RIETI Discussion Paper Series 26-E-067

Long-term Scarring or Emerging Recovery? The decade-long trajectory of the GEJE

IZUMI, Yutaro
GRIPS

NAKATA, Hiroyuki
RIETI

SAWADA, Yasuyuki
University of Tokyo

SEKIGUCHI, Kunio
RIETI

RIETI
Research Institute of Economy, Trade & Industry, IAA

The Research Institute of Economy, Trade and Industry
https://www.rieti.go.jp/en/

RIETI Discussion Paper Series 26-E-067
September 2026

Long-Term Scarring or Emerging Recovery? The decade-long trajectory of the GEJE[1]

Yutaro IZUMI (GRIPS), Hiroyuki NAKATA (University of Tokyo), Yasuyuki SAWADA (University of Tokyo), Kunio SEKIGUCHI (Cabinet Office)

Abstract

This paper evaluates the long-term economic impacts of the 2011 Great East Japan Earthquake (GEJE) on small and medium-sized enterprises (SMEs). We combine a proprietary firm survey, which measures disaster exposure, with an 18-year panel of financial statements (fiscal 2007--2024), which measures outcomes. Firms that report having experienced the earthquake or tsunami fall durably behind comparable respondents in the scale of their activity and exhibit the following: annual sales among affected firms average roughly a quarter below control firms over the fourteen post-disaster years, the gap widens rather than closes over time, and it persists through 2024 with no sign of reversal. The scarring is concentrated outside the construction sector, where reconstruction demand was not a significant factor. Exposure assigned based on geocoded physical hazard data and maps or from firm location instead of the self-report yields no evidence of scarring, suggesting that location-based measures conflate damage with presence in the reconstruction zone. The findings indicate that post-disaster trajectories unfold over a multi-decadal horizon, which is well beyond standard three-to-five-year policy frameworks.

Keywords: Disaster resilience; Economic scarring; Small and medium-sized enterprises (SMEs); Disasters triggered by natural hazards.

JEL classification : Q54, L25, R11, D22

[1]This study is conducted as a part of the Project "SME Resilience to Disasters: An empirical investigation" undertaken at the Research Institute of Economy, Trade and Industry (RIETI).
The draft of this paper was presented at the RIETI DP seminar for the paper. We would like to thank participants of the RIETI DP Seminar for their helpful comments.

1. Introduction

As climate change has moved to the center of the policy agenda, a growing body of research in economics and related disciplines has investigated disaster resilience in the context of climate change adaptation (Carvalho et al., 2021; Noy and duPont IV, 2018; Carleton et al., 2024). While a well-established literature examines the impacts of natural hazards and man-made disasters over different horizons, evidence on long-term recovery is mixed: spatial economic structures typically exhibit resilience to physical damage, whereas subsequent economic and political shocks often trigger significant reconfigurations (Glaeser, 2022; Takeda and Yamagishi, 2026). Empirical evidence that employs granular microdata to capture the long-run dynamics of resilience, however, remains scarce: existing firm-level assessments are heavily skewed toward short- or medium-term windows of three to five years, leaving the multi-decadal evolution of disaster shocks poorly understood.

This study fills this gap for the Great East Japan Earthquake (hereafter, GEJE) of March 11, 2011, which exerted a severe and prolonged shock on the Japanese economy: production fell sharply not only in the directly devastated Tōhoku region but nationwide, as supply chain disruptions propagated through the production network to firms far removed from the disaster zone. The GEJE has accordingly generated a large multi-disciplinary literature, reviewed below, spanning macroeconomic growth, household welfare, and production networks. Yet firm-level resilience and recovery remain under-investigated, particularly over long horizons. To the best of our knowledge, no existing study tracks the economic consequences of the GEJE at the firm level for more than a decade after the event.

We track affected small and medium-sized enterprises (SMEs) for fourteen post-disaster years by combining two data sources with distinct roles: a proprietary RIETI survey fielded in early 2025, whose disaster-experience module defines exposure, and the Tokyo Shoko Research (TSR) database of annual financial statements for fiscal 2007–2024, which measures every outcome. Our treatment is therefore self-reported exposure; Section 4 characterizes what that measure captures relative to exposure assigned by firm location. The central finding is persistent scarring: firms that report having experienced the March 2011 earthquake or tsunami lose on average 447 million yen in annual sales relative to comparable respondents (about 27 percent of the pre-disaster control mean), with parallel losses in assets and fixed assets, and the gap deepens over the post-disaster horizon rather than closing. We find no significant pre-trends, supporting the parallel-trends assumption. The post-2011 losses are concentrated among firms outside the construction sector, which reconstruction demand did not reach: non-construction treated firms decline significantly in proportional as well as absolute terms, while construction firms show no significant change. The losses are also concentrated among larger treated firms, while the difference between larger and smaller firms is not precisely estimated. We find no recovery within the sample period.

A natural question is how far these estimates depend on the way exposure is measured. Our treatment indicator is the firm's own report, whereas much of the disaster literature assigns exposure from firm location (Carvalho et al., 2021). Taking exposure instead from geocoded physical hazard or from the officially designated disaster area, we find no evidence of scarring. We interpret the contrast as reflecting what each measure identifies: nearly all physically exposed firms are located within the designated disaster area, where reconstruction spending was concentrated, so exposure defined by location might capture the reach of the reconstruction program along with physical damage, while the self-report identifies operational disruption that a geocoded headquarters address does not reveal.

This paper contributes to the literature on the economic consequences of natural hazards, which has documented effects at several levels of aggregation. At the macro level, catastrophic disasters need not depress long-run growth (Skidmore and Toya, 2002; Cavallo et al., 2013); at the household and individual level, the GEJE left measurable and persistent marks on well-being, preferences, and behavior (Rehdanz et al., 2015; Kuroishi and Sawada, 2024; Sawada et al., 2025), part of a broader research program on the economics of disasters and risk (Kunreuther, 1978; Camerer and Kunreuther, 1989; Sawada, 2007; Kellenberg and Mobarak, 2011; Chuang and Schechter, 2015; Schildberg-Hörisch, 2018; Sawada, 2022; Haushofer and Salicath, 2023). At the firm and network level, the GEJE became the canonical natural experiment for shock propagation through production networks: supply chain disruptions cascaded both upstream and downstream, and their aggregate cost — estimated at 0.35 to 0.47 percentage points of GDP growth — exceeded the direct physical damage (Carvalho et al., 2021; Inoue and Todo, 2019; Tokui et al., 2017). Recent micro-level evidence likewise points to prolonged negative hysteresis rather than creative destruction (Leiter et al., 2009), and studies of the reconstruction period document both the stabilizing role of relief programs and the frictions that public reconstruction demand introduced into local markets (Kameda et al., 2021; Miyazaki, 2025). Within this literature, we provide a fourteen-year, firm-level tracking of disaster scarring, together with an explicit account of what self-reported exposure measures relative to exposure assigned from a firm's location.

The rest of this paper is structured as follows. Section 2 describes the setting, the two data sources, the construction of the estimation samples, and the econometric framework. Section 3 reports the results: the average post-disaster effect, the event-study trajectory, the heterogeneity by industry and firm size (Section 3.3), and the robustness checks (Section 3.4). Section 4 discusses what the self-reported treatment measures and re-estimates the design under alternative treatment definitions, assigning exposure instead from geocoded physical hazard and from the officially designated disaster area. Section 5 concludes.

2. Data and empirical strategy

2.1. 2011 Great East Japan Earthquake (GEJE)

The Great East Japan Earthquake of March 11, 2011 (moment magnitude 9.0) was the largest earthquake instrumentally recorded in Japan. Strong ground shaking extended over most of eastern Japan, and the tsunami it triggered inundated some 560 km^2 of the Pacific coast across six prefectures, destroying whole coastal districts together with the ports, plants, and small firms clustered in them. The tsunami also disabled the Fukushima Dai-ichi nuclear power station, causing a severe accident that displaced the surrounding population under evacuation orders, some of which remain in force today. More than 18,000 people were killed or went missing (National Police Agency, 2026), and the Cabinet Office valued the direct damage to the capital stock at approximately 16.9 trillion yen — on the order of 3.5 percent of Japan's GDP at the time, and close to twice its corresponding estimate for the 1995 Kobe earthquake (Cabinet Office, 2011).[1]

2.2. Data

We estimate the effect of the GEJE with a difference-in-differences design, comparing firms exposed to the disaster with unexposed firms before and after March 2011; Section 2.4 states the specification. The design therefore turns on how exposure is defined, and our two data sources play distinct roles in it: a firm survey supplies the treatment indicator, and a registry of annual financial statements supplies every outcome.

Treatment: Self-reported exposure to GEJE

Our treatment is a self-report: a firm is *treated* if it reports having experienced the March 2011 earthquake or tsunami at either its primary or secondary establishment. The report comes from the disaster-experience module of the "Survey on Small and Medium-Sized Enterprises' Disaster Countermeasures" (hereafter, the RIETI survey), whose fiscal year 2024 wave was administered by Tokyo Shoko Research (TSR) on behalf of the Research Institute of Economy, Trade and Industry (RIETI) between January and March 2025. In the module, reproduced in Appendix A.1.2, each respondent reports up to five disaster episodes per establishment over the preceding fifteen years, recording each episode's type, year and month, write-in amounts of direct and indirect damage, and yes/no items for tangible-asset damage and operation suspension. A firm is treated if any episode, at either establishment, records an earthquake or tsunami dated March 2011; the indicator

[1] The police tabulation counts deaths directly attributed to the earthquake and tsunami; the Reconstruction Agency additionally recognizes some 3,800 disaster-related deaths among evacuees in the years that followed. The Cabinet Office estimate covers buildings, lifelines, social infrastructure, and agricultural, forestry, and fishery assets, and excludes damage from the Fukushima Dai-ichi accident; the share of GDP is our own calculation against 2010 nominal GDP.

is defined at the firm level and is time-invariant, and by this definition 343 of the 3,510 respondent firms (9.8%) are treated, with the remaining respondents serving as controls. Since the damage items are recorded separately from the episode itself, a firm can report the March 2011 event while marking no damage category: 166 of the 343 write in a positive direct-damage amount, 128 a positive indirect-damage amount, 157 flag tangible-asset damage, and 135 operation suspension (Appendix Figure A.4).[2]

Treatment thus conditions on the firm's own account of having been affected, whether through physical destruction, operational disruption, or losses transmitted through suppliers and customers. A location-based alternative, classifying as treated every firm inside the strong-shaking zone or the inundation footprint, would draw the line differently: it would include firms that experienced intense shaking but suffered little damage, and it would miss firms harmed at establishments or through business relationships that a geocoded address cannot see. The control group under our definition accordingly includes many firms that felt the earthquake but did not report it as a disaster experience. Appendix Figure A.2 maps the treated and control firms in the estimation samples, and Section 4 discusses what each exposure measure captures and re-estimates the design under alternative treatment definitions on the same firms.

The RIETI survey was designed to measure SMEs' disaster exposure, countermeasures, and resilience in general, and the module asks about eleven disaster types over the preceding fifteen years. Two implications follow. For sampling, the roster of 15,000 SMEs is purposive rather than a random draw from the register: half is drawn from firms certified under the Small and Medium Enterprise Agency's Business Continuity Enhancement Plan certification scheme[3] and half from the TSR register at large, matched on industry, with further targeting of prefectures containing municipalities affected by officially designated severe disasters and oversampling of firms near the size thresholds of the industry-specific SME definitions. The roster is therefore purposive along four dimensions — certification status, industry, geography, and firm size — and the external validity of the estimates should be read with that design in mind; Appendix A.1.1 details the sampling frame and fielding procedure. For measurement, respondents were not prompted to recall the GEJE in particular, so a March 2011 report reflects the firm's unprompted ranking of that event among its disaster experiences. 3,510 valid responses were returned, achieving a response rate of 23.4%; these constitute the respondent sample.

Finally, as the survey was fielded in early 2025, the indicator rests on a recall of nearly fourteen years. Imperfect recall and the five-event cap may suppress true 2011 reports,[4]

[2]Establishments are those at the time of the survey, so a firm that relocated after 2011 reports the event as experienced at its then-current location. We measure such relocation from the head-office addresses in successive vintages of the TSR registry and re-estimates the design excluding relocated firms.

[3]The certification scheme, introduced by the Small and Medium Enterprise Agency in 2019, certifies a simplified form of business continuity plan (BCP) designed for smaller firms.

[4]The cap almost never binds: only 15 respondents (0.4% of the 3,510 usable responses) date five

while the reverse error, an unaffected firm misdating another event to March 2011,[5] is made rare by the disaster's salience. The plausible errors therefore run mostly toward placing genuinely affected firms in the control group, and with a binary treatment such misclassification shrinks the treated–control contrast, so we read the measurement error as attenuating.

Firm Outcomes: TSR Database

Every outcome is measured in the TSR database, which gives annual financial statements for fiscal years 2007–2024, and which also provides the industry classification used throughout the paper. The headline outcome variables are total sales, the number of employees, total assets, and total fixed assets; Appendix Table A.1 gives their definitions.[6]

We designate these four as the headline outcomes because they are the financial-statement items least exposed to reporting discretion. Privately held firms report primarily for tax purposes and face well-documented incentives to manage book income downward; in Japan, SME owner-managers commonly extract earnings as executive compensation, so that reported profit is depressed through discretionary costs even at otherwise healthy firms. Profit-based measures therefore reflect tax-motivated reporting choices as well as economic performance. Sales and balance-sheet stocks are far harder to understate: revenue is anchored by transactions with counterparties and cross-checked through consumption-tax filings, and assets accumulate from recorded acquisitions. Profits and the main financing positions of the balance sheet (total liabilities, net worth, and long-term debt) are examined in the additional analyses of Appendix B.4, which read the profit outcomes with this caveat in mind. firm-year, either the financial statement is available and all statement variables are recorded, or none is, so the panel-completeness requirement of the next section screens on the availability of the statement itself rather than on any particular outcome.

Estimation Sample

Figure 1 traces the construction of the estimation samples: out of 3,510 respondents, 2,157 can be linked to at least one year of TSR financial statements; of these, 103 firms

disaster events at their primary establishment. Of these, four report the March 2011 event among the five, and only one dates all five before March 2011, the sole case in which a further 2011 event mechanically could not have been reported; that firm does not enter either estimation sample.

[5] The events most plausibly conflated with the mainshock are the April 7, 2011 aftershock off Miyagi (magnitude 7.2) and the earthquakes off the coast of Fukushima on February 13, 2021 (magnitude 7.3) and March 16, 2022 (magnitude 7.4), each of which reached a maximum JMA intensity of 6+ in parts of Miyagi and Fukushima, the strongest shaking in the region since March 2011.

[6] TSR statements are assigned to the fiscal year in which the accounting period ends, and closing months differ across firms, so fiscal 2011 contains between roughly nine and twenty-one post-disaster months depending on the firm's year-end. Fiscal 2011 is therefore only partially treated for some firms, which attenuates the event-study coefficient at $k = 0$ (Section 2.4).

