigned to multiple firms.
Standard errors are clustered at the firm level.
Wide event study sample. All covariates.
* p<0.10, ** p<0.05, *** p<0.01

if we assume a division of labour between firms in distinct chronological stages, some firms that have already provided their services and have been paid before the suspension might not actually be treated, though counting officially as such. In this case, the suspension treatment effect applied to consortia would be harder to interpret, because it would not arise from a pure liquidity shock to the treated firms. Reassuringly, my suspension channel effects are instead driven by single-firms only, who are unambiguously hit by the cash flow freeze, while the responsiveness of consortia is insignificant (table 6.4).

6.3 Knock-on effects

In this section I investigate one explanation for the large and long-lasting effects of weather suspensions - namely, the spillover of the disruption to other portions of a firm's portfolio of construction projects, i.e. to unsuspended contracts. This spillover means that an even larger share of a firm's revenues would be affected by contract suspensions. To explore this theme, I maintain the same matched DiD design and event study samples as before, and build dependent variables that measure the disruptions experienced across all public contracts held, i.e. both the suspended and unsuspended ones.

My hypothesis is that to the extent that firms have limited spare capacity, cash and ability to reallocate factors of production, one contract being suspended might delay the work on other contracts, giving rise to a knock-on effect of delayed payments on non-suspended contracts as well.

The intuition underlying this hypothesis comes from company operations' research, which suggests that when construction firms are involved with multiple concurrent construction projects, even if these are operationally independent, the delays in one may transmit to the others, thereby forming a "project network" (Chen et al., 2018). Constraints on the production capacity of constructions firms are common, e.g. because of limited availability of extra machinery and labour (O'Brien and Fischer, 2000). This means that when one project faces delays, factors of production may not be easily reallocated to other construction sites, and that a firm might not have enough spare capacity to bid for a new project to work on in the interim. This would delay the completion times and payments' schedule of other projects as well.

Furthermore, the cash flow squeeze arising from one contract suspension may disrupt other ongoing contracts held by the affected firm, insofar as concurrent projects may require significant upfront payments for machinery, equipment and materials at specific completion stages (He et al., 2023). As firm liquidity constraints bind and these costly purchases cannot be made in a timely manner, completion times of unsuspended contracts will also be lengthened. Both constrained spare capacity in factors of production and limited cash availability should then both lead to delayed payment tranches for firms, and thus ultimately to an even greater worsening in firms' financial positions after a weather suspension.

To measure contract disruptions, I use high-frequency data on contract payment times, again from the BDNCP database. This means I only observe payment times and delays on firms' ongoing public contracts, and not on their private ones. This data source is nonetheless valuable, since it tells me for each stage of a public construction project, and for each associated payment to the firm, if its completion was delayed or disrupted, and by how many days. I use this dataset to build two measures of disruption, and to test how the suspension treatment affects them: the number of contract payments being delayed (figure 8) and the probability that a contract held by the firm deviates from schedule, which is defined slightly differently in the dataset (figure 9).

The results in figures 8 and 9 are unambiguous: before the suspension, treated and control groups exhibit similar levels of contract delays; after the suspension, and already from the first year, there is a sharp and persistent rise in both disruption measures. More specifically, payment delays rise by about 15%, and the probability of a contract deviating from schedule by 7-8%. Both effects become slightly less pronounced over time, but remain sizeable and highly significant four years after the first suspension.

These results point to limited cash flows and spare capacity acting to transmit the effects of weather suspensions across a firm's portfolio of projects. Consistently with this mechanism, figure 10 indicates that the accumulation of inventories of raw materials and finished goods, which would be the primary form of self-insurance available to firms against production process disruption, is negatively affected, though with a lag. Indeed, inventories drop by about 10% two years after the first weather suspension experienced by a firm, i.e. they are substantially used up. This delayed depletion does not seem to be driven by the initial liquidity effects of suspensions, but rather by their medium-term knock-on effects on other contracts, which force firms to draw on their stock of inventories in order to mitigate disruption across their ongoing projects.