Figure 1: Construction of the Estimation Samples

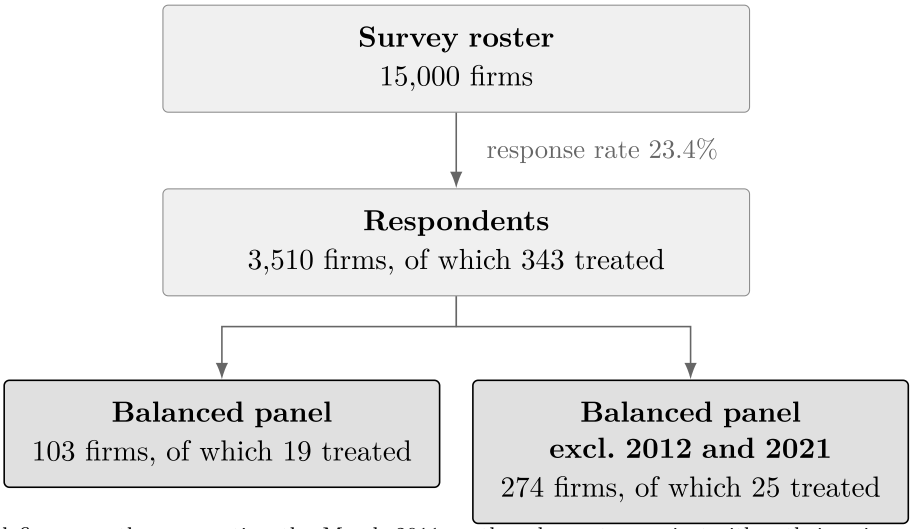

Notes: Treated firms are those reporting the March 2011 earthquake or tsunami at either their primary or secondary establishment in the survey's disaster-experience module. Of the 15,000 firms on the survey roster, 3,510 returned the questionnaire with a usable firm identifier (a 23.4 percent response rate) and form the respondent sample; 2,157 of these can be linked to at least one fiscal year of TSR financial statements. The balanced panel comprises firms with financial statements for every fiscal year 2007–2024. The second panel drops the 2012 and 2021 fiscal years, whose TSR data vintages have known coverage discontinuities, and requires financial statements in every one of the sixteen remaining years, a weaker requirement, so the panel is larger and the balanced panel is a subset of it. Financial-statement availability is evaluated jointly for the four outcome variables (total sales, employees, total assets, and total fixed assets) and for sales per employee, the fifth financial variable used in the sample screening.

are observed in every fiscal year 2007–2024 (the *balanced panel*; 19 treated). We require a balanced panel so that the dynamic estimates of Section 2.4 track a fixed set of firms: with entry and exit excluded by construction, changes in the estimated treatment–control difference over the event horizon reflect within-firm trajectories rather than a changing sample composition. The cost of this requirement is that the estimates describe firms that survived and remained continuously observable throughout. We examine firm exit directly and asks whether this survival conditioning selects treated and control firms differently. The natural concern is that the most severely damaged firms are the ones that drop out of continuous observation, so that the balanced panels would retain the treated firms that fared best and the estimated losses would be a lower bound on that account. The registry-based evidence points the other way: physically exposed firms were *less* likely to exit after 2011 than observably similar unexposed firms — a retention effect that is strongest in 2011–2012 and among the smallest firms, and that survives a battery of checks against mechanical explanations — consistent with reconstruction demand, insurance payouts, and disaster-related lending keeping affected firms in operation. The survival conditioning is therefore unlikely to have purged the treatment group of its worst-hit members; if anything, the panels may retain treated firms that public and private support kept alive, which works against finding losses on the intensive margin rather than manufacturing them.

A major limitation of the TSR data interacts with this balancing requirement. The

2012 and 2021 data vintages are missing a significant share of firms; the gap is a discontinuity in the coverage of the data files themselves rather than a wave of exits, and it is concentrated in western Japan, whose firms largely disappear from those two vintages while coverage of the Tohoku region and the Tokyo area is spared. Any sample that requires continuous observation through 2012 and 2021 therefore inherits a geographic skew from the data, and this is precisely what happens to the balanced panel: its control group concentrates in eastern Japan (Appendix Figure A.2), and its treated share rises to 18.4 percent, well above the 9.8 percent of the respondent pool, because treated firms cluster in the regions the two vintages spare (Appendix Figure A.3). Our second estimation sample therefore drops the two years and reconstructs the balance requirement within the sixteen remaining fiscal years: 274 firms (25 treated), of which the balanced panel is a subset. This panel restores control firms across western Japan and carries a treated share (9.1 percent) that mirrors the respondent pool, so the two samples bracket the trade-off between temporal completeness and geographic representativeness. Nevertheless, the estimation results do not differ substantially between the two panels; thus, We only report the estimation results using the panel including 2012 and 2021 in the main text, and report the results using the panel excluding the two years in online Appendix B.1.

2.3. Descriptive Statistics and Trends

Table 1 compares treated and control firms before the disaster. It reports pre-GEJE (2007–2010) means of sales, employment, assets, and fixed assets, together with the industry composition in four coarse groups, separately for the balanced panel (19 treated, 84 control firms) and the balanced panel excluding 2012 and 2021 (25 treated, 249 control firms). Treated firms are somewhat smaller than controls on every financial measure, but no difference is statistically distinguishable: the p-values from two-sample t-tests exceed conventional significance thresholds throughout. We note that with 19 to 25 treated firms these tests have limited power, and that our design does not rest on balance in levels: firm fixed effects absorb any time-invariant difference between the groups, and identification comes from the parallel-trends assumption, assessed visually below and formally through the pre-trend tests of the event studies in Section 3. The table also shows that the panel firms are substantially larger than the typical Japanese SME, pre-disaster mean sales of 1.6 to 2.1 billion yen and 40 to 50 employees, reflecting the selection toward larger firms at each screening margin, documented in Appendix A.3 and a feature that bears on the size heterogeneity of Section 3.3.2.

Table 1 also shows how the two estimation samples differ from each other. The larger panel excluding 2012 and 2021 skews toward somewhat larger firms — control-group mean sales of 2.1 billion yen against 1.6 billion in the balanced panel, with employees, assets,

Table 1. Summary Statistics by Sample and Treatment Status, Pre-GEJE Period (2007–2010)

	Treatment			Control			
	Firms	Mean	SD	Firms	Mean	SD	p-value
Balanced panel							
Total sales	19	1,662.1	2,354.0	84	1,643.1	2,253.9	0.974
Number of employees	19	36.8	50.3	84	42.1	74.4	0.771
Total assets	19	1,101.3	1,649.4	84	1,213.4	2,023.9	0.818
Fixed assets	19	427.3	756.0	84	388.1	613.3	0.811
Industry share (2010 classification)							
Construction	19	0.47	0.51	84	0.60	0.49	0.338
Manufacturing	19	0.37	0.50	84	0.18	0.39	0.069
Wholesale and retail trade	19	0.11	0.32	84	0.15	0.36	0.585
Other industries	19	0.05	0.23	84	0.07	0.26	0.771
Balanced panel excluding 2012 and 2021							
Total sales	25	1,756.8	2,211.4	249	2,090.2	3,316.4	0.619
Number of employees	25	44.6	50.9	249	50.2	65.9	0.679
Total assets	25	1,475.4	1,913.3	249	1,668.1	2,784.4	0.732
Fixed assets	25	633.9	926.1	249	709.0	1,658.4	0.822
Industry share (2010 classification)							
Construction	25	0.44	0.51	249	0.45	0.50	0.926
Manufacturing	25	0.44	0.51	249	0.27	0.44	0.071
Wholesale and retail trade	25	0.08	0.28	249	0.18	0.39	0.190
Other industries	25	0.04	0.20	249	0.10	0.30	0.352

Notes: Pre-GEJE (2007–2010) statistics by treatment status. For the financial variables, Firms is the number of distinct firms with at least one non-missing pre-GEJE observation, and means and SDs are over pre-GEJE firm-year observations. Industry shares are firm-level shares of each group, with industry measured by the large industry classification pinned to the firm's last pre-treatment (2010) classification; less frequent divisions and firms with no usable classification are grouped into Other industries. The p-values are from two-sample t-tests at the firm level (pre-GEJE firm means for the financial variables, group indicators for the industry shares) and test the null of equal treated and control means. All monetary variables are measured in 1 million yen.

and fixed assets higher in step — and toward manufacturing: the manufacturing share rises from 37 to 44 percent among treated firms and from 18 to 27 percent among controls, while the construction share of the control group falls from 60 to 45 percent. Recalling that the two samples also differ in geography and treated share, they should be read as drawn from overlapping but distinct populations of respondents, and the compositional differences are one reason their estimates can differ.

The industry rows show one marginal exception: in both panels, treated firms are more often manufacturers than controls (37 versus 18 percent on the balanced panel, $p = 0.07$; 44 versus 27 percent excluding 2012 and 2021, $p = 0.07$). This imbalance matters because manufacturing was differentially exposed to post-2011 aggregate shocks unrelated to a firm's own disaster experience, the nationwide supply-chain disruptions and the sharp yen appreciation of 2011–12, which the construction-by-year fixed effects of the main specification do not absorb. Section 3.4 therefore re-estimates the DD with manufacturing-by-year in place of construction-by-year fixed effects and reports a manu-

Figure 2: Industry composition by sample

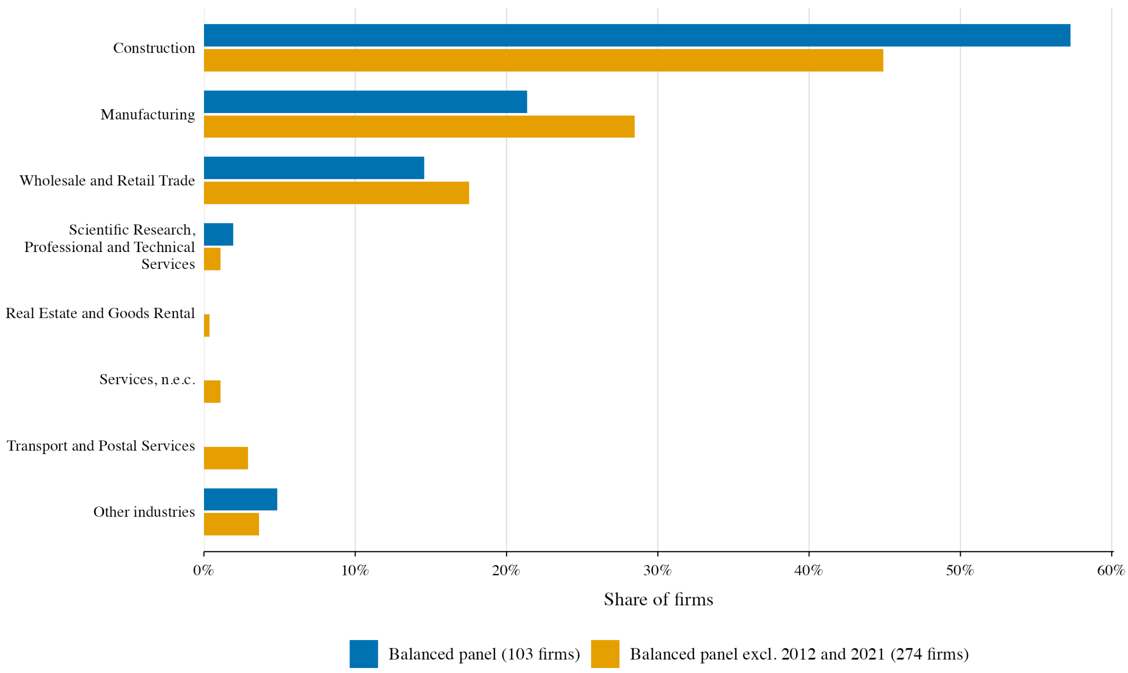

Notes: Share of firms in each industry division, for the balanced panel (103 firms) and the balanced panel excluding 2012 and 2021 (274 firms) side by side. Industries are measured at the firm level using the large industry classification, pinned to the firm's last pre-treatment (2010) classification — the same definition as the construction indicator in the main specification. The eight most frequent divisions are shown individually; less frequent divisions and firms with no usable classification are grouped into "Other industries."

facturing versus non-manufacturing split of the treatment effect, and Section 3.3 examines the construction margin, the industry dimension most directly tied to the reconstruction boom, in the main text.

Figure 2 reports the finer industry composition of the two estimation samples. The balanced panel has a higher share of construction firms and a lower share of manufacturing firms than the panel excluding 2012 and 2021. The construction share matters in its own right, because reconstruction demand after the GEJE supported the construction sector (Appendix A.4 documents the size and duration of the reconstruction boom): differences in composition across the two samples are one reason their estimates can differ, and they motivate the construction-by-year fixed effects of the main specification as well as the construction split of Section 3.3.

Figure 3 plots yearly mean outcomes by treatment group on the balanced panel.[7] Before 2011, the two groups move roughly in parallel. After the disaster, the control group grows steadily on every outcome, while treated firms remain close to their pre-disaster

[7]Monetary variables are nominal throughout; the year fixed effects in the estimation absorb price movements at the national level, which are common to both groups in any year, though not price movements specific to a region.

Figure 3: Trends in mean outcomes for treatment and control groups: *Balanced panel*

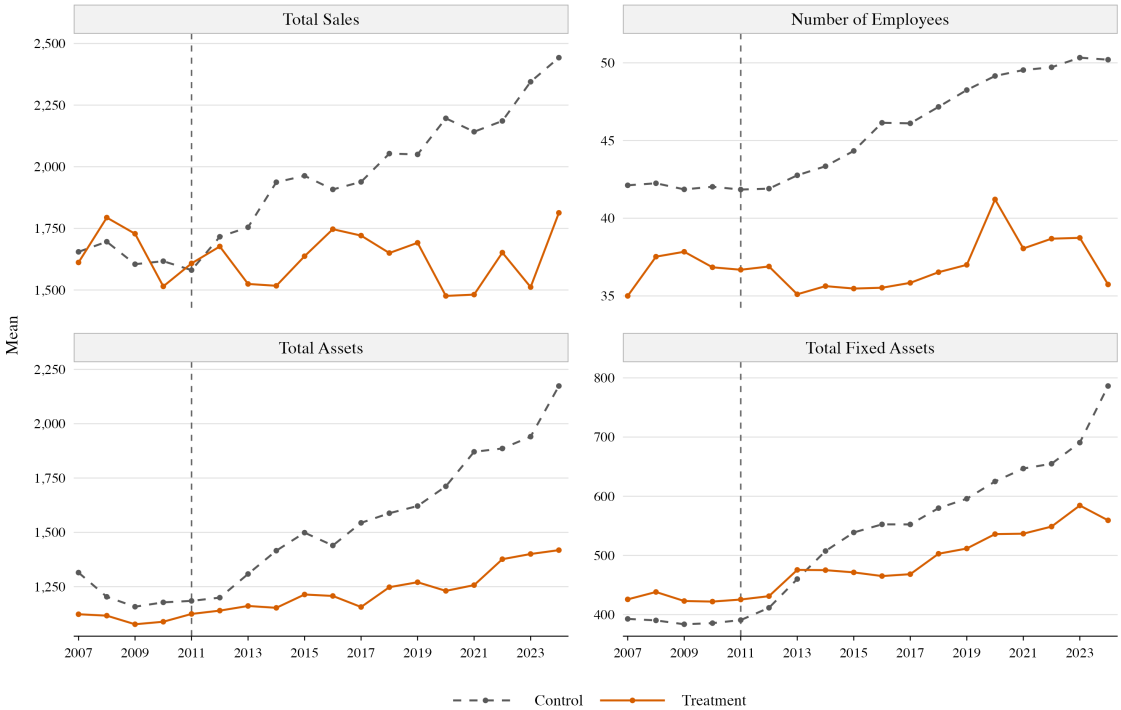

Notes: Yearly means by treatment group, balanced panel (103 firms, of which 19 treated). The dashed vertical line marks 2011, the year of the GEJE. Monetary variables are measured in 1 million yen; the number of employees is in persons.

level, so the gap widens year by year rather than closing. Appendix Figure A.1 isolates this pattern, plotting the treated − control difference in yearly means relative to its 2010 value: the gap is close to zero through the pre-disaster years, sales move transitorily around the 2008–09 global financial crisis but return to zero by 2010, and then declines nearly monotonically through 2024, with no sign of reversal. The corresponding figures for the balanced panel excluding 2012 and 2021 (Appendix Figure B.1) show the same pattern.

2.4. Empirical Strategy

We begin with the difference-in-differences specification in Equation (1).

$$Y_{it} = \alpha_i + \lambda_{c(i),t} + \beta\,(Treat_i \times Post_t) + \varepsilon_{it} \tag{1}$$

where Y_{it} is an outcome for firm i in year t, $Treat_i$ is a dummy equal to one if firm i reported experiencing an earthquake or tsunami in March 2011 at either its primary or secondary establishment and zero otherwise, and $Post_t$ is a dummy equal to one from 2011 on. The term α_i is a firm fixed effect, absorbing time-invariant differences across firms, including sector and location. The term $\lambda_{c(i),t}$ is a construction-status-by-year fixed effect, where $c(i)$ indicates whether firm i's industry as of 2010 is construction:[8] it allows aggregate shocks to differ between construction and non-construction firms in every year, thereby absorbing the publicly financed reconstruction boom documented in Appendix A.4, which is the first-order confound in this setting, since reconstruction demand may have disproportionately benefited the construction sector, and it subsumes ordinary year fixed effects.[9] Standard errors are clustered at the firm level, allowing for arbitrary serial correlation within a firm over time.[10]

Identification rests on the assumption that, in the absence of the GEJE, the outcomes of treated and control firms would have followed parallel paths within the same sector (construction or non-construction). Two features of the setting support the assumption. First, the disaster was unanticipated, so treated firms' pre-2011 outcomes cannot embed responses to the coming shock; the flat pre-disaster gaps of Section 2.3 and the formal pre-trend tests reported with the event studies in Section 3 assess the assumption empirically. Second, because treatment occurs at a single date and control firms are never treated, the two-way fixed-effects estimator involves no comparisons between earlier- and later-treated units, and the negative-weighting and contamination problems of staggered adoption designs do not arise.[11] Under this assumption, β identifies the average effect of *reported* disaster exposure on treated firms. The relevant population is the respondents that remained continuously observed in the financial data (Section 2.2). Section 4 re-estimates the design with exposure assigned from the firm's location instead, and compares what the two measures capture.

The average specification summarizes the entire post-GEJE period in a single coefficient and therefore does not reveal how the impact evolved over time. To trace the trajectory of the effect, we estimate a fully flexible event-study design,

$$Y_{it} = \alpha_i + \lambda_{c(i),t} + \sum_{k \neq -1} \beta_k \left(\mathbf{1}_k \times Treat_i \right) + \varepsilon_{it} \qquad (2)$$

where $k := t - 2011$ denotes event time, i.e., the number of years relative to the GEJE, running from $k = -4$ (fiscal 2007) to $k = 13$ (fiscal 2024). This specification estimates a separate coefficient for each year, with the year immediately before the disaster ($k = -1$,

[8] Every firm in both estimation samples, and every one of the 3,510 respondents, carries a 2010 industry classification, and the registry industry code is time-invariant in our panel, so the 2010 pinning discards no variation.

[9] The year component of $\lambda_{c(i),t}$ also absorbs price movements at the national level, so the nominal measurement of the monetary outcomes (Section 2.3) does not affect the estimates through aggregate inflation; region-specific price movements are not absorbed.

[10] Since only 19 to 25 firms are treated, cluster-robust inference could over-reject even though the panels contain many clusters (MacKinnon and Webb, 2018); Section 3.4 therefore re-evaluates every headline estimate with a wild cluster bootstrap and with randomization inference.

[11] See, for instance, Goodman-Bacon (2021), Callaway and Sant'Anna (2021), and Sun and Abraham (2021). With a single treatment date and a never-treated control group, the event-study coefficients of Equation (2) are free of the cross-cohort contamination these papers analyze.

Table 2. Average DD Estimates (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t$	−447.338*	−4.799	−275.766*	−136.825
	(236.272)	(3.314)	(151.866)	(93.276)
Observations	1854	1854	1854	1854
Firms	103	103	103	103
Pre-treatment mean	1,646.58	41.09	1,192.75	395.35
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. The estimation sample is the balanced panel. Appendix Table B.1 reports the same specification on the balanced panel excluding 2012 and 2021.

fiscal 2010) omitted as the reference, so that each β_k measures the difference between treated and control firms in event-year k relative to that baseline. The fixed effects and clustering are as in Equation (1).

3. Estimation Results

3.1. Static Difference-in-Differences

Table 2 reports the average difference-in-differences estimates of Equation (1) with outcomes in levels, on the balanced panel. Treated firms lose 447.3 million yen in annual total sales over the post-disaster period, about 27 percent of the pre-treatment mean, and 275.8 million yen in total assets (23 percent), both significant at the 10 percent level. The fixed-asset estimate, −136.8 million yen (35 percent of the pre-treatment mean), is of comparable proportional magnitude but imprecise, and the employment estimate is negative but statistically indistinguishable from zero.[12]

Appendix Table A.3 re-estimates the same specification with the outcomes in logs. None of the estimates is statistically significant. We interpret this discrepancy between the negative level estimates and the null log estimates as reflecting a pooled coefficient that averages over offsetting groups: as Section 3.3 shows, treated non-construction firms decline proportionally as well as in levels, while the coefficients for treated construction firms are insignificant and positive in sign, consistent with reconstruction demand offsetting the disaster's damage in that sector. The contrast also indicates that the yen losses are borne disproportionately by larger treated firms, a pattern Section 3.3.2 examines directly.

[12] Online Appendix B.1 reproduces the tables and figures of Sections 3 and 4 on the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, a larger cross-section of 274 rather than 103 firms, where the estimates are nearly unchanged and more precisely estimated. As stated above, unless stated otherwise, we report results for the balanced panel throughout the main text.

3.2. Dynamic Effects

Figure 4: Dynamic Effects on Firm Outcomes: *Balanced panel*

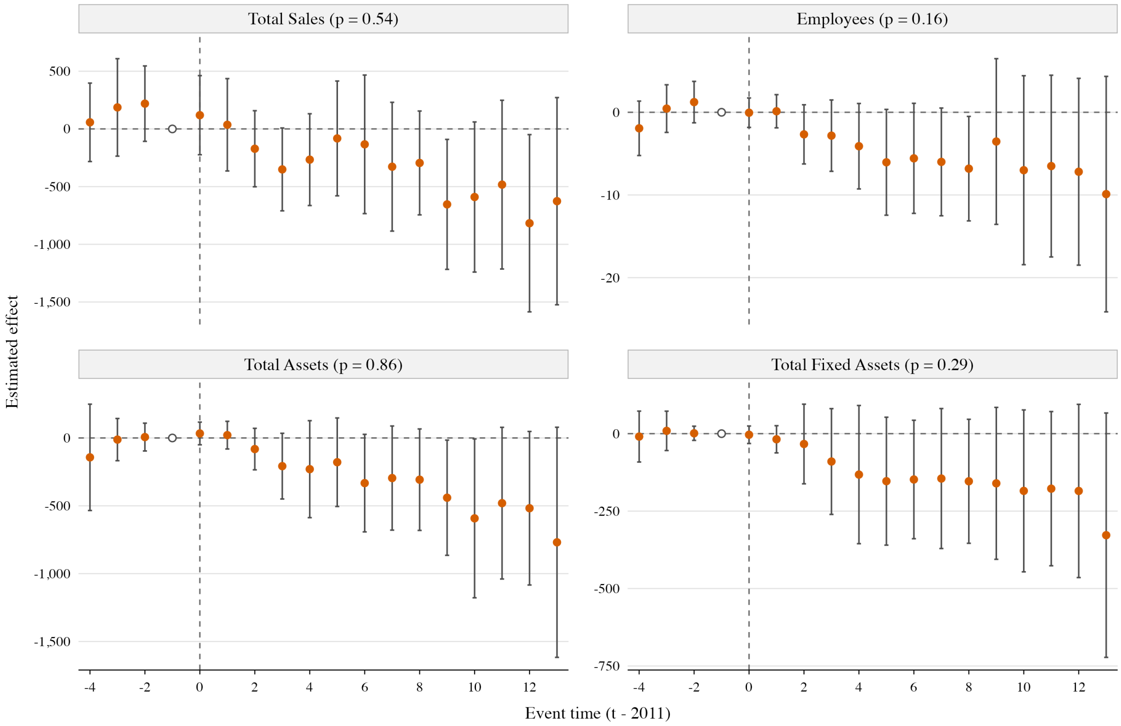

Notes: Event-study estimates from Equation (2), level outcomes, balanced panel (103 firms, of which 19 treated). The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) × year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure 4 reports the event-study estimates of Equation (2) in levels for the balanced panel. Two patterns stand out. First, the pre-treatment coefficients are close to zero and statistically insignificant throughout, and the joint Wald tests do not reject the null of no differential pre-trend for any outcome in either sample ($p = 0.54$, 0.16, 0.86, and 0.29 for total sales, employment, total assets, and fixed assets on the balanced panel, respectively), lending support to the parallel-trends assumption. Second, the post-2011 coefficients are mostly negative and become progressively more so over the event horizon, so the gap between treated and control firms widened rather than narrowed over the full post-disaster horizon.

The impact of the GEJE was not immediate. The coefficients for the first years after the disaster are small and statistically insignificant for every outcome, and two features of the setting bear on this null. Fiscal 2011 is only partially treated for many firms (Section 2.2), which attenuates the earliest coefficients, and relief and reconstruction programs concentrated in the immediate aftermath may have sustained affected firms while they were in place. The losses instead emerge several years after the disaster and deepen thereafter.

Two features of the timing indicate that this deepening is not an artifact of the pandemic. First, the pre-2020 coefficients are already negative and growing in magnitude, so the trajectory is established before COVID-19 rather than by it. Second, the gap continues to widen across the pandemic years themselves rather than appearing as a level shift in 2020.

3.3. Heterogeneity

3.3.1. Construction versus Non-Construction

The average and dynamic results indicate persistent losses among GEJE-treated firms. However, these average effects may mask important heterogeneity across industries. In particular, construction firms may have followed a different post-disaster trajectory from firms in other sectors. On the one hand, construction firms may have directly suffered from physical damage, operational disruption, and local demand shocks. On the other hand, they may also have benefited from reconstruction demand after the earthquake (Appendix A.4). To examine this possibility, we divide firms into construction and non-construction firms and estimate the following specification:

$$\begin{aligned} Y_{it} = \alpha_i + \lambda_{c(i),t} + \beta_C \left(Treat_i \times Post_t \times Construction_i\right) \\ + \beta_N \left(Treat_i \times Post_t \times NonConstruction_i\right) + \varepsilon_{it}, \end{aligned} \tag{3}$$

where $Construction_i$ is a dummy equal to one for firms whose industry as of 2010 is construction and $NonConstruction_i$ is a dummy equal to one for all other firms. The fixed effects are those of the main specification, Equation (1); the construction-by-year cells $\lambda_{c(i),t}$ already allow construction and non-construction firms to follow different average paths in every year, so they absorb the $Post_t \times Construction_i$ main effect, and no separate term is needed.

Table 3 presents the estimates of Equation (3). The results indicate clear heterogeneity between construction and non-construction firms. For construction firms, the treatment coefficient β_C is statistically indistinguishable from zero on every outcome, with positive point estimates throughout — a pattern consistent with post-disaster reconstruction demand offsetting the disaster's damage in this sector.

By contrast, the coefficient for non-construction firms (β_N) is large and sharply negative. Treated non-construction firms lose 1,028.0 million yen in total sales per year (about 41% of the group's pre-treatment mean, significant at the 1% level), 12.2 employees (19%, 1% level), 670.5 million yen in total assets (36%, 1% level), and 338.2 million yen in fixed assets (49%, 5% level). The table reports the p-value of the test that $\beta_C = \beta_N$. The two groups are formally distinguishable on every outcome: $p < 0.01$ for total sales and total

Table 3. Heterogeneous DD Estimates: construction vs. non-construction (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t \times Construction_i$	140.980	2.693	124.155	67.162
	(149.151)	(3.808)	(141.352)	(77.681)
$Treat_i \times Post_t \times NonConstruction_i$	−1028.031***	−12.193***	−670.503***	−338.167**
	(386.909)	(4.613)	(221.199)	(150.272)
p-value: $\beta_C = \beta_N$	< 0.01	0.01	< 0.01	0.02
Observations	1854	1854	1854	1854
Firms (construction)	59	59	59	59
Firms (non-construction)	44	44	44	44
Pre-treatment mean (construction)	1,017.53	23.00	691.99	171.74
Pre-treatment mean (non-construction)	2,490.07	65.35	1,864.22	695.19
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications; the construction × year cells absorb the $Post_t \times Construction_i$ main effect. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment means pool treated and control firm-years 2007–2010 within each group. The p-value row tests $\beta_C = \beta_N$, the equality of the two group coefficients, from the same specification re-parameterised with construction firms as the baseline, so that the non-construction term is the between-group difference. The estimation sample is the balanced panel. Appendix Table B.2 reports the same specification on the balanced panel excluding 2012 and 2021.

assets, $p = 0.01$ for employment, and $p = 0.02$ for total fixed assets. This pronounced divergence indicates that the aggregate scarring documented earlier is driven entirely by the non-construction firms.

The sectoral divergence also appears in proportional terms. Re-estimating Equation (3) in logs (Appendix Table A.4), treated non-construction firms decline by about 25 log points in total sales and 14 log points in employment, both significant at the 5 percent level, while every construction coefficient is positive and insignificant. Where reconstruction demand did not reach, the losses are proportional as well as absolute, and the scarring documented in Section 3 does not simply reflect the weight of large firms' yen losses.

Two qualifications apply. The split rests on 44 non-construction and 59 construction firms, so the magnitudes are imprecisely estimated. More important, while each group's coefficient separates a significant non-construction decline from an insignificant construction one, the log specification does not resolve the difference between them: the between-group test fails to reject equality for every outcome (p between 0.12 and 0.19), against $p \leq 0.02$ in levels. The proportional evidence for the sectoral divergence is thus weaker than the level evidence.

Figure 5 traces the dynamics behind this divergence. It reports event-study estimates of Equation (2) on the construction and non-construction subsamples of the balanced panel separately, with each group's coefficients scaled by that group's pooled 2010 mean of the outcome so that the two are comparable despite their different sizes. Construction firms show no persistent decline: their post-2011 coefficients remain near zero and turn positive in the later years. Non-construction firms diverge steadily, the treated–control

Figure 5: Event-study: Construction versus Non-Construction: *Balanced panel*

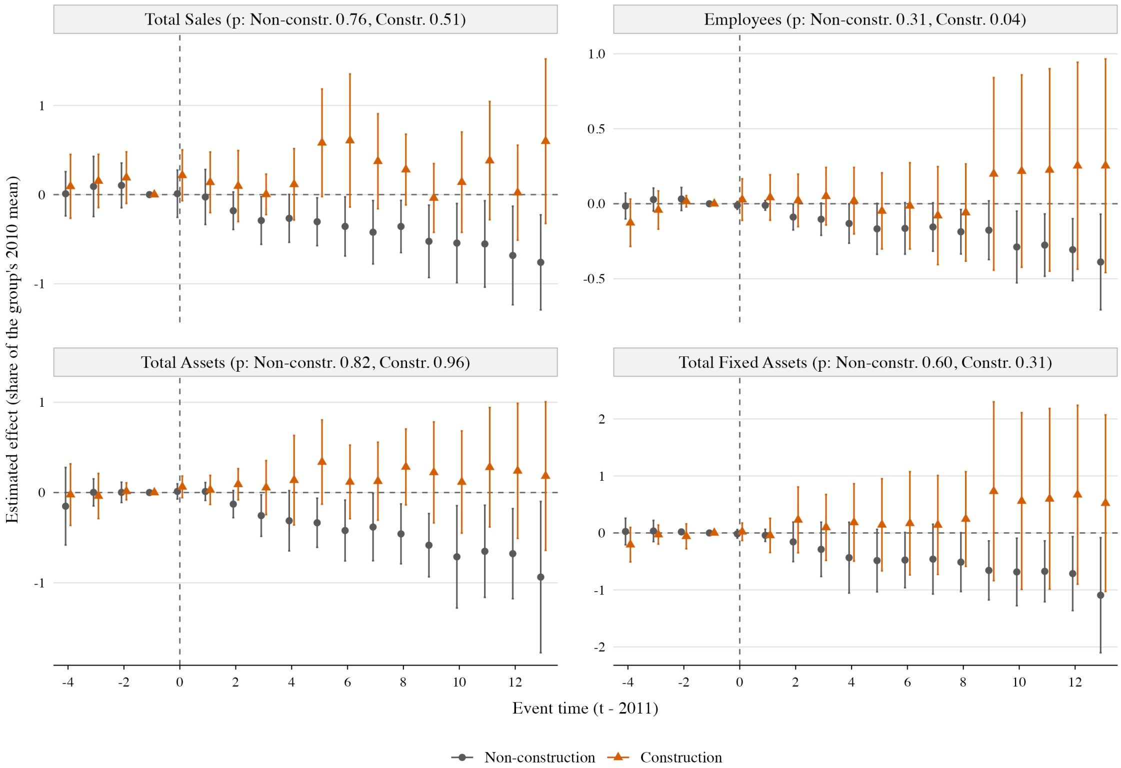

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the construction (59 firms) and non-construction (44 firms) subsamples of the balanced panel with the main-specification fixed effects. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

gap in sales widening to roughly half of the group's 2010 mean by the end of the horizon. The contrast is consistent with reconstruction demand sustaining construction firms while firms outside the sector absorbed the disruption.