Figure 8: Contract Payments Delayed

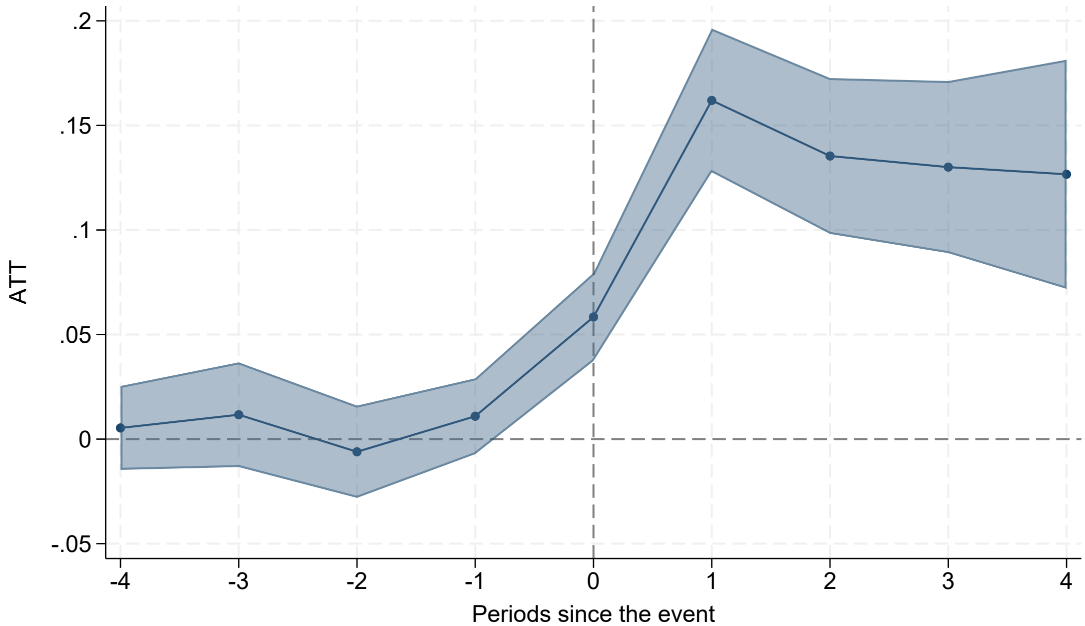

Note: all covariates, wide event study sample

Figure 9: Contracts Deviating from Schedule

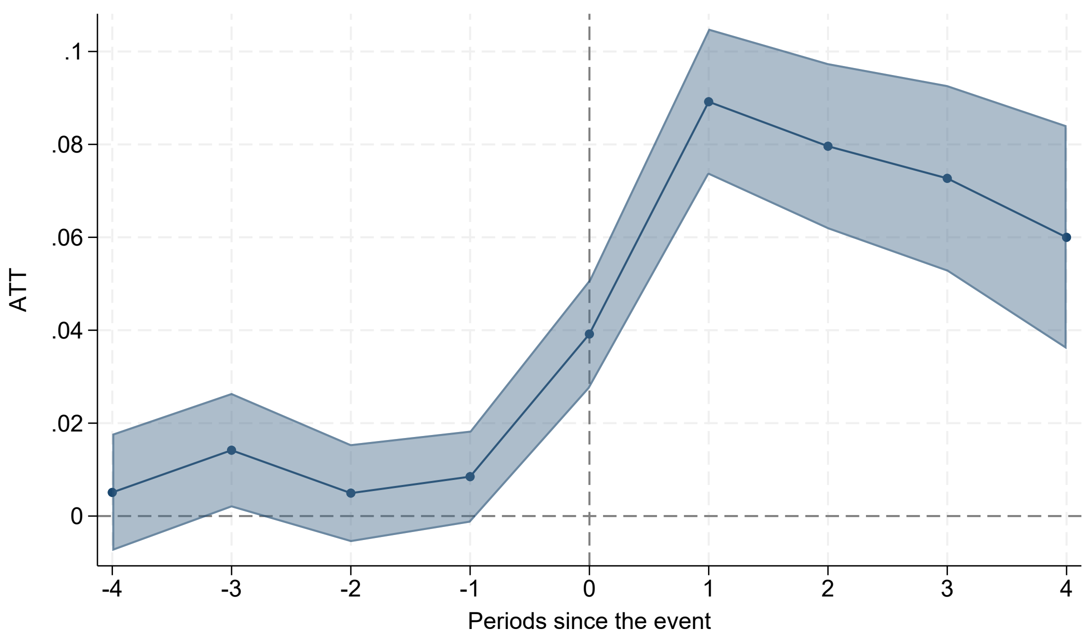

Note: all covariates, wide event study sample

Figure 10: log(Inventories)

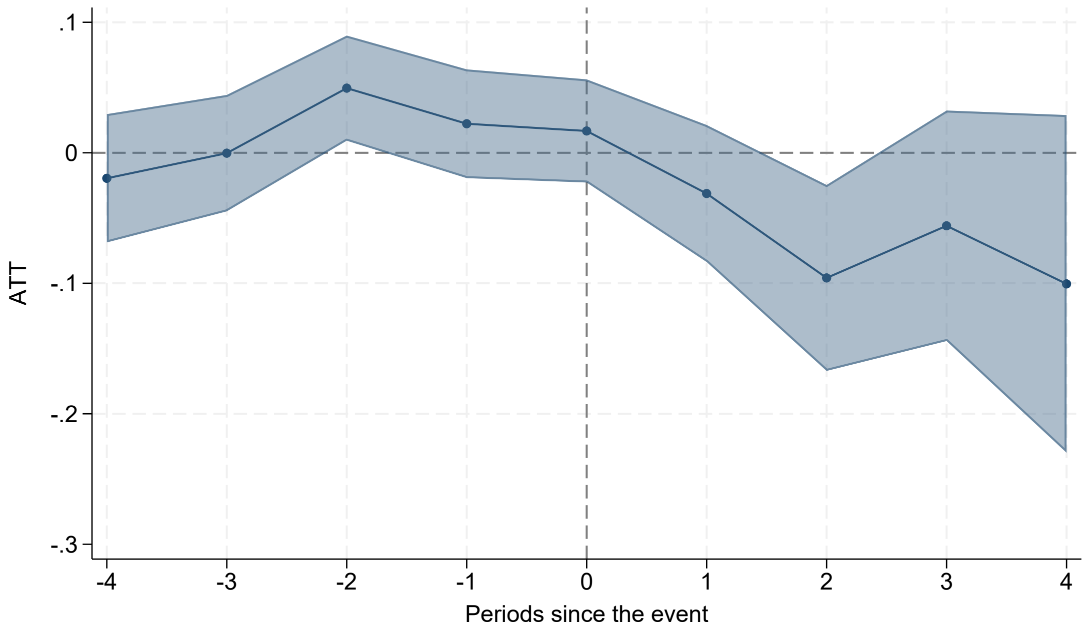

Note: all covariates, wide event study sample

7 Discussion

7.1 External validity

The suspension channel induced by adverse weather events has large, robust and persistent negative liquidity effects on firms. Concurrently, the disruption of suspended contracts has a knock-on effect on firms' other ongoing (unsuspended) construction projects, which are themselves delayed. This amplifies the scale of the cash flow freeze experienced by firms, and rationalises its remarkable scale and duration.

Crucially, the impact of the suspension channel does not correspond to the total effect of climate shocks on firms, which may well be less severe. For instance, construction firms may benefit from the increased reconstruction and repair work needed post-event. This mechanism is however excluded in my matched DiD design, since I compare firms operating in the same local market.

Moreover, the substantial suspension effects I pick up should be treated as an upper bound, particularly in relation to other countries and sectors. This is for several reasons. Firstly, because of the specific Italian regulatory context, which places all liquidity risk on firms in the event of a suspension (section 3). In other contexts, e.g. in private sector contracts, contract interruptions may also occur, but with some degree of payments from clients to firms continuing, depending on their relative bargaining power. Secondly, since the firms in my sample are highly exposed to the construction projects that end up being suspended (table 4.7), suggesting that this a group of firms with a high reliance on public sector demand. While this is common for construction firms, those operating in other weather-exposed sectors, like airlines, might rely much more on their private sector clients. Thirdly, construction itself as a sector is by design directly exposed to adverse weather events, with most infrastructure work occurring outdoors. Conversely, other important sectors that may be vulnerable to climate, like hospitality, should still be able to largely function indoors, notwithstanding some weather-driven disruptions to their operations.

Simply put then, my estimates of the impact of the suspension channel arise from a setting with minimal mitigating factors shielding firms from the suspension treatment. While this has been qualitatively important in examining how this channel disrupts economic activity, other contexts with different sectors, legal and contractual arrangements may well exhibit quantitatively smaller effects.

7.2 Implications for insurance

Though I do not observe granular data on insurance packages held by firms, my results indicate that on average, insurance coverage must be highly incomplete in my sample, at least against this particular type of climate-driven business interruption risk.

In principle, an insurance package covering cases of contract interruptions that are not necessarily caused by physical damage to the firm, and that are exogenously driven, could fully offset the negative impact on firm liquidity of the cash flow freeze, assuming speedy payments coming from the insurance company. This is however not what can be observed in my empirical results, which suggests that such packages are not widely taken up by Italian construction firms, and/or not widely offered in the insurance market.

Indeed, while insurance on physical damage to tangible assets is commonplace among Italian firms, that against specific classes of business interruption is not (Frigo and Venturini, 2024; Gallo et al., 2022). As a consequence, my results support the notion that the development of sophisticated insurance markets to withstand climate risk is still incomplete.

7.3 Long-run trend

The liquidity and knock-on effects of the weather suspension channel do not exist in a vacuum. The weather distribution underlying these disruptions - climate - is changing, becoming warmer and more volatile. Here I provide illustrative evidence in this direction by aggregating the average treatment effects produced in my DiD design by *calendar time*, rather than event time. In other words, in figure 11 I plot the cumulative effects of suspensions on total assets by (calendar) year of first treatment.

Notice that for better inter-temporal comparisons, in figure 11 I only include suspension treatments on total assets lasting for at most one year (narrow event study sample), meaning that over time there is limited variation in the duration of the suspension episodes, and that the direct impact of each suspension episode lasts only one period. The different coefficients we observe in figure 11 should thus arise chiefly from differences in the underlying climate

Figure 11: Cumulative suspension effects by first treatment year

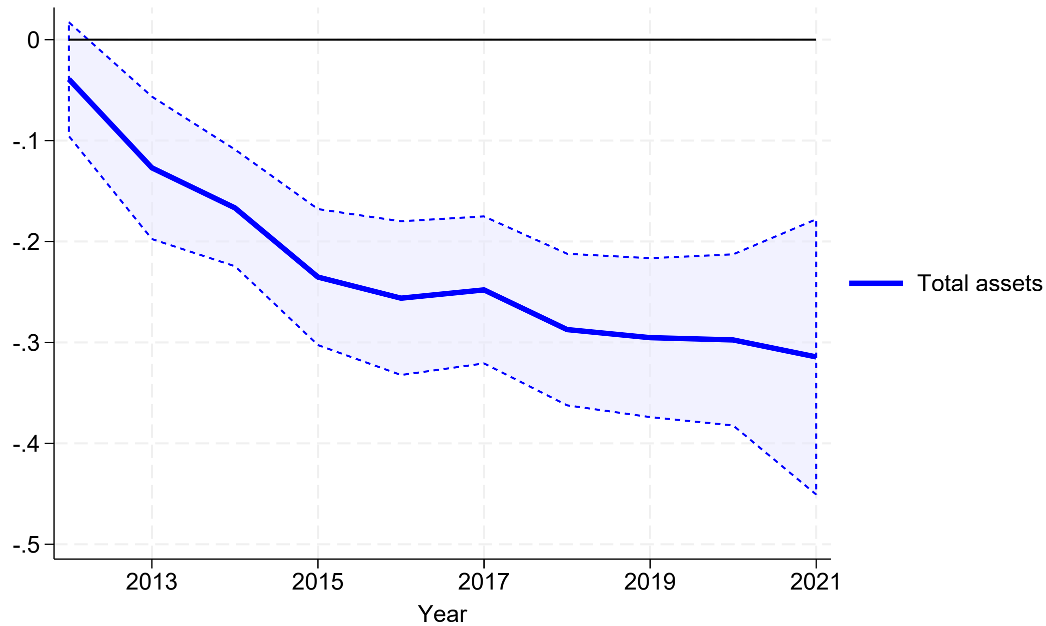

Note: all covariates, narrow event study sample

shocks.[9]

Average cumulative effects of the weather suspension channel are worsening over 2012-21, albeit at a decreasing rate. At the beginning of the sample period these are insignificant, before reaching -30% of firm total assets by the early 2020s. This worsening long-run trend is consistent with the long-run evolution of climate patterns, and is highly concerning. Given that the suspension channel does not capture the physical damage arising from weather events, but only the liquidity costs and disruptions arising from contract interruptions, this downward trend could arise from factors such as higher climate and climate policy uncertainty, and stricter bank lending policies in light of transition risk. Both of these would act to exacerbate the impact of weather suspensions on firms.

[9] The limitations of this statement are twofold. Firstly, heterogeneous firms might be treated at different times in the narrow sample, and thus give rise to different coefficients. Secondly, although the narrow sample leaves little scope for adaptation, as each firm is treated only once, in fact firm adaptation to climate risks should actively mitigate these negative effects.

8 Conclusion

In this paper, I study the effects of weather-driven interruptions of economic activity using data on Italian public construction contracts and on firm financials. I can isolate this *suspension channel* through which climate change is disrupting economic activity by deploying a matched DiD design comparing firms first treated with a suspension with those not yet treated.

I find that weather suspensions deplete firm liquidity, giving rise to substantial and persistent drops in firm sales, earnings, employment and total assets. However, I find limited evidence of an impact on investment, leverage and maturity. Concurrently, I document that suspensions on some contracts have knock-on effects on the other contracts held by firms, inducing delays to their completion times. The postponement of the payment schedules of unsuspended contracts exacerbates the costs of suspensions for firms.

Overall, I contend that the scale of the firm-level effects of weather suspensions mean that this is a key channel through which the disruptive effects of climate change materialise. Furthermore, my results from the Italian construction market provide indirect evidence that firms are on average inadequately insured against it.

References

Viral Acharya, Abhishek Bhardwaj, and Tuomas Tomunen. Do Firms Mitigate Climate Impact on Employment? Evidence from US Heat Shocks. *NBER Working Papers*, w31967, 12 2023. doi: 10.3386/w31967.

Jawad M. Addoum, David T. Ng, and Ariel Ortiz-Bobea. Temperature shocks and establishment sales. *Review of Financial Studies*, 33(3):1331–1366, 3 2020. ISSN 14657368. doi: 10.1093/rfs/hhz126.

Matej Bajgar, Giuseppe Berlingieri, Sara Calligaris, Chiara Criscuolo, and Jonathan Timmis. Coverage and representativeness of Orbis data. *OECD Science, Technology and Industry Working Papers*, 6, 2020.

Andrew C. Baker, David F. Larcker, and Charles C.Y. Wang. How much should we trust staggered difference-in-differences estimates? *Journal of Financial Economics*, 144(2): 370–395, 5 2022. ISSN 0304405X. doi: 10.1016/j.jfineco.2022.01.004.

Clare Balboni, Johannes Boehm, and Mazhar Waseem. Firm Adaptation in Production Networks: Evidence from Extreme Weather Events in Pakistan, 2023.

Audinga Baltrunaite, Cristina Giorgiantonio, Sauro Mocetti, and Tommaso Orlando. Discretion and Supplier Selection in Public Procurement. *The Journal of Law, Economics, and Organization*, 37(1):134–166, 4 2021. ISSN 8756-6222. doi: 10.1093/jleo/ewaa009.

Audinga Baltrunaite, Enza Maltese, Tommaso Orlando, and Gabriele Rovigatti. Procurement managers and effective tendering: The case of Italian public works contracts. *Bank of Italy Occasional Papers*, 803, 10 2023.

Henrique S. Basso, Myroslav Pidkuyko, and Omar Rachedi. Opening the Black Box: Aggregate Implications of Public Investment Composition, 2 2024.

Andrea Caggese, Andrea Chiavari, Sampreet S. Goraya, and Carolina Villegas-Sanchez. Climate change, firms, and aggregate productivity. *CEPR Discussion Papers*, DP19164, 6 2024.

Brantly Callaway and Pedro H.C. Sant'Anna. Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 12 2021. ISSN 03044076. doi: 10.1016/j.jeconom.2020.12.001.

Matilde Cappelletti, Leonardo Maria Giuffrida, and Gabriele Rovigatti. Procuring survival. *CEPR Discussion Papers*, DP18796, 1 2024.