3.3.2. Firm Size

We next examine whether the long-run effects of the GEJE differ by pre-disaster firm size. Firm size is defined using the firm's 2007–2010 mean of total sales: firms at or above the median across all respondent firms with pre-disaster accounting data are classified as large, and firms below the median as small. This heterogeneity is motivated by Figure 6, which compares reported damage among GEJE-treated firms by firm size (63 small and 56 large treated firms with pre-disaster data). The share reporting any damage is nearly identical across the two groups (about 70 versus 68 percent), but its composition differs: large treated firms more often report operational disruption, most visibly operation suspension (about 46 versus 30 percent), and tangible-asset damage (48 versus 40), while

Figure 6: Damage distribution among GEJE-treated firms by pre-disaster firm size

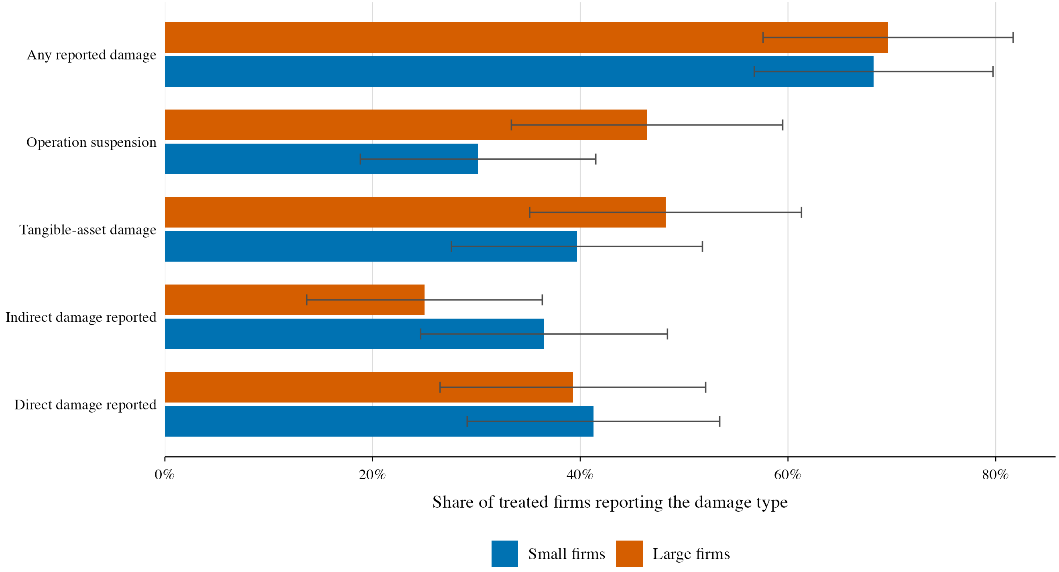

Notes: Share of GEJE-treated respondent firms reporting each damage type in the survey's disaster-experience module, separately for small and large firms (63 and 56 firms). Firm size is defined by the firm's 2007–2010 mean of total sales relative to the median across all respondent firms with pre-disaster accounting data. Tangible-asset damage and operation suspension are yes/no items; direct and indirect damage are flagged when the firm writes in a positive monetary amount (in thousand yen). "Any reported damage" is the union of the four categories. Error bars are approximate 95 percent confidence intervals based on the binomial standard error.

small treated firms write in monetary damage amounts at least as often as large ones (indirect damage 37 versus 25 percent).

This descriptive contrast suggests that larger firms may have been more exposed to operational disruption conditional on reporting GEJE exposure. We therefore estimate a pooled DD specification with separate treatment terms for small and large firms, in the form of the construction split of Equation (3):

$$Y_{it} = \alpha_i + \lambda_{c(i),t} + \beta_S\,(Treat_i \times Post_t \times Small_i) + \beta_L\,(Treat_i \times Post_t \times Large_i) + \gamma\,(Post_t \times Large_i) + \varepsilon_{it}, \quad (4)$$

where $Large_i$ is a dummy equal to one for firms whose 2007–2010 mean total sales lies at or above the median across all respondent firms with pre-disaster accounting data, and $Small_i = 1 - Large_i$. The fixed effects are those of the main specification. The $Post_t \times Large_i$ term, with coefficient γ reported in the table, controls for differential post-disaster trends of large firms; the $Treat_i \times Large_i$ interaction is constant within firm and absorbed by the firm fixed effects.

Table 4 reports the estimates. The point estimates are concentrated among larger treated firms: large treated firms lose 586.1 million yen in annual total sales (β_L, significant at the 10 percent level) and 335.2 million yen in total assets (10 percent), while every

Table 4. Heterogeneous DD Estimates: small vs. large firms (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t \times Small_i$	−59.410	1.004	9.213	−20.136
	(184.210)	(3.372)	(126.858)	(90.166)
$Treat_i \times Post_t \times Large_i$	−586.074*	−6.881	−335.225*	−159.453
	(323.101)	(4.700)	(190.527)	(109.194)
$Post_t \times Large_i$	323.585***	4.816***	397.350***	169.161***
	(95.479)	(1.477)	(74.596)	(46.159)
p-value: $\beta_L = \beta_S$	0.14	0.18	0.09	0.22
Observations	1854	1854	1854	1854
Firms (small)	26	26	26	26
Firms (large)	77	77	77	77
Pre-treatment mean (small)	191.46	11.32	146.66	63.44
Pre-treatment mean (large)	2,137.91	51.15	1,545.97	507.42
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. Firm and construction (2010) × year fixed effects are included; the $Post_t \times Large_i$ interaction controls for differential trends of large firms, and its coefficient is reported. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment means pool treated and control firm-years 2007–2010 within each group. The p-value row tests $\beta_L = \beta_S$ (Equation (4)), the equality of the large- and small-firm treatment coefficients, from the same specification re-parameterised with small firms as the baseline, so that the large-firm term is the between-group difference. The estimation sample is the balanced panel. Appendix Table B.3 reports the same specification on the balanced panel excluding 2012 and 2021.

small-firm coefficient β_S is close to zero and insignificant. The contrast is not statistically sharp, however. The null that $\beta_L = \beta_S$ is rejected only for total assets ($p = 0.09$). The estimates therefore establish that the losses are detectable among large treated firms and not among small ones, which is weaker than a difference between the two. The pattern is nonetheless consistent with the descriptive evidence in Figure 6.[13]

3.4. Robustness

Leave-One-Firm-Out Estimates The paper's central estimate, the non-construction coefficient β_N of Section 3.3.1, rests on 44 surviving firms, of which only 10 are treated, so a single anomalous trajectory could in principle account for it. To examine the sensitivity of this estimate to individual firms, we re-estimate the average DD specification of Equation (1) on this subsample, dropping one firm at a time.

Figure 7 illustrates the distributions of these iterative treatment coefficients across the four outcomes. The leave-one-out point estimates cluster tightly around the baseline reference line, and statistical significance is unaffected in every case: all 176 estimates — 44 firms × four outcomes — remain negative and significant at the 5 percent level, with a largest p-value of 0.048, and for total assets every estimate is significant at the 1 percent level. The non-construction scarring is therefore not the artifact of one firm's trajectory: it survives the removal of any one of the ten treated firms, and of any one of

[13] Appendix B.3 examines this size heterogeneity further.

Figure 7: Leave-one-firm-out robustness check: balanced panel excluding construction firms, level outcomes

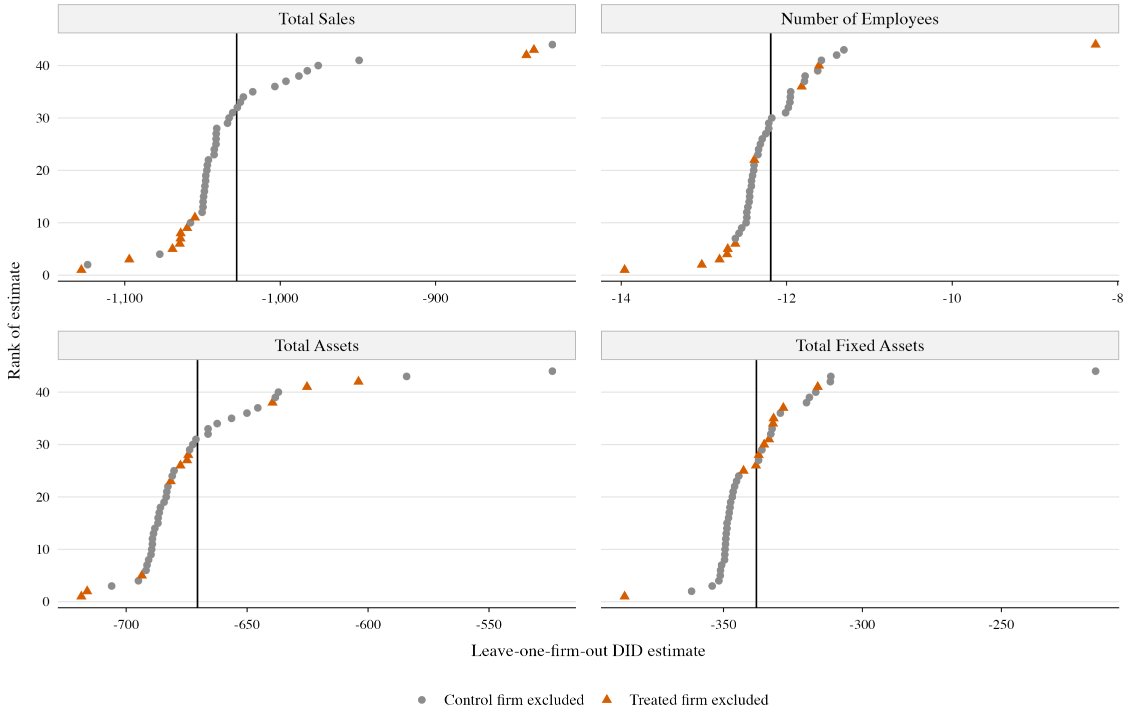

Notes: Leave-one-firm-out check for the average DD estimates on the balanced non-construction subsample (44 firms), level outcomes. For each outcome, each point is the treatment effect re-estimated after dropping one firm, with the main-specification fixed effects and firm-clustered standard errors; points are ordered vertically by estimate size, and their shape and color record whether the excluded firm belongs to the treatment (orange triangles) or the control group (grey circles). Every point is significant at the 5 percent level, so significance is not encoded in the figure. The vertical line marks the baseline estimate on the full non-construction subsample.

the thirty-four controls.

Inference with Few Treated Clusters The estimation panels contain many clusters — 103 firms in the balanced panel and 274 in the panel excluding 2012 and 2021 — but only 19 and 25 of them, respectively, are treated. With few *treated* clusters, cluster-robust standard errors can over-reject even when the total number of clusters is large (MacKinnon and Webb, 2018). We therefore re-evaluate the headline estimates with two procedures that do not rely on that asymptotic approximation: a restricted wild cluster bootstrap with firm-level Rademacher weights (Cameron et al., 2008; MacKinnon and Webb, 2018), and randomization inference against the sharp null of no effect for any firm, re-assigning treatment within 2010-construction strata (Young, 2019); both use 4,999 draws. The bootstrap p-values closely match the analytic ones, indicating that the cluster-asymptotic approximation is not the binding constraint at these cluster counts, while the randomization p-values are generally somewhat larger, as expected under the sharp null with few treated firms. The substantive conclusions are unaffected (Tables A.5–

A.7): the average sales and asset declines remain significant under both procedures, at the 10 percent level on the balanced panel and at or near the 5 percent level on the panel excluding 2012 and 2021; the concentration of the losses among non-construction firms is robust to all three procedures; and the joint pre-trend tests continue not to reject for any outcome.

Construction strata reproduce the treated firms' sectoral but not their size composition, so a small p-value could still reflect the differential trends of firms carrying the treated group's size mix. Table A.8 repeats the exercise with construction $\times$ pre-disaster median size strata (Section 3.3.2), matching each placebo group to the treated firms on both margins. The composition-matched p-values are modestly larger, but the pattern of results is unaffected: fewer than one placebo group in ten produces average sales or asset losses as large as those observed. The estimated scarring is therefore not an artifact of the treated firms' sectoral or size mix.

Manufacturing-by-Year Fixed Effects Treated firms are more often manufacturers than controls in both panels (37 versus 18 percent on the balanced panel and 44 versus 27 percent excluding 2012 and 2021, $p \approx 0.07$; Table 1), and manufacturing was differentially exposed to post-2011 aggregate shocks unrelated to a firm's own disaster experience: the nationwide supply-chain disruptions and the sharp yen appreciation of 2011–12, neither of which the construction-by-year fixed effects absorb. This is the one imbalance in the estimation panels that bears directly on the identifying assumption, since it pairs an imbalanced covariate with a large aggregate shock that moves with it.

We address it in two ways. Replacing the construction-by-year cells with manufacturing-by-year fixed effects, the manufacturing analogue of the main specification (Table A.9), leaves the estimates close to their baseline values with the same significance pattern, and the corresponding event studies (Figures A.12 and A.13) reproduce the flat pre-trends and the lagged, deepening losses of the headline figures. Estimating separate treatment effects for manufacturing and non-manufacturing firms (Table A.10) yields negative sales and asset point estimates in both groups, neither precisely estimated, and the test of their equality does not reject. Neither exercise indicates that a manufacturing-specific shock accounts for the estimated scarring.

Survival Conditioning and Firm Exit Section 2.2 noted that the panels condition on continuous observation and that the registry evidence on exit implies that survivorship bias attenuates the estimates rather than inflating them. This subsection reports the supporting evidence. The survey panels cannot settle the question, since exit is excluded from them by construction. We therefore turn to the TSR registry, following the roughly 1.5 million firms first observed by 2010 with a geocodable address and estimating a discrete-time exit hazard with pre-2011 industry-by-year and size-bin-by-year fixed effects; treatment is assigned from the physical exposure of Section 4.1, since the self-report exists only for survey respondents. Exposed firms were *less* likely to exit than observably similar unexposed firms (Table A.11), an effect largest immediately after the disaster, fading over time, and concentrated among the smallest firms. It is robust to the mechanical explanations, the most substantial of which is that TSR intensified its surveying of the affected region after 2011: restricting the cohort to firms already present in at least three pre-2011 vintages shrinks the estimate only modestly. Panel conditioning does not select the two groups differently either: once the 2012 and 2021 vintage discontinuities are removed, exposed and unexposed firms enter a registry-based balanced panel at essentially the same rate.

Head-Office Relocation The survey records establishments as of the survey date, so a firm that moved after 2011 reports the March 2011 event as experienced at its then-current location (Section 2.2). We measure relocation directly from the head-office municipality recorded in successive TSR registry vintages. Relocation is rare in both panels but more common among treated firms, and they are spread across the whole 2011–2024 window rather than concentrated immediately after the disaster — a pattern itself consistent with relocation as one margin of adjustment to the shock. Re-estimating the average DD without the movers (Table A.12) leaves the delayed and persistent profile unchanged (Figures A.17 and A.18), with coefficients somewhat below their baseline magnitudes and less precisely estimated once three of the treated firms are removed. Because the excluded firms plausibly relocated *because of* the disaster, dropping them removes some genuinely affected firms, so the modest attenuation is expected and does not indicate a weaker underlying effect.

4. Alternative Treatment Definitions

The treatment indicator of this paper is a self-report: a firm is treated because it reports, fourteen years later, having experienced the March 2011 earthquake or tsunami at either of its establishments. Section 2.2 discussed the recall properties of that indicator. This section considers what it measures. A self-report carries no severity threshold: a firm enters the treatment group if the disaster affected it in any way it recognized — physical destruction, operational disruption, or losses transmitted through suppliers and customers — and the module records the accompanying damage separately from the event itself. The indicator is therefore best read as broad exposure by the firm's own account.