Vasco M Carvalho, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. Supply Chain Disruptions: Evidence from the Great East Japan Earthquake*. *The Quarterly Journal of Economics*, 136(2):1255–1321, 3 2021. ISSN 0033-5533. doi: 10.1093/qje/qjaa044.

Michele Cascarano, Filippo Natoli, and Andrea Petrella. Entry, exit, and market structure in a changing climate. *Bank of Italy Working Papers*, 1418(7):1–63, 7 2023.

Juanma Castro-Vincenzi. Climate hazards and resilience in the global car industry, 6 2024.

Juanma Castro-Vincenzi, Gaurav Khanna, Nicolas Morales, and Nitya Pandalai-Nayar. Weathering the storm: supply chains and climate risk. *NBER Working Papers*, w32218, 5 2024.

Serhan Cevik and Fedor Miryugin. Rogue Waves: Climate Change and Firm Performance. *IMF Working Papers*, 102(5), 5 2022.

Weiwei Chen, Lei Lei, Zhengwei Wang, Mingfei Teng, and Junming Liu. Coordinating supplier selection and project scheduling in resource-constrained construction supply chains. *International Journal of Production Research*, 56(19):6512–6526, 10 2018. ISSN 0020-7543. doi: 10.1080/00207543.2018.1436782.

Stefano Clò, Francesco David, and Samuele Segoni. The impact of hydrogeological events on firms: Evidence from Italy. *Journal of Environmental Economics and Management*, 124: 102942, 3 2024. ISSN 00950696. doi: 10.1016/j.jeem.2024.102942.

Francesco Paolo Conteduca and Ludovic Panon. Natural Disasters and Markups, 12 2024.

Cláudia Custódio, Miguel A Ferreira, Emilia Garcia-Appendini, and Adrian Lam. How Does Climate Change Affect Firm Sales? Identifying Supply Effects, 2022.

Jamie R. Daw and Laura A. Hatfield. Matching and Regression to the Mean in Difference-in-Differences Analysis. *Health Services Research*, 53(6):4138–4156, 12 2018. ISSN 0017-9124. doi: 10.1111/1475-6773.12993.

Julian di Giovanni, Manuel García-Santana, Priit Jeenas, Enrique Moral-Benito, and Josep Pijoan-Mas. *Buy Big or Buy Small?: Procurement Policies, Firms' Financing, and the Macroeconomy*. World Bank, Washington, DC, 7 2023. doi: 10.1596/1813-9450-10522.

Mitch Downey, Nelson Lind, and Jeffrey G. Shrader. Adjusting to Rain Before It Falls. *Management Science*, 69(12):7399–7422, 12 2023. ISSN 0025-1909. doi: 10.1287/mnsc.2023.4697.

Robert J.R. Elliott, Yi Liu, Eric Strobl, and Meng Tong. Estimating the direct and indirect impact of typhoons on plant performance: Evidence from Chinese manufacturers. *Journal of Environmental Economics and Management*, 98, 11 2019. ISSN 10960449. doi: 10.1016/j.jeem.2019.102252.

Thijs Endendijk, W.J. Wouter Botzen, Hans de Moel, Kymo Slager, Matthijs Kok, and Jeroen C.J.H. Aerts. Enhancing resilience: Understanding the impact of flood hazard and vulnerability on business interruption and losses. *Water Resources and Economics*, 46: 100244, 4 2024. ISSN 22124284. doi: 10.1016/j.wre.2024.100244.

Serena Fatica, Gabor Katay, and Michela Rancan. Floods and firms: vulnerabilities and resilience to natural disasters in Europe. *JRC Working Papers in Economics and Finance*, 13, 2022.

Adriano Fernandes and Rodolfo Rigato. K Wasn't Built in a Day: Investment with Endogenous Time to Build, 7 2024.

Claudio Ferraz, Frederico Finan, and Dimitri Szerman. Procuring Firm Growth: The Effects of Government Purchases on Firm Dynamics. *NBER Working Papers*, w21219, 5 2015. doi: 10.3386/w21219.

Annalisa Frigo and Andrea Venturini. La copertura assicurativa contro i rischi naturali: un'analisi preliminare. *Bank of Italy Occasional Papers*, 830(2), 2 2024.

Ricardo Duque Gabriel. The Credit Channel of Public Procurement. *Journal of Monetary Economics*, 147:103601, 10 2024. ISSN 03043932. doi: 10.1016/j.jmoneco.2024.103601.

Raffaele Gallo, Giovanni Guazzarotti, Valentina Nigro, and Marco Cosconati. Le coperture assicurative contro i rischi operativi delle imprese italiane: alcune evidenze dai risultati dell'indagine sulle imprese industriali e dei servizi della Banca d'Italia. *Banca d'Italia - Note di stabilità finanziaria e vigilanza*, 31(10):1–15, 10 2022.

Seth Goodman, Ariel BenYishay, Zhonghui Lv, and Daniel Runfola. GeoQuery: Integrating HPC systems and public web-based geospatial data tools. *Computers & Geosciences*, 122: 103–112, 1 2019. ISSN 00983004. doi: 10.1016/j.cageo.2018.10.009.

Andrew Goodman-Bacon. Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277, 12 2021. ISSN 03044076. doi: 10.1016/j.jeconom. 2021.03.014.

Arti Grover and Matthew E. Kahn. Firm adaptation to climate change. *NBER Working Papers*, w32848, 8 2024.

Yukang He, Tao Jia, and Weibo Zheng. Tabu search for dedicated resource-constrained multiproject scheduling to minimise the maximal cash flow gap under uncertainty. *European Journal of Operational Research*, 310(1):34–52, 10 2023. ISSN 03772217. doi: 10.1016/j.ejor.2023.02.029.

Graham J. Ive and Stephen L. Gruneberg. *The Economics of the Modern Construction Sector*. Palgrave Macmillan UK, London, 2000. ISBN 978-0-333-62662-7. doi: 10.1057/9780230510913.

Sebnem Kalemli-Ozcan, Bent E Sørensen, Carolina Villegas-Sanchez, Vadym Volosovych, and Sevcan Yesiltas. How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration, 2023.

Agnes Norris Keiller and John Van Reenen. Disaster Management, 2023.

Jaroslav Mysiak, Silvia Torresan, Francesco Bosello, Malcolm Mistry, Mattia Amadio, Sepehr Marzi, Elisa Furlan, and Anna Sperotto. Climate risk index for Italy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 376(2121):20170305, 6 2018. ISSN 1364-503X. doi: 10.1098/rsta.2017.0305.

William J. O'Brien and Martin A. Fischer. Importance of Capacity Constraints to Construction Cost and Schedule. *Journal of Construction Engineering and Management*, 126(5):366–373, 10 2000. ISSN 0733-9364. doi: 10.1061/(ASCE)0733-9364(2000)126:5(366).

Nora Pankratz, Rob Bauer, and Jeroen Derwall. Climate Change, Firm Performance, and Investor Surprises. *Management Science*, 69(12):7352–7398, 12 2023. ISSN 15265501. doi: 10.1287/mnsc.2023.4685.

Nora M C Pankratz and Christoph M Schiller. Climate Change and Adaptation in Global

Supply-Chain Networks. *The Review of Financial Studies*, 37(6):1729–1777, 5 2024. ISSN 0893-9454. doi: 10.1093/rfs/hhad093.

Jacopo Ponticelli, Qiping Xu, and Stefan Zeume. Temperature, Adaptation, and Local Industry Concentration. *NBER Working Papers*, w31533, 8 2023. doi: 10.3386/w31533.

Daniel Runfola, Austin Anderson, Heather Baier, Matt Crittenden, Elizabeth Dowker, Sydney Fuhrig, Seth Goodman, Grace Grimsley, Rachel Layko, Graham Melville, Maddy Mulder, Rachel Oberman, Joshua Panganiban, Andrew Peck, Leigh Seitz, Sylvia Shea, Hannah Slevin, Rebecca Youngerman, and Lauren Hobbs. geoBoundaries: A global database of political administrative boundaries. *PLOS ONE*, 15(4):e0231866, 4 2020. ISSN 1932-6203. doi: 10.1371/journal.pone.0231866.

Paola Sakai and Zeyu Yao. Financial losses and flood damages experienced by SMEs: Who are the biggest losers across sectors and sizes? *International Journal of Disaster Risk Reduction*, 91:103677, 6 2023. ISSN 22124209. doi: 10.1016/j.ijdrr.2023.103677.

Pedro H.C. Sant'Anna and Jun Zhao. Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122, 11 2020. ISSN 03044076. doi: 10.1016/j.jeconom.2020.06.003.

Steven J. Schuldt, Mathew R. Nicholson, Yaquarri A. Adams, and Justin D. Delorit. Weather-Related Construction Delays in a Changing Climate: A Systematic State-of-the-Art Review. *Sustainability*, 13(5):2861, 3 2021. ISSN 2071-1050. doi: 10.3390/su13052861.

Zakia Sultana, Tobias Sieg, Patric Kellermann, Meike Müller, and Heidi Kreibich. Assessment of Business Interruption of Flood-Affected Companies Using Random Forests. *Water*, 10(8):1049, 8 2018. ISSN 2073-4441. doi: 10.3390/w10081049.

Zhanhan Yu and Ying Shi. Heat and Productivity: Evidence From Flight On-Time Performance, 5 2024.

9 Appendix

9.1 Data on firm balance sheets and income statements (Orbis)

I merge three *Orbis Historical* datasets, and only keep firms that appear in all three: *Identifiers*, with information on the legal IDs of the firms, *Firmographics*, to obtain data on their location and sector, and *Financials History*, for variables related to their income statements and balance sheets.

To begin with, I drop firms with no information on the sector in which they operate, I drop all firms that are not non-financial corporations (NFCs), and only keep firms that are legally registered in Italy. Out of the different legal identifiers available for Italian firms in the *Identifiers* dataset, I use the fiscal code (*codice fiscale*), which has a one-to-one match with Orbis's own identifiers (*bvd-id-number*), and which appears also in BDNCP, and is thus needed to match the firm with the public contracts' data. I thus drop observations with missing fiscal code information.

I then implement some standard cleaning of firm financials, largely following the recommended steps in Kalemli-Ozcan et al. (2023):

- drop observations where total assets, revenues, sales and number of employees are all concurrently missing;
- drop observations with missing balance sheet closing date;
- drop observations if the number of employees is greater than two million;
- drop observations with negative total assets;
- drop whole firms with negative sales and/or negative tangible fixed assets;

- check exact balance sheet consistency, i.e. that for each firm-year observation, $assets = liabilities + equity$.

Duplicate firm observations arising as different balance sheets are submitted on the same day for both subsidiaries and parent companies (*unconsolidated*, *consolidated*) are cleaned (i.e. made unique) following the multi-step algorithm described in Bajgar et al. (2020), who argue that for studies focused on single countries or industries, focusing as much as possible on unconsolidated accounts is preferable, so as to avoid duplicating accounts of subsidiaries within those of the parent company. Further duplicates are dropped by choosing observations from local registry filings over annual reports.

Duplicate firm observations arising as different balance sheets are submitted on different days of the same year may arise due to "the presence of both quarterly and annual reports" and to "firms switching from presenting their end of accounting year balance sheet information in one month to some other month" (Kalemli-Ozcan et al., 2023). To clean these I follow two steps, largely following Kalemli-Ozcan et al. (2023). Firstly, when one of the duplicate observations is from the month of December, I drop any non-December duplicate observations, so as to prioritise end-of-year financials. Secondly, when none or all of the duplicates are from the month of December, I pick the observation with the highest value of firm sales, since this is likely to be that coming from a yearly report, while those with lower reported sales are likely to come from quarterly reports.

After these cleaning steps, 98.8% of my observations come from balance sheets submitted on December 31st, right at the end of the year, and I thus choose not to make specific assumptions about the non-end-of-the-year observations. Conversely, Kalemli-Ozcan et al. (2023) are confronted with more of these, and therefore assign balance sheet data preceding June 1st to the previous year, while keeping the post-June 1st ones to the current year. This is an artificial reallocation I do not conduct.

In line with Kalemli-Ozcan et al. (2023), nominal variables are deflated using World Bank GDP deflators (2015 is the base year for Italy).

9.2 Data on public contracts (BDNCP)

I only keep public contracts with:

- known start and end dates;
- auctions successfully won by one or more firms;
- a strictly positive contract value for the winning firm(s);
- as contract type public works (infrastructure work) rather than procurement of supplies or service contracts;
- known location and identity of the procuring entity;
- known identity of the winning firm(s).

I treat public contracts as ongoing between their *de jure* start date and *de facto* end date. In cleaning the data on contract suspensions, I only keep suspension episodes with known start and end dates, and I require the latter to be greater than or equal to the former.

9.3 Further results: indirect financial effects

Figure 12: log(Working Capital)