The geographic evidence supports this reading. Among self-reporting firms, 82 percent experienced instrumental seismic intensity of at least 4.5, so the self-report does capture genuine shaking; but only 27 percent reach the severe 5.5 threshold that defines physical exposure below, and only about 21 percent of the survey-treated firms on the balanced

Table 5. Alternative treatment definitions: treated and total firm counts

Treatment definition	Treated	Firms
Baseline: self-reported exposure	19	103
Physical (GIS-based) exposure	9	103
Designated disaster area (41 municipalities)	12	103

Notes: Each row swaps the treatment indicator on the balanced panel, holding the specification fixed. Baseline: the firm reports experiencing the March 2011 earthquake or tsunami at either establishment. Physical exposure: the geocoded pre-disaster (2010-pinned) address lies within 500 m of the GSI tsunami-inundation footprint or in the zone of interpolated instrumental seismic intensity of at least 5.5; self-reported exposure plays no role, and firms without a geocodable address are dropped. Disaster area: the 2010-pinned address lies in one of the 41 officially designated disaster-area municipalities of Carvalho et al. (2021); firms without a geocodable address are dropped. Appendix Table B.7 reports the same counts on the balanced panel excluding 2012 and 2021.

panel are physically treated (Appendix Figure A.25). The baseline treatment group thus pools severely and mildly affected firms under one label, which attenuates the estimates if the mildly affected suffered little.

To assess how far our main estimates depend on the self-report, we re-estimate the design under two alternative treatment definitions that assign exposure solely from the firm's geocoded pre-disaster location: physical hazard (the tsunami footprint and severe-shaking zone), and the officially designated disaster area. Under both, self-reporting firms outside the exposed zone serve as controls, and firms without a geocodable address are dropped. Table 5 summarizes the three definitions and their treated counts on the balanced panel; Appendix Table B.7 reports the same counts on the panel excluding 2012 and 2021.

Because the definitions classify the same firms, their overlap is itself informative. Table 6 cross-tabulates the three treated groups on the balanced panel: a diagonal cell counts the firms treated under one definition, an off-diagonal cell the firms treated under both. The overlap is limited on the self-reported margin and nearly complete on the location-based one. Of the 19 self-reported treated firms on the balanced panel, only 4 are physically treated and 8 lie in the designated disaster area; of the 9 physically treated firms, by contrast, 8 lie in the disaster area. The location-based definitions thus largely nest one another, while the self-report captures a much wider set of firms than either. Appendix Table B.8 reports the same cross-tabulation on the panel excluding 2012 and 2021.

4.1. Physical Exposure

The first alternative measure sets the self-report aside and assigns exposure from location. A firm is *physically* treated if its pre-disaster (fiscal 2010) geocoded address lies in or within 500 meters of the official tsunami-inundation footprint mapped by the Geospatial Information Authority of Japan, or in the zone of strong ground shaking, an instrumental

Table 6. Alternative treatment definitions: cross-tabulation of the treated groups

	Self-reported	Physical	Disaster area
Self-reported exposure	19	4	8
Physical (GIS-based) exposure	4	9	8
Designated disaster area	8	8	12

Notes: The three treatment definitions of Table 5, cross-tabulated on the balanced panel (103 firms); every definition is defined for every firm in the panel. A diagonal cell reports the number of firms treated under the definition; an off-diagonal cell the number treated under both definitions, so the full two-by-two contingency table of any pair follows from these counts and the panel's firm count. Self-reported exposure is the baseline definition; physical exposure and the designated disaster area are the location-based definitions of Sections 4.1 and 4.2. Appendix Table B.8 reports the same cross-tabulation on the balanced panel excluding 2012 and 2021.

seismic intensity of at least 5.5 at the firm's address. We construct the intensity following Hanaoka et al. (2018), as an inverse-distance-weighted average of the intensities recorded for the mainshock at the three nearest Japan Meteorological Agency stations, measured at the firm's own geocoded address rather than at the city hall of its municipality. Appendix A.6 gives the construction in full.

The physical-exposure measure yields markedly different estimates compared to the self-reported measure. Swapping the self-reported treatment for the physical one on the same balanced panel, holding the sample and the specification fixed, turns the average level estimates from negative and significant to positive and imprecise (Panel A of Table 7), and the log total-assets estimate from near zero to significantly positive (0.38; Appendix Table A.13). Only nine firms in the 103-firm panel are physically treated, so these estimates pin down direction rather than magnitude. What the design finds therefore depends on which measure defines exposure, even when the sample, the period, and the estimating equation are held fixed.

4.2. The Designated Disaster Area

The second alternative definition replaces the physical-hazard footprint with an administrative one: a firm is treated if its 2010-pinned geocoded address lies in one of the 41 municipalities that constitute the officially designated disaster area, the definition Carvalho et al. (2021) use to identify disaster-hit firms, combining the MLIT decree of 28 April 2011 designating severely affected municipalities with the government evacuation decrees around the Fukushima Dai-ichi plant. It is the coarsest defensible location-based definition, and therefore a natural external benchmark: the designation treats whole municipalities, including inland cities far from the coast and all five wards of Sendai, so its treated group is differently composed and much less exposed on average than the physical-hazard group. Twelve firms on the balanced panel (13 on the panel excluding 2012 and 2021) fall in the disaster area.

Panel B of Table 7 reports the estimates. Every coefficient is small and statistically indistinguishable from zero: average total sales move by +66.7 million yen on the balanced

Table 7. Location-based treatment definitions: average DD estimates (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Physical (GIS-based) exposure				
$Treat_i \times Post_t$	495.134	6.063	366.227	166.171
	(351.583)	(4.359)	(223.922)	(132.927)
Treated firms	9	9	9	9
Panel B: Designated disaster area				
$Treat_i \times Post_t$	66.713	−2.950	73.065	27.315
	(367.407)	(5.087)	(217.791)	(114.217)
Treated firms	12	12	12	12
Observations	1854	1854	1854	1854
Firms	103	103	103	103
Pre-treatment mean	1,646.58	41.09	1,192.75	395.35
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Both panels re-estimate Equation (1) on the balanced panel, replacing the self-reported treatment indicator with a location-based one; the sample, specification, fixed effects, and pre-treatment means are common to the two panels and are reported once at the foot of the table. *Panel A*: a firm is treated if its geocoded pre-disaster address (registry address pinned to 2010, with the earliest available registry address as a fallback) lies within 500 m of the tsunami-inundation footprint mapped by the Geospatial Information Authority of Japan, or in the zone of instrumental seismic intensity of at least 5.5 interpolated from Japan Meteorological Agency station records. *Panel B*: a firm is treated if that address lies in one of the 41 officially designated disaster-area municipalities of Carvalho et al. (2021). Under both definitions self-reported exposure plays no role: survey-treated firms outside the exposed zone serve as controls, and firms without a geocodable address are excluded. Firm and construction (industry classification as of 2010) × year fixed effects are included throughout. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010. Appendix Tables B.4 and B.6 report the same specifications on the balanced panel excluding 2012 and 2021, and Appendix Table A.13 the log estimates under physical exposure.

panel and by −20.5 million on the panel excluding 2012 and 2021 (Appendix Table B.6). Neither location-based measure therefore yields evidence of scarring among the panel firms, while the self-report yields losses that are significant in levels.

4.3. Discussion

These comparisons characterize what each measure captures. The self-reported and physically treated groups overlap only partially and neither nests the other (Table 6), so the two definitions estimate the effects of different treatments on the same firms. Nearly all physically treated firms lie in the designated disaster area, where the reconstruction program of Appendix A.4 was concentrated and where exposed firms were less likely to exit after 2011; location-based exposure therefore proxies for presence in the reconstruction zone as much as for damage. The self-report instead identifies firms whose operations were disrupted, including at establishments and through business relationships that a geocoded headquarters address cannot observe. We view that the two types of measures complement each other: reconstruction inflows may have outweighed physical damage at the typical small firm's pre-disaster address, while the exposure firms report themselves corresponds to the shortfall the headline estimates quantify.

5. Concluding Remarks

This paper evaluates the long-term economic impacts of the 2011 Great East Japan Earthquake on small and medium-sized enterprises. Combining a proprietary survey, which measures disaster exposure, with an 18-year panel of financial statements (fiscal 2007–2024), which measures outcomes, we document persistent economic scarring: firms that report having experienced the disaster fall durably behind comparable respondents in sales, assets, and fixed assets, the losses emerge with a lag, deepen over time, and persist through 2024 with no sign of reversal. The scarring is concentrated among firms outside the construction sector, where reconstruction demand did not reach, the decline is proportional as well as absolute, and among larger treated firms. Comparing self-reported with physical exposure on the same firms, we show that the estimated effect changes sign with the choice of exposure measure: the two identify overlapping but distinct groups of firms, and location-based exposure captures presence in the reconstruction zone as much as damage.

Two policy implications follow. First, on the time dimension: post-disaster interventions are typically framed over three-to-five-year horizons, yet the losses we document were still deepening a decade after the event. Recovery policy, and the monitoring that informs it, needs multi-decadal framing. Second, on targeting: reconstruction demand shielded the construction sector, while the firms it did not reach bore losses that widened for over a decade. Support that flows through reconstruction spending alone therefore misses the firms with the most persistent scarring, and instruments that reach firms outside the construction supply chain, credit lines, demand-side support, adaptation subsidies, deserve a larger role. Third, exposure assigned from administrative boundaries conflates damage with the reconstruction inflows that followed, and registers no lasting loss whereas firms' own reports register a widening shortfall. Monitoring by location alone will therefore likely understate the persistence of the shock, and firm-reported exposure is an important complement.

References

Cabinet Office. 2011. "Estimate of the Damage Caused by the Great East Japan Earthquake." press release, Cabinet Office, Government of Japan, Director-General for Disaster Management, In Japanese. Available at https://www.bousai.go.jp/2011daishinsai/pdf/110624-1kisya.pdf.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics* 225 (2): 200–230. 10.1016/j.jeconom.2020.12.001.

Camerer, Colin F, and Howard Kunreuther. 1989. "Decision Processes for Low Probability Events: Policy Implications." *Journal of Policy Analysis and Management* 8 (4): 565–592.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics* 90 (3): 414–427. 10.1162/rest.90.3.414.

Carleton, Tamma, Esther Duflo, B. Kelsey Jack, and Guglielmo Zappalà. 2024. "Adaptation to Climate Change." In *Handbook of the Economics of Climate Change*, edited by Barrage, Lint, and Solomon Hsiang Volume 1. 143–248, Amsterdam: North-Holland.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. 2021. "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake." *Quarterly Journal of Economics* 136 (2): 1255–1321.

Cavallo, Eduardo, Sebastian Galiani, Ilan Noy, and Juan Pantano. 2013. "Catastrophic Natural Disasters and Economic Growth." *Review of Economics and Statistics* 95 (5): 1549–1561.

Chuang, Yanyan, and Laura Schechter. 2015. "Stability of Experimental and Survey Measures of Risk, Time, and Social Preferences: A Review and Some New Results." *Journal of Development Economics* 117 151–170.

Glaeser, Edward L. 2022. "Urban Resilience." *Urban Studies* 59 (1): 3–35. 10.1177/00420980211052230.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–277. 10.1016/j.jeconom.2021.03.014.

Hanaoka, Chie, Hitoshi Shigeoka, and Yasutora Watanabe. 2018. "Do Risk Preferences Change? Evidence from the Great East Japan Earthquake." *American Economic Journal: Applied Economics* 10 (2): 298–330. 10.1257/app.20170048.

Haushofer, Johannes, and Daniel Salicath. 2023. "The Psychology of Poverty: Where Do We Stand?" *Social Philosophy and Policy* 40 (1): 150–184. 10.1017/S0265052523000419.

Inoue, Hiroyasu, and Yasuyuki Todo. 2019. "Firm-Level Propagation of Shocks through Supply-Chain Networks." *Nature Sustainability* 2 (9): 841–847. 10.1038/s41893-019-0351-x.

Kameda, Keigo, Hiroyoshi Miwa, and Jun-ichi Nagamine. 2021. "Effects of Reconstruction Works on Private Employment after a Natural Disaster." *International Journal of Disaster Risk Reduction* 52 101968. 10.1016/j.ijdrr.2020.101968.

Kellenberg, Derek, and A Mushfiq Mobarak. 2011. "The Economics of Natural Disasters." *Annual Review of Resource Economics* 3 (1): 297–312.

Kunreuther, Howard. 1978. *Disaster Insurance Protection: Public Policy Lessons.* New York: Wiley.

Kuroishi, Yusuke, and Yasuyuki Sawada. 2024. "Natural Disasters and Preference Instability: An Evaluation of the Stability of Time Preferences." *European Economic Review* 161 104618.

Leiter, Andrea M, Harald Oberhofer, and Paul A Raschky. 2009. "Creative Disasters? Flooding Effects on Capital, Labour and Productivity within European Firms." *Environmental and Resource Economics* 43 (3): 333–350.

MacKinnon, James G., and Matthew D. Webb. 2018. "The Wild Bootstrap for Few (Treated) Clusters." *Econometrics Journal* 21 (2): 114–135. 10.1111/ectj.12107.

Miyazaki, Tomomi. 2025. "Disaster Relief and Regional Employment: The Case of the Great East Japan Earthquake." *Applied Economics Letters* 32 (1): 78–86. 10.1080/13504851.2023.2257025.

National Police Agency. 2026. "Police Countermeasures and Damage Situation Associated with the 2011 off the Pacific Coast of Tōhoku Earthquake." emergency disaster countermeasures headquarters report, National Police Agency of Japan, In Japanese. Cumulative tabulation as of February 28, 2026. Available at `https://www.npa.go.jp/news/other/earthquake2011/pdf/2026higaijyoukyou.pdf`.

Noy, Ilan, and William duPont IV. 2018. "The Long-Term Consequences of Disasters: What Do We Know, and What We Still Don't." *International Review of Environmental and Resource Economics* 12 (4): 325–354.

Rehdanz, Katrin, Heinz Welsch, Daiju Narita, and Toshihiro Okubo. 2015. "Well-Being Effects of a Major Natural Disaster: The Case of Fukushima." *Journal of Economic Behavior & Organization* 116 500–517. 10.1016/j.jebo.2015.05.014.

Sawada, Yasuyuki. 2007. "The Impact of Natural and Manmade Disasters on Household Welfare." *Agricultural Economics* 37 (s1): 59–73.

Sawada, Yasuyuki. 2022. "Preferences, Behavior, and Welfare Outcomes against Disasters: A Review." In *Handbook on the Economics of Disasters*, edited by Skidmore, M. Edward Elgar Publishing.

Sawada, Yasuyuki, Yusuke Kuroishi, Toyo Ashida et al. 2025. "Unhealthy Behaviours in Disaster Survivors Are Associated with Scarcity and Present Bias." *Communications Psychology* 3 168. 10.1038/s44271-025-00344-0.

Schildberg-Hörisch, Hannah. 2018. "Are Risk Preferences Stable?" *Journal of Economic Perspectives* 32 (2): 135–154.

Skidmore, Mark, and Hideki Toya. 2002. "Do Natural Disasters Promote Long-Run Growth?" *Economic Inquiry* 40 (4): 664–687.

Sun, Liyang, and Sarah Abraham. 2021. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics* 225 (2): 175–199. 10.1016/j.jeconom.2020.09.006.

Takeda, Kohei, and Atsushi Yamagishi. 2026. "The Economic Dynamics of City Structure: Evidence from Hiroshima's Recovery." *Journal of Political Economy*. 10.1086/743334, Forthcoming.

Tokui, Joji, Kazuyasu Kawasaki, and Tsutomu Miyagawa. 2017. "The Economic Impact of Supply Chain Disruptions from the Great East-Japan Earthquake." *Japan and the World Economy* 41 59–70. 10.1016/j.japwor.2016.12.005.

Young, Alwyn. 2019. "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results." *Quarterly Journal of Economics* 134 (2): 557–598. 10.1093/qje/qjy029.

A. Appendix

Contents of the Appendix

A.1. The RIETI Survey

A.1.1. Design and Fielding

The roster of 15,000 firms is divided evenly between two arms of 7,500 firms each: one drawn from firms certified under the Small and Medium Enterprise Agency's Business Continuity Enhancement Plan certification scheme (single-firm type), and one drawn from the TSR register at large. Within each arm the draw was allocated across seventeen industry divisions in numbers that are identical between the two arms to within a single firm, so that the certified and non-certified halves of the roster are matched on industry. The allocation embodies several further criteria: (i) restricting the target to the thirty-two prefectures containing municipalities affected by disasters designated as severe disasters, including localized designations;[14] (ii) accounting for the distribution of firms at the municipal level; (iii) concentrating the industry quotas where certified firms are numerous; and (iv) deliberately focusing on firms operating near the thresholds of the industry-specific SME definitions, particularly around the upper and lower limits of medium-sized enterprises, through allocation across cells defined by capital and employment size.

The questionnaires were distributed by mail, and responses were collected via mail, a dedicated website, and email, with a response deadline of February 21, 2025. To enhance the response rate, reminders were issued by mail and telephone three weeks after the initial distribution.

[14] In Japan, a "severe disaster" is an administrative classification of a disaster, determined by the Cabinet under the Act on Special Financial Support to Deal with the Designated Disaster of Extreme Severity. The designation attaches to the disaster, either nationwide or as a localized designation naming the municipalities it covers, and triggers special financial measures including enhanced subsidy rates and special disaster-recovery credit guarantees and loans targeting affected SMEs.

A.1.2. Survey Questionnaire: Disaster-Experience Module

The treatment indicator of Section 2.2 is constructed from Q1 of the questionnaire, reproduced below in the authors' English translation of the original Japanese instrument. The translation preserves the structure of the original answer grid; blank rules (____) mark write-in fields. The grid was printed twice, once per establishment, and is shown once. The remaining questions — on financial coping, changes in business strategy, mitigation practices, and business continuity planning, are analyzed in a companion paper and are omitted.

Q1. Please report your experience of disaster damage over the past 15 years, separately for each establishment. For the type of disaster, choose the applicable number from the choice list below and enter it in the answer column.

Type of disaster (choice list)

1. Earthquake	2. Wind damage (typhoon, storm, tornado, etc.)	3. Flood (heavy rain, river flooding, etc.)
4. Lightning	5. Heavy snow	6. Storm surge and other abnormal tides, wave damage
7. Tsunami	8. Volcanic eruption	9. Landslide
10. Drought / water shortage		11. Other (specify: ________________)

Episode	Year and month	Disaster type (enter no.)	Direct damage*1 (thousand yen)	Indirect damage*2 (thousand yen)	Tangible assets damaged	Operations suspended
1st	____/___	_____	_______	_______	1. Yes 2. No	1. Yes 2. No
2nd	____/___	_____	_______	_______	1. Yes 2. No	1. Yes 2. No
3rd	____/___	_____	_______	_______	1. Yes 2. No	1. Yes 2. No
4th	____/___	_____	_______	_______	1. Yes 2. No	1. Yes 2. No
5th	____/___	_____	_______	_______	1. Yes 2. No	1. Yes 2. No

*1: For the direct damage amount, report the damage to stocks, such as buildings and equipment.

*2: For the indirect damage amount, report the damage to flows caused by the suspension of business.

[An identical grid was provided for a second establishment; respondents with no second establishment were instructed to leave it blank. "Operations suspended" options read "1. Suspended; 2. Not suspended" in the original.]

A.2. Variable Definitions and Additional Descriptions

Table A.1. Variable Definitions

Variable	**Definition**	**Unit**
Panel A: Main outcomes		
Total sales	Sum of sales, completed construction revenue, and operating revenue, minus sales discounts and returns.	Million JPY
Employees	Number of regular full-time employees, excluding part-time and temporary workers.	Persons
Assets	Total assets reported on the balance sheet.	Million JPY
Fixed assets	Sum of total tangible fixed assets, total intangible fixed assets, and total investments and other assets.	Million JPY
Panel B: Additional outcomes		
Operating profit	Operating profit from the income statement.	Million JPY
Gross profit	Gross profit (total sales minus cost of sales).	Million JPY
Total liabilities	Total liabilities reported on the balance sheet.	Million JPY
Net worth	Total net assets reported on the balance sheet; can be negative.	Million JPY
Long-term debt	Long-term borrowings plus corporate bonds.	Million JPY

Notes: All variables are constructed from the TSR registry financial statements (fiscal years 2007–2024); monetary amounts are converted from thousand to million yen. Where a table or figure header indicates a log transformation ($\log(1+Y)$), it is applied to the value defined here.

Figure A.1: Treated − control gap in mean outcomes, relative to 2010: *Balanced panel*

Notes: For each outcome, the difference in yearly means between treated and control firms on the balanced panel, minus the same difference in 2010, so that the gap is zero in 2010 by construction. The dashed vertical line marks 2011, the year of the GEJE. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure A.2: Location of the estimation samples

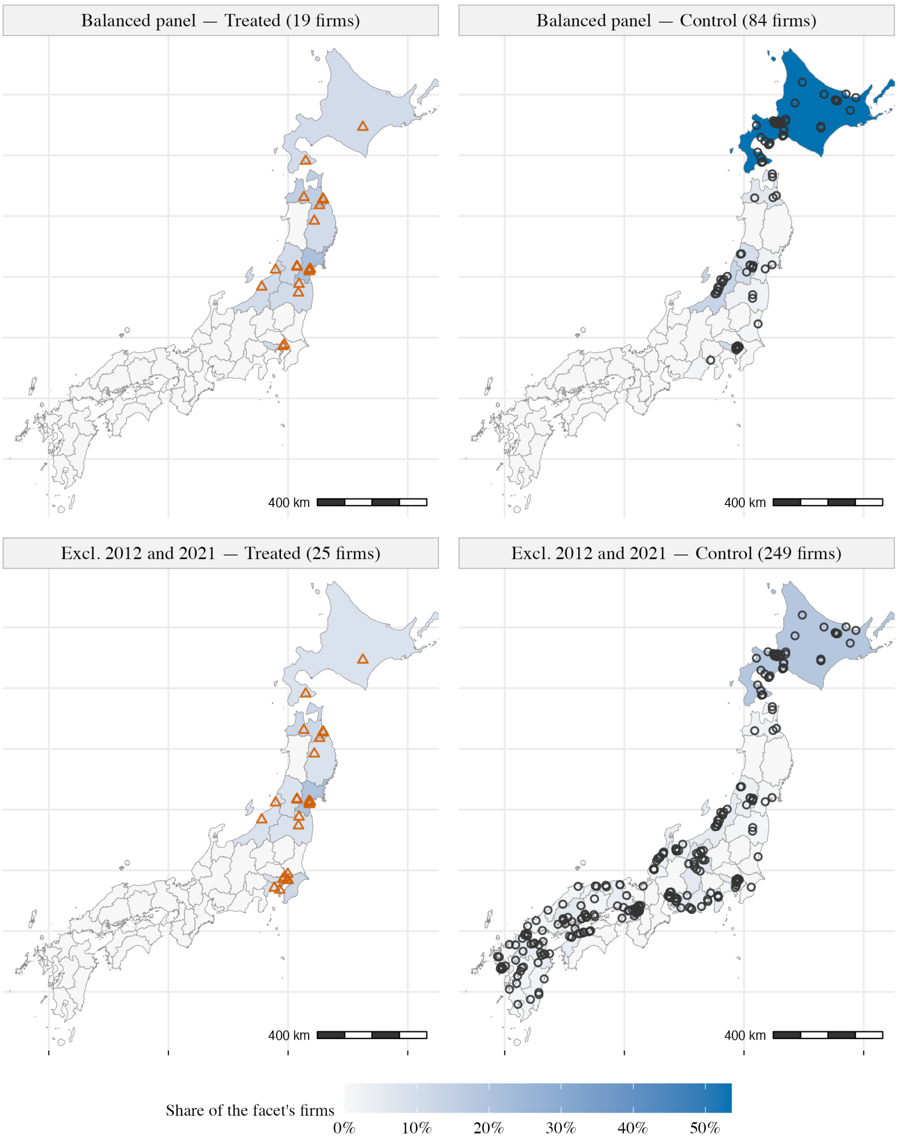

Notes: Each point is a firm at its geocoded pre-disaster address (registry address pinned to 2010, with the earliest available vintage as a fallback), drawn as an open triangle (treated) or open circle (control) so that overlapping firms remain distinguishable; the scale bar gives distance in kilometers. Prefecture shading gives the share of the facet's firms located in that prefecture, so each facet sums to 100 percent. Top row: the balanced panel (19 treated, 84 control firms); bottom row: the balanced panel reconstructed excluding the 2012 and 2021 fiscal years (25 treated, 249 control firms). Treated firms cluster in the Tohoku region and the Tokyo area in both samples. The balanced panel's control group is heavily concentrated in Hokkaido, whereas the ex-2012/2021 control group is spread across western Japan as well — the 2012 and 2021 TSR vintage gaps are concentrated in western Japan, so requiring balance across all eighteen years disproportionately removes western firms.

Figure A.3: Treated and control shares by sample

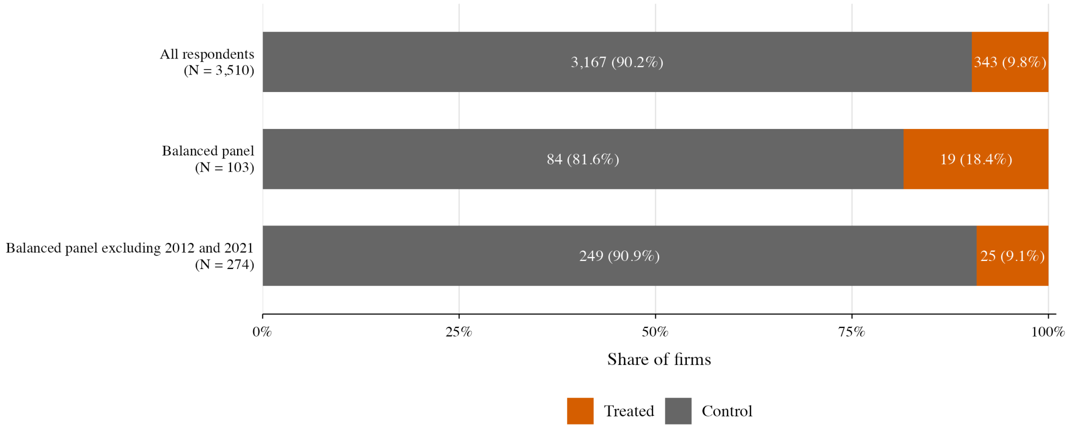

Notes: Share of firms reporting GEJE exposure (treated) among all 3,510 usable respondents and within each estimation sample. Bar segments are labeled with firm counts and shares. The treated share is markedly higher in the 103-firm balanced panel than among all respondents, while the panel excluding 2012 and 2021 mirrors the respondent pool.

Figure A.4: Reported damage among the firms reporting the March 2011 event

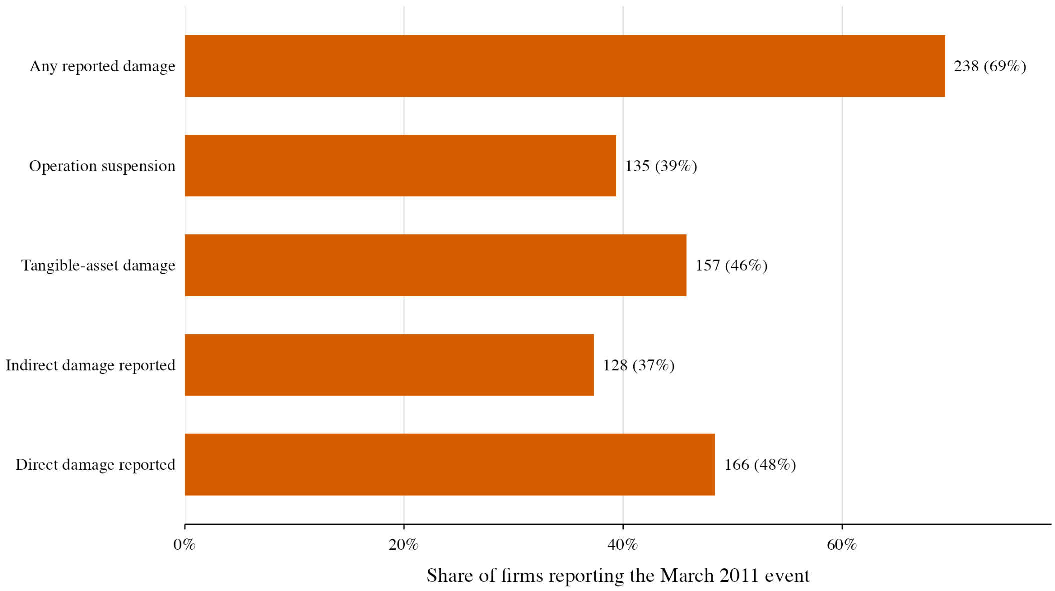

Notes: Share of the 343 respondent firms reporting the March 2011 earthquake or tsunami in the survey's disaster-experience module that report each damage type; bars are labeled with the firm count and share. Tangible-asset damage and operation suspension are yes/no items; direct and indirect damage are flagged when the firm writes in a positive monetary amount. "Any reported damage" is the union of the four categories. The damage items are recorded separately from the episode itself, so 105 of the 343 firms (31 percent) report the event without marking any damage category. Figure 6 plots the same margins for the treated firms with pre-disaster accounting data, split by firm size.

A.3. Sample Selection into the Estimation Panels

The estimation panels are not a random draw from the survey respondents, and this appendix characterizes the selection into them. Selection operates at three margins: response to the survey, linkage to at least one year of TSR financial statements, and panel completeness. At the response margin, respondent firms have larger initial capital (by about 6 log points) but roughly four fewer employees than non-respondents on the 15,000-firm roster (Table A.2), with a broadly similar industry mix. The two subsequent margins tilt unambiguously toward larger firms.

Figure A.5 compares respondent firms included in each sample with respondent firms excluded from it: included firms are larger in both roster employment and initial capital at every margin, and the gap widens sharply from the accounting-linked sample (about half a log point of capital and ten employees) to the two balanced panels (over a full log point of capital and roughly twenty-five to forty employees).

Figure A.6 shows the industry counterpart. The most pronounced difference is the rising representation of construction: construction firms make up about 19 percent of all respondents, 30 percent of the accounting-linked sample, 45 percent of the panel excluding 2012 and 2021, and 57 percent of the balanced panel, with manufacturing and the residual group falling correspondingly. These patterns carry two implications drawn on in the main text. The overrepresentation of construction is a further reason the main specification interacts construction status with year (Section 2.4), and the selection toward larger firms matters for interpreting the estimates, because proportional treatment effects decline steeply with firm size (Section 4): the estimates apply to firms that remain in the linked accounting and panel samples, which are larger and more concentrated in construction than the respondents excluded from them.

Table A.2. Response selection

	(1)	(2)
Outcome:	log(Initial Capital)	Number of Employees
Responded	0.058***	−4.361***
	(0.022)	(1.090)
Observations	14 992	14 836

Notes: Cross-section of the 15,000-firm survey roster. Each column regresses the roster characteristic on an indicator for responding to the survey; heteroskedasticity-robust standard errors in parentheses. Initial capital is in thousand yen (log); employees are in persons.