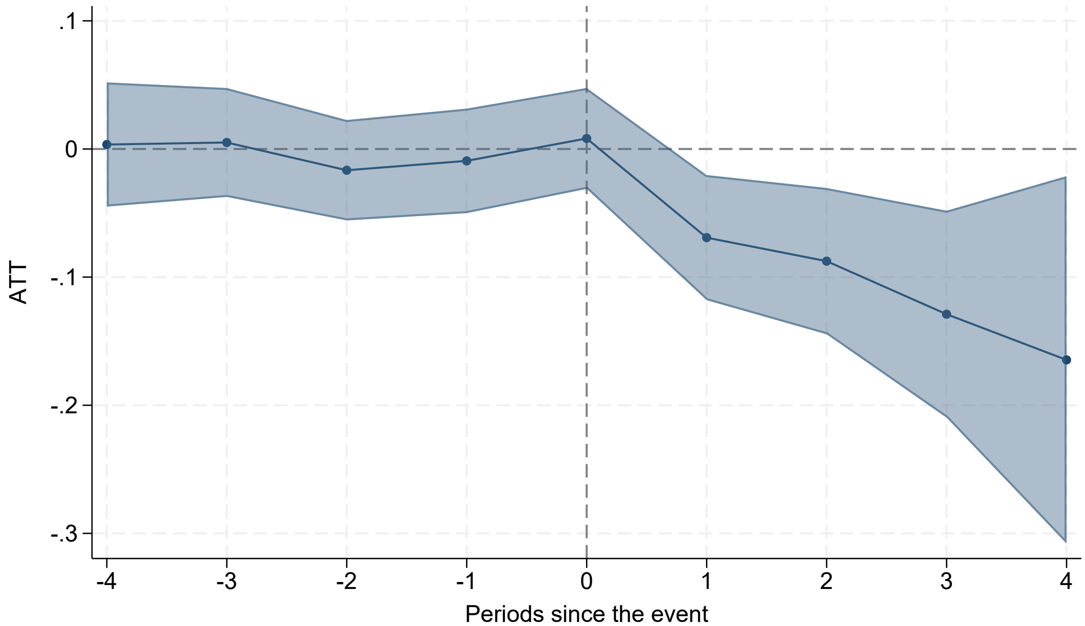

Note: all covariates, wide event study sample

Figure 13: Investment

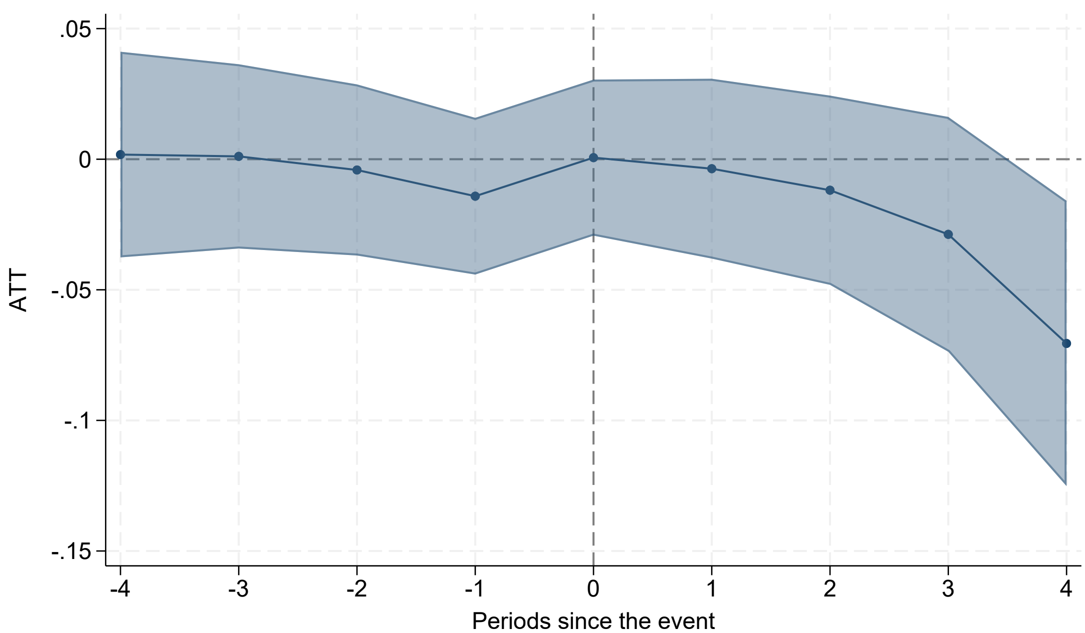

Note: all covariates, wide event study sample

Figure 14: log(Leverage)

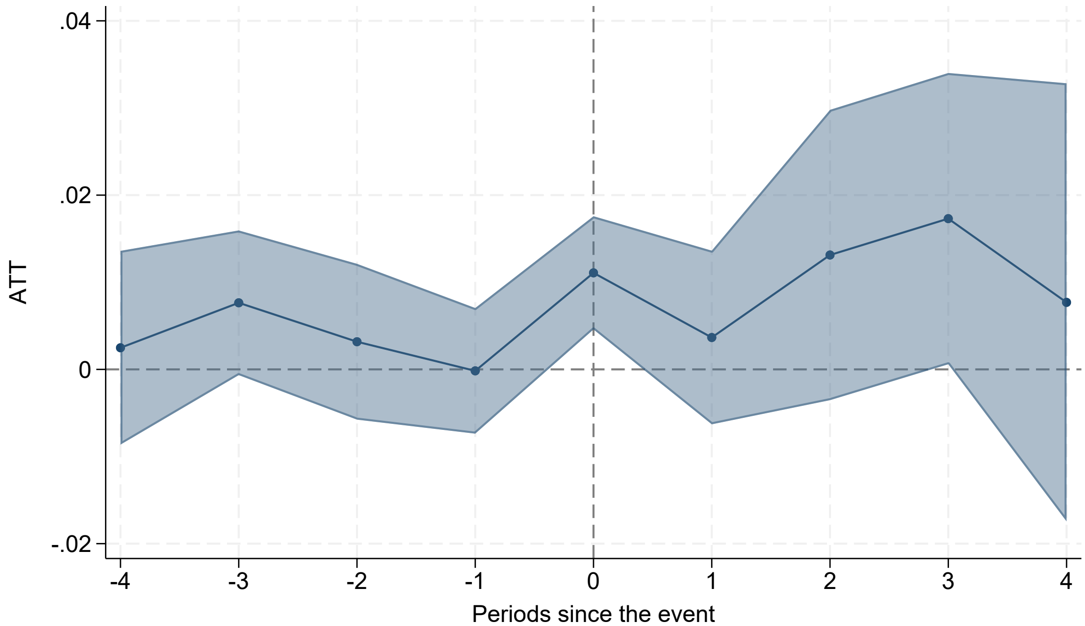

Note: all covariates, wide event study sample

Figure 15: log(Debt Maturity)