Figure A.5: Sample selection by firm size

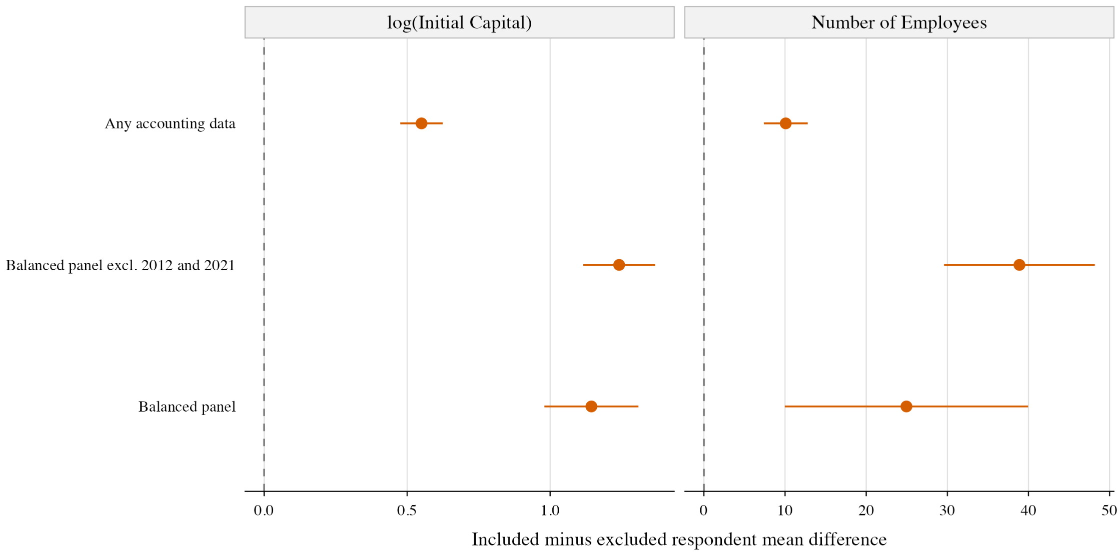

Notes: Each point is the mean difference in the roster size measure between respondent firms included in the indicated sample and respondent firms excluded from it, with 95 percent confidence intervals from heteroskedasticity-robust standard errors. Initial capital and employment are the roster (pre-survey) measures. See Figure 1 for the flow of the sample selection.

Figure A.6: Sample selection by industry

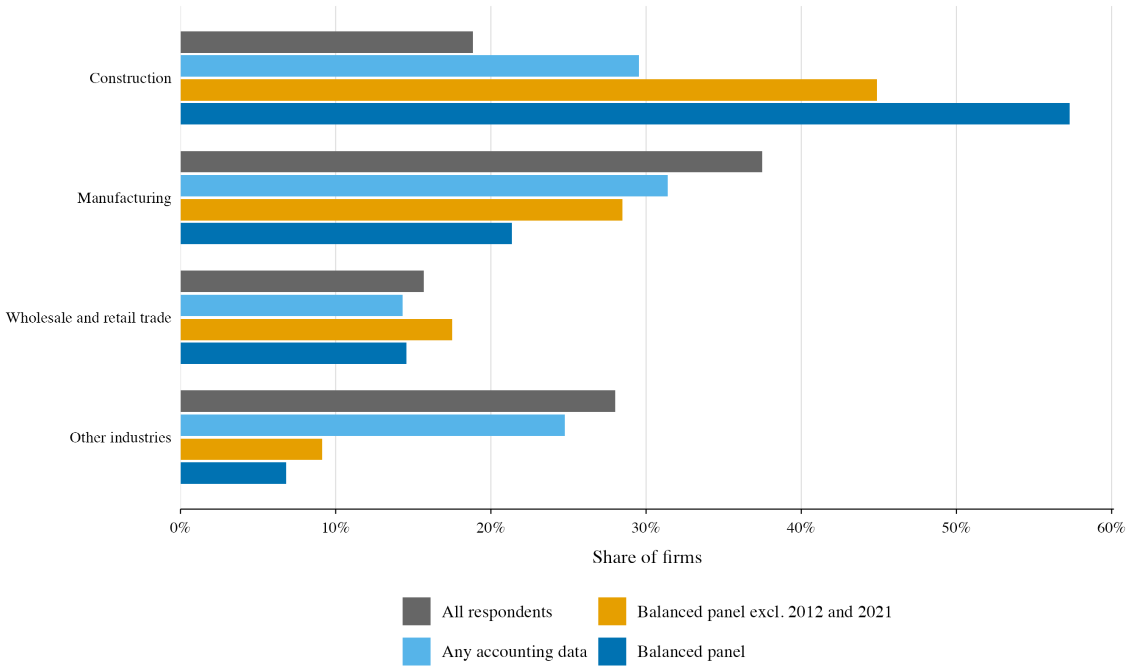

Notes: Industry composition, in four coarse groups based on the roster's primary industry code, of all respondents and of each estimation sample. See Figure 1 for the flow of the sample selection.

A.4. Reconstruction Demand after the GEJE

The main specification's construction-by-year fixed effects, and the construction split of Section 3.3, are motivated by the publicly financed reconstruction boom that followed the GEJE. This appendix documents that boom's size and duration directly. Figure A.7 plots earthquake-reconstruction-related construction orders (*shinsai fukkyū kankei juchūdaka*) from the national construction-orders statistics: the 12-month rolling sum of reconstruction-related orders (bars) and its share of all construction orders received (line). Reconstruction work amounted to roughly one-fifth of all construction orders at its 2012–2013 peak; the volume of orders peaked near 560 billion yen on a 12-month basis around 2014–2015, remained elevated through 2018–2019, and faded toward zero only in the early 2020s. Reconstruction demand was thus a large and persistent sectoral shock, sustained for close to a decade, and any comparison of treated and control firms that pools construction with other industries would confound a firm's own disaster experience with this sector-wide demand inflow, which is why every specification in the paper allows construction firms their own year effects.

Figure A.7: Reconstruction-related construction orders: size and duration

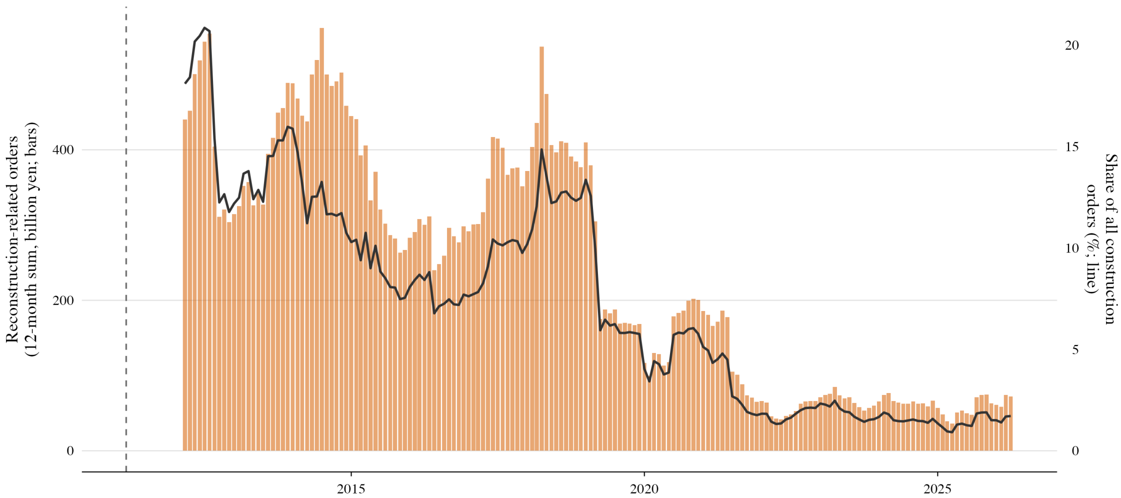

Notes: Earthquake-reconstruction-related construction orders reported in the Current Survey on Orders Received for Construction (national aggregates; Ministry of Land, Infrastructure, Transport and Tourism, via e-Stat). Bars show the 12-month rolling sum in billion yen (left axis); the line shows its share of all construction orders received (right axis, percent). The dashed vertical line marks March 2011; the underlying series begins in April 2011, so the first 12-month sum is available in March 2012.

A.5. Additional Figures and Tables

Table A.3. Average DD Estimates (log)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t$	-0.098	0.025	0.010	-0.047
	(0.099)	(0.122)	(0.122)	(0.139)
Observations	1854	1854	1854	1854
Firms	103	103	103	103
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t$	-0.051	0.040	0.001	-0.024
	(0.072)	(0.089)	(0.091)	(0.104)
Observations	4384	4384	4384	4384
Firms	274	274	274	274
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Outcomes are $\log(1+Y)$, so coefficients can be approximately interpreted as proportional changes. In Panel B, the 2012 and 2021 fiscal years are excluded, so the post period covers twelve rather than fourteen years.

Figure A.8: Dynamic Effects on Log Outcomes: *Balanced panel*

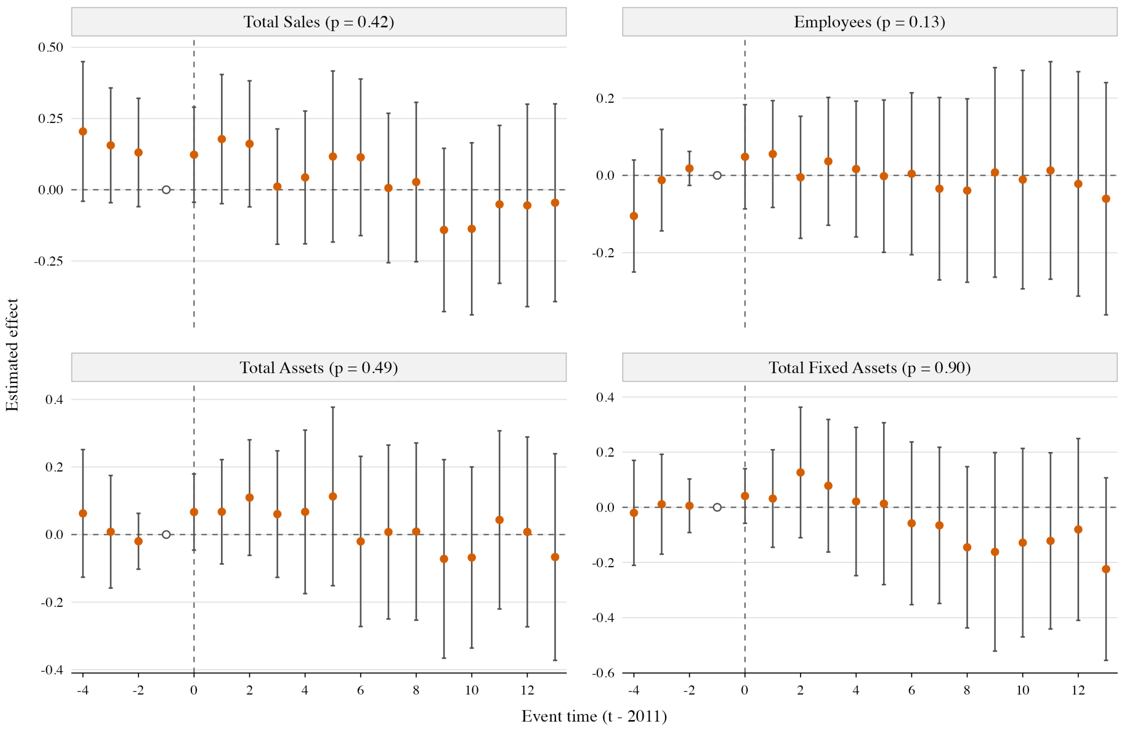

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, balanced panel (103 firms, of which 19 treated). The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) × year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Coefficients can be approximately interpreted as proportional changes.

Figure A.9: Dynamic Effects on Log Outcomes: *Balanced panel excluding 2012 and 2021*

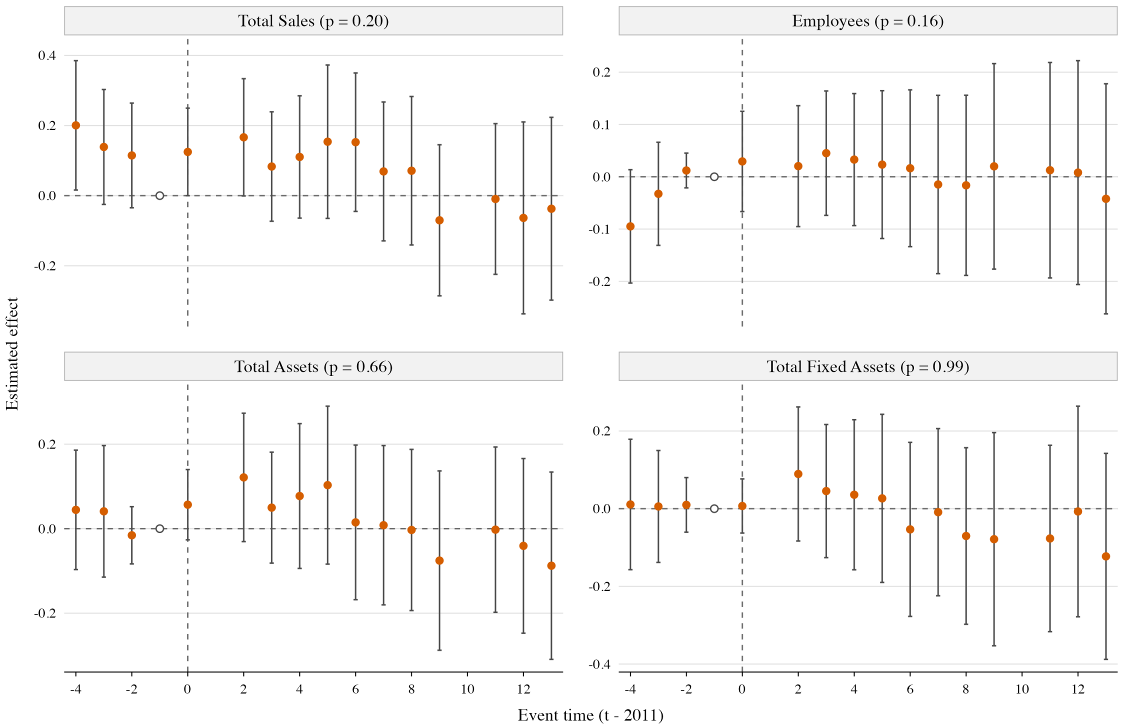

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, balanced panel reconstructed excluding the 2012 and 2021 fiscal years (274 firms, of which 25 treated); event years $k=1$ and $k=10$ are absent. The reference year is 2010 ($k=-1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Coefficients can be approximately interpreted as proportional changes.