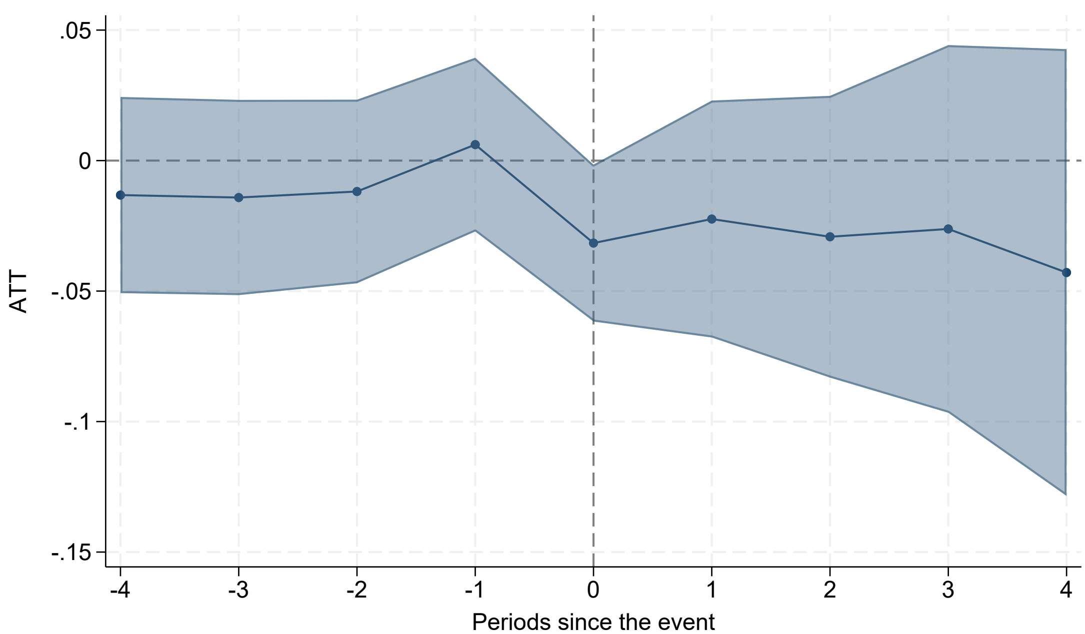

Note: all covariates, wide event study sample

9.4 Extension: checking for physical damage

In this section, I provide evidence that the large, negative effects of weather suspensions on firms are not driven by weather events inducing greater physical damage to the productive capacity of treated firms compared to control ones. In figures 16 and 17 I adopt the same matched DiD design as in section 6, and apply it to variables related to the physical assets held by firms.

I argue that were a firm to be subject to great physical damage due to the weather event underlying the suspension, this would be reflected in her greater demand for materials to rebuild her facilities, and in an immediate drop in the balance sheet value of her (newly damaged) fixed assets.

Instead, material costs follow exactly the same evolution as firm sales (figure 16): firstly, a small jump on the suspension year, as automatic payments of the client are received as compensation for the work completed by the firm up to the suspension start. Secondly, a steadily growing decline, continuing into the fourth year after the first contract suspension. Tangible fixed assets display a very similar pattern (figure 17), dominated by the steady down-scaling of the firm as the suspension effects reverberate. The time zero drop in fixed assets is only marginally significant, and negligible compared to the later drops undergone by this variable.

Overall, this evidence indicates that the liquidity effects of suspensions dominate the responses of firms, and that the worry that physical damage due to weather conditions might be driving this result is unfounded.

Figure 16: log(Material Costs)

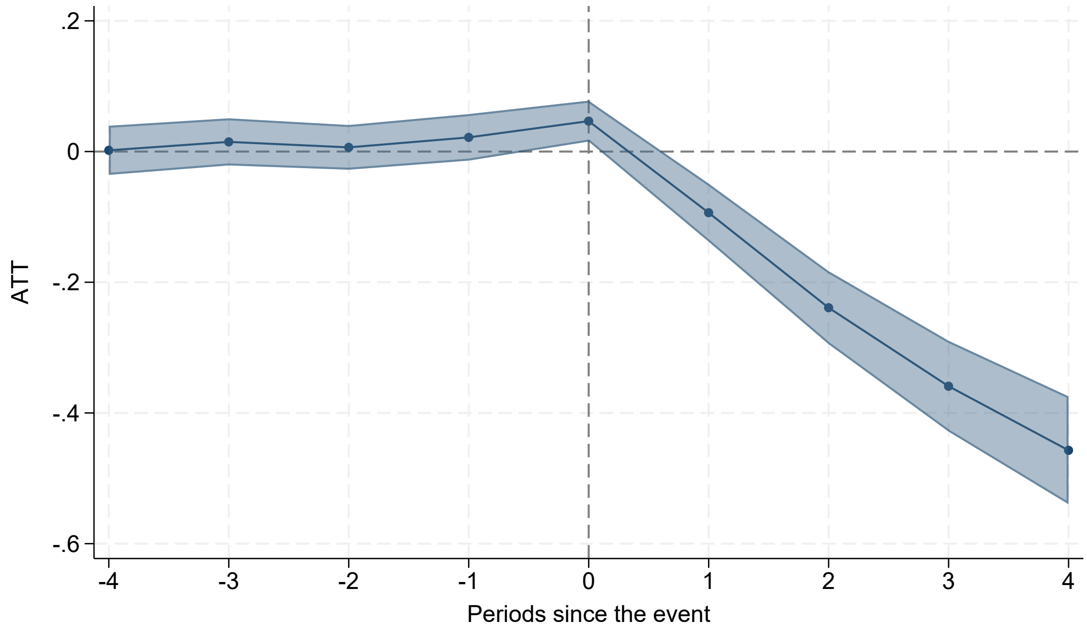

Note: all covariates, wide event study sample

Figure 17: log(Tangible Fixed Assets)

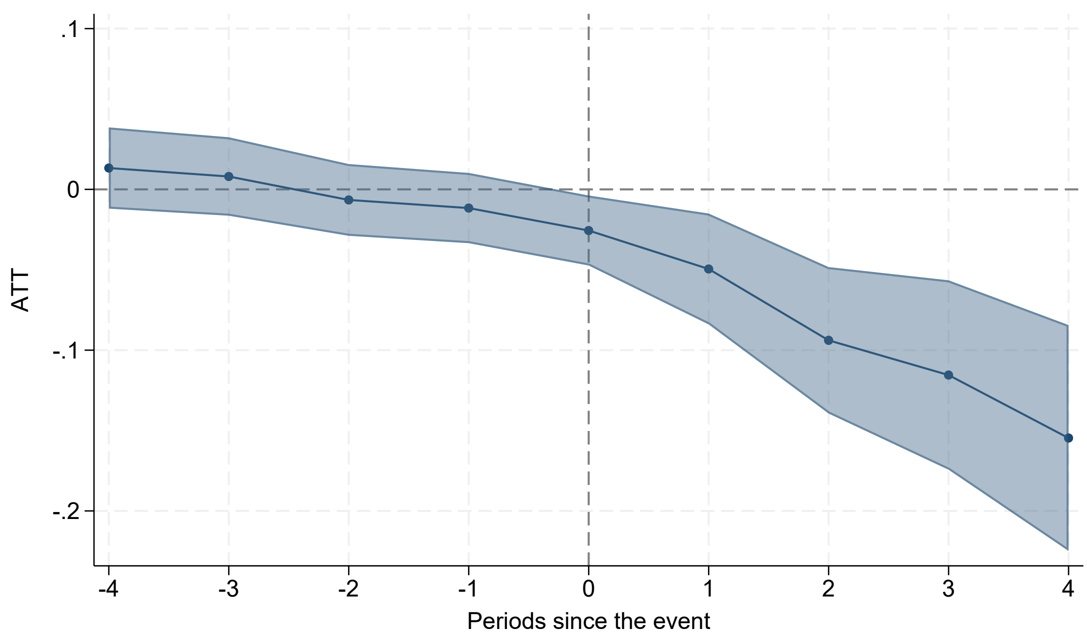

Note: all covariates, wide event study sample

9.5 Robustness: DiD plots for the narrow event study sample

Figure 18: log(Sales)

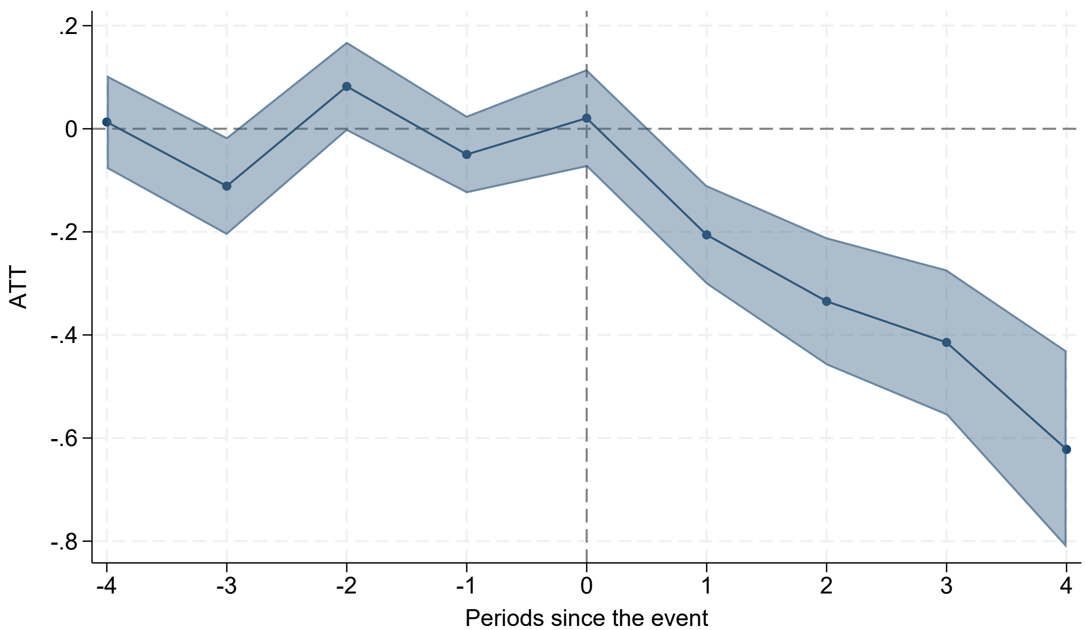

Note: all covariates, narrow event study sample

Figure 19: log(Number of Employees)

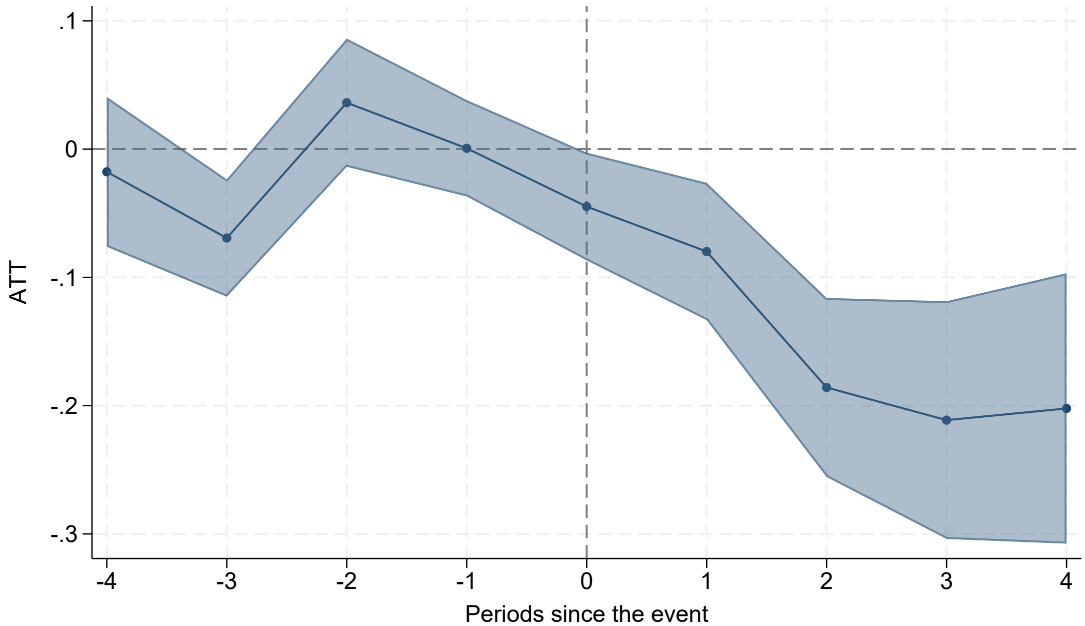

Note: all covariates, narrow event study sample

Figure 20: log(Total Assets)

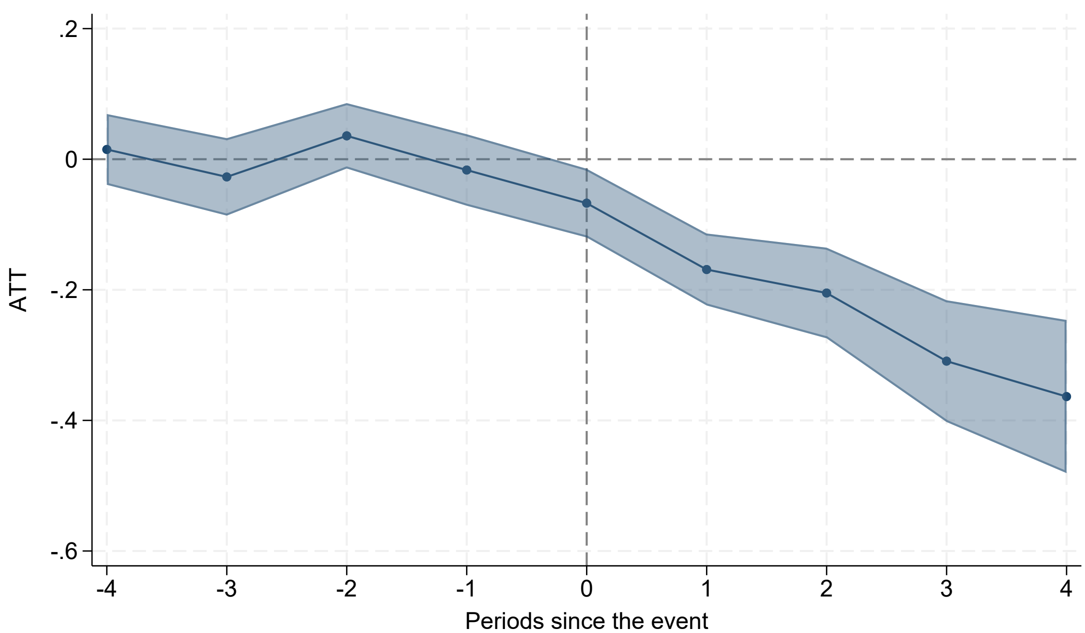

Note: all covariates, narrow event study sample