Figure A.10: Event-study: Construction versus Non-Construction: *Balanced panel excluding 2012 and 2021*

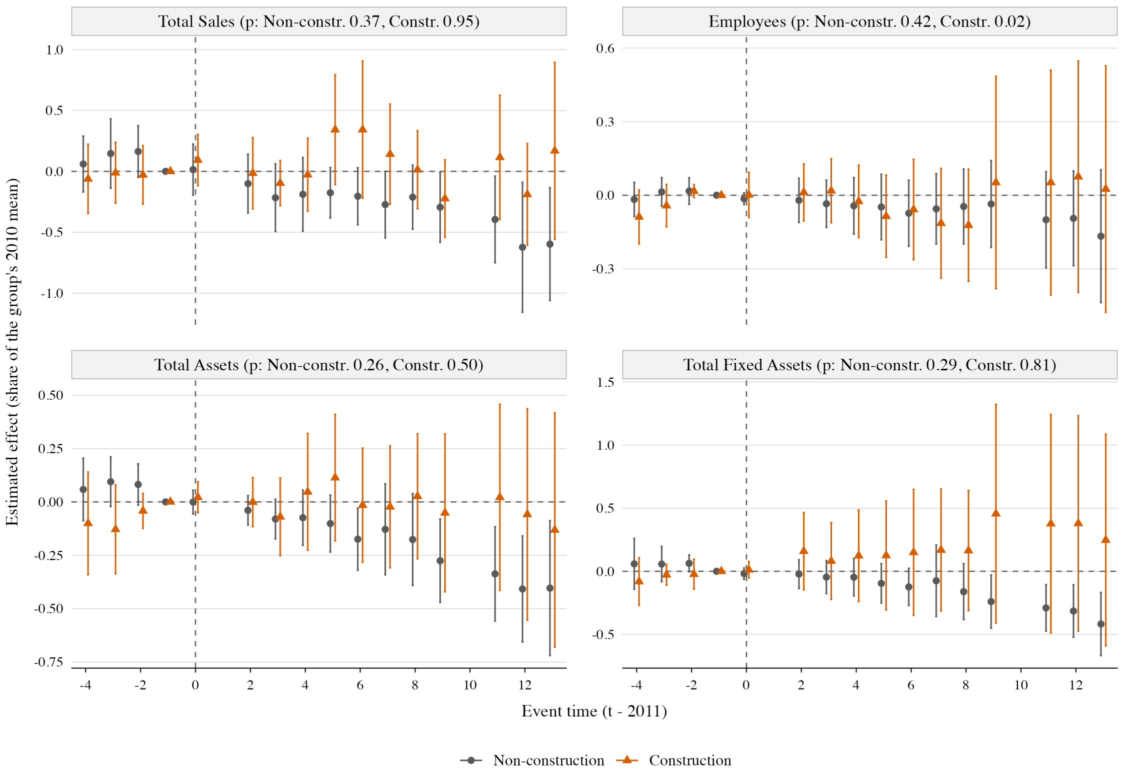

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the construction (123 firms) and non-construction (151 firms) subsamples of the balanced panel reconstructed excluding the 2012 and 2021 fiscal years; event years $k = 1$ and $k = 10$ are therefore absent. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Table A.4. Heterogeneous DD Estimates: construction vs. non-construction (log)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t \times Construction_i$	0.053	0.190	0.194	0.133
	(0.158)	(0.228)	(0.197)	(0.214)
$Treat_i \times Post_t \times NonConstruction_i$	−0.246**	−0.138**	−0.171	−0.225
	(0.105)	(0.066)	(0.126)	(0.165)
p-value: $\beta_C = \beta_N$	0.12	0.17	0.12	0.19
Observations	1854	1854	1854	1854
Firms (construction)	59	59	59	59
Firms (non-construction)	44	44	44	44
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t \times Construction_i$	0.054	0.120	0.178	0.178
	(0.129)	(0.188)	(0.160)	(0.174)
$Treat_i \times Post_t \times NonConstruction_i$	−0.133*	−0.023	−0.138	−0.184
	(0.074)	(0.054)	(0.087)	(0.112)
p-value: $\beta_C = \beta_N$	0.21	0.47	0.08	0.08
Observations	4384	4384	4384	4384
Firms (construction)	123	123	123	123
Firms (non-construction)	151	151	151	151
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Outcomes are $\log(1+Y)$, so coefficients can be approximately interpreted as proportional changes. Firm and construction (2010) × year fixed effects are included; the construction × year cells absorb the $Post_t \times Construction_i$ main effect. Standard errors clustered at the firm level are reported in parentheses. The p-value row tests $\beta_C = \beta_N$, the equality of the two group coefficients, from the same specification re-parameterised with construction firms as the baseline, so that the non-construction term is the between-group difference. In Panel B, the 2012 and 2021 fiscal years are excluded.

Table A.5. Inference with Few Treated Clusters: average DD (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t$	−447.338*	−4.799	−275.766*	−136.825
	(236.272)	(3.314)	(151.866)	(93.276)
p-values	0.061 [0.066] (0.053)	0.151 [0.168] (0.165)	0.072 [0.070] (0.093)	0.145 [0.148] (0.159)
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t$	−476.833**	−2.260	−275.062**	−94.189
	(213.075)	(3.192)	(120.334)	(61.219)
p-values	0.026 [0.019] (0.049)	0.479 [0.508] (0.563)	0.023 [0.028] (0.061)	0.125 [0.123] (0.164)

Notes: Each cell reports three p-values for the same null hypothesis: the analytic p-value based on standard errors clustered at the firm level, in square brackets the restricted (null-imposed) wild cluster bootstrap p-value (Rademacher weights drawn at the firm level, bootstrap-t, 4,999 draws), and in parentheses the randomization-inference p-value from 4,999 permutations of the firm-level treatment indicator within 2010-construction strata, holding the number of treated firms in each stratum fixed (studentized statistic; the randomization p-value tests the sharp null of no effect for any firm). Specification, samples, and estimates are those of the average DD table (levels): firm and construction (2010) × year fixed effects; 103 firms (19 treated) in Panel A and 274 firms (25 treated) in Panel B. Monetary variables are measured in 1 million yen.

Figure A.11: Event-study estimates by construction status: balanced panel, log outcomes

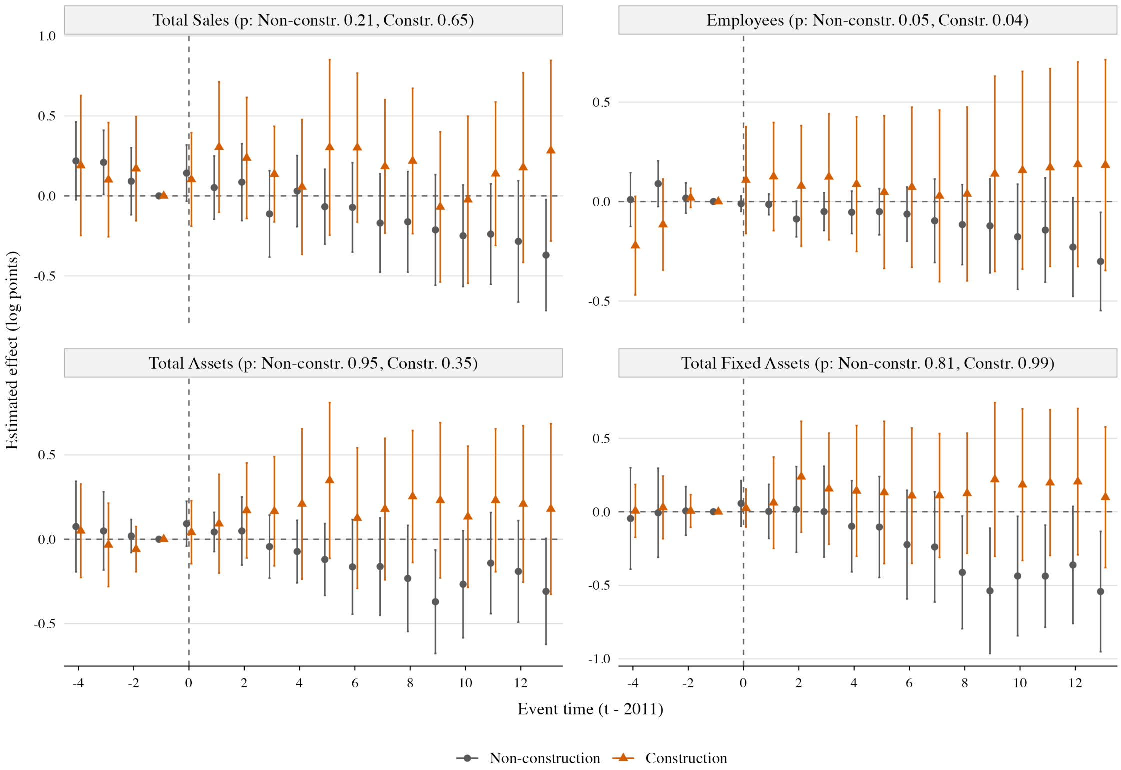

Notes: Event-study estimates from Equation (2), $\log(1 + Y)$ outcomes, estimated separately on the construction (59 firms) and non-construction (44 firms) subsamples of the balanced panel with the main-specification fixed effects. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero. Since outcomes are in logs, coefficients can be approximately interpreted as proportional changes relative to the control group.

Table A.6. Inference with Few Treated Clusters: construction vs. non-construction average DD (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t \times Construction_i$	140.980	2.693	124.155	67.162
	(149.151)	(3.808)	(141.352)	(77.681)
p-values	0.347 [0.414] (0.390)	0.481 [0.517] (0.530)	0.382 [0.478] (0.446)	0.389 [0.517] (0.509)
$Treat_i \times Post_t \times NonConstruction_i$	−1028.031***	−12.193***	−670.503***	−338.167**
	(386.909)	(4.613)	(221.199)	(150.272)
p-values	0.009 [0.010] (0.006)	0.010 [0.004] (0.014)	0.003 [0.004] (0.006)	0.027 [0.020] (0.039)
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t \times Construction_i$	88.071	0.404	52.927	63.022
	(133.539)	(3.247)	(129.734)	(65.347)
p-values	0.510 [0.549] (0.564)	0.901 [0.908] (0.918)	0.684 [0.760] (0.729)	0.336 [0.475] (0.420)
$Treat_i \times Post_t \times NonConstruction_i$	−922.293***	−4.361	−533.701***	−218.159**
	(327.456)	(5.044)	(163.647)	(85.141)
p-values	0.005 [0.005] (0.010)	0.388 [0.425] (0.471)	0.001 [0.005] (0.022)	0.011 [0.021] (0.042)

Notes: Each cell reports three p-values for the same null hypothesis: the analytic p-value based on standard errors clustered at the firm level, in square brackets the restricted (null-imposed) wild cluster bootstrap p-value (Rademacher weights drawn at the firm level, bootstrap-t, 4,999 draws), and in parentheses the randomization-inference p-value from 4,999 permutations of the firm-level treatment indicator within 2010-construction strata, holding the number of treated firms in each stratum fixed (studentized statistic; the randomization p-value tests the sharp null of no effect for any firm). Specification and samples are those of the construction-split average DD table (levels). Monetary variables are measured in 1 million yen.

Table A.7. Inference with Few Treated Clusters: joint pre-trend tests

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
All firms	0.537 [0.684] (0.613)	0.155 [0.217] (0.202)	0.855 [0.936] (0.968)	0.291 [0.367] (0.357)
Construction firms	0.514 [0.680] (0.654)	0.041 [0.093] (0.077)	0.958 [0.980] (0.978)	0.307 [0.404] (0.404)
Non-construction firms	0.759 [0.947] (0.844)	0.315 [0.517] (0.401)	0.815 [0.915] (0.955)	0.595 [0.705] (0.692)
Panel B: Balanced panel excluding 2012 and 2021				
All firms	0.380 [0.462] (0.461)	0.278 [0.335] (0.364)	0.611 [0.655] (0.624)	0.484 [0.567] (0.653)
Construction firms	0.949 [0.960] (0.979)	0.018 [0.040] (0.054)	0.502 [0.623] (0.658)	0.814 [0.894] (0.936)
Non-construction firms	0.371 [0.516] (0.503)	0.422 [0.561] (0.568)	0.263 [0.347] (0.289)	0.290 [0.380] (0.465)

Notes: Joint test that all pre-2011 event-study coefficients (event years −4 to −2; 2010 is the reference) equal zero, in the event-study specification of the main text (levels). Each cell reports three p-values for the same null hypothesis: the analytic p-value based on standard errors clustered at the firm level, in square brackets the restricted (null-imposed) wild cluster bootstrap p-value (Rademacher weights drawn at the firm level, bootstrap-t, 4,999 draws), and in parentheses the randomization-inference p-value from 4,999 permutations of the firm-level treatment indicator within 2010-construction strata, holding the number of treated firms in each stratum fixed (studentized statistic; the randomization p-value tests the sharp null of no effect for any firm). For the construction and non-construction rows the event study is estimated on the subsample, as in the construction event-study figures, and the treatment permutation is within the subsample. The analytic p-values for All firms match those reported in the event-study figure panels.

Table A.8. Randomization Inference with Composition-Matched Placebo Groups

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t$	0.053 (0.079)	0.165 (0.197)	0.093 (0.181)	0.159 (0.284)
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t$	0.049 (0.069)	0.563 (0.575)	0.061 (0.115)	0.164 (0.234)

Notes: Each cell reports two randomization-inference p-values for the same null hypothesis (the sharp null of no treatment effect for any firm), both from 4,999 permutations of the firm-level treatment indicator holding the number of treated firms in each stratum fixed (studentized statistic): first the baseline p-value with 2010-construction strata, as in the preceding tables, and in parentheses the composition-matched p-value with construction (2010) × firm-size strata, where the size split is the pre-disaster median split of 2007–2010 mean Total Sales used in the size tables and firms without pre-disaster sales data form a separate stratum. Every placebo treatment group therefore replicates the treated firms' construction and size composition. Specification, samples, and estimates are those of the average DD table (levels): firm and construction (2010) × year fixed effects; 103 firms (19 treated) in Panel A and 274 firms (25 treated) in Panel B; in Panel B the 2012 and 2021 fiscal years are excluded. In the balanced panel the treated firms fall into the strata as follows: 4 non-construction small, 6 non-construction large, 3 construction small, 6 construction large.

Table A.9. Average DD Estimates with Manufacturing × Year Fixed Effects (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t$	−386.345*	−4.241	−203.708*	−80.000
	(207.233)	(3.041)	(120.138)	(68.779)
Observations	1854	1854	1854	1854
Firms	103	103	103	103
Pre-treatment mean	1,646.58	41.09	1,192.75	395.35
Firm FE	✓	✓	✓	✓
Manufacturing × Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t$	−418.354**	−2.009	−264.543**	−78.682
	(183.008)	(3.175)	(112.909)	(54.333)
Observations	4384	4384	4384	4384
Firms	274	274	274	274
Pre-treatment mean	2,059.81	49.71	1,650.52	702.12
Firm FE	✓	✓	✓	✓
Manufacturing × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and manufacturing (industry classification as of 2010) × year fixed effects are included in all specifications; the manufacturing × year cells subsume ordinary year fixed effects and absorb aggregate shocks specific to manufacturing in every year. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. In Panel B, the 2012 and 2021 fiscal years are excluded, so the post period covers twelve rather than fourteen years.

Figure A.12: Dynamic Effects with Manufacturing × Year Fixed Effects: *Balanced panel*

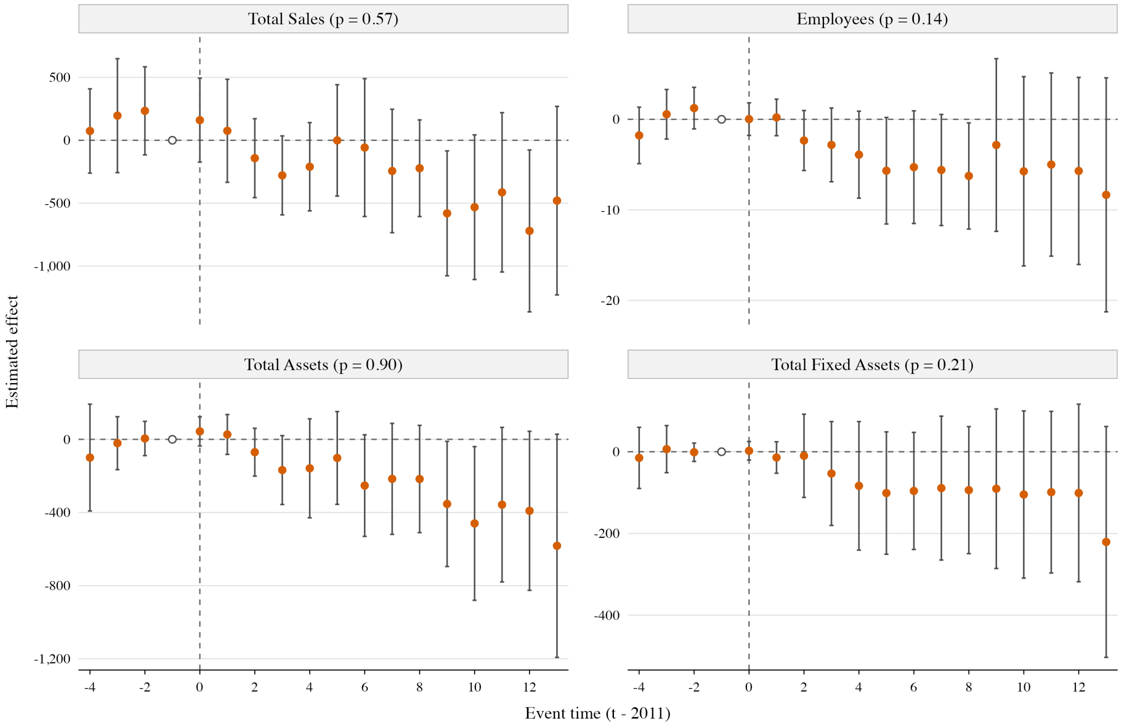

Notes: Event-study estimates from Equation (2) with the construction-by-year fixed effects replaced by manufacturing (industry classification as of 2010) × year fixed effects, level outcomes, balanced panel (103 firms, of which 19 treated). The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure A.13: Dynamic Effects with Manufacturing × Year Fixed Effects: *Balanced panel excluding 2012 and 2021*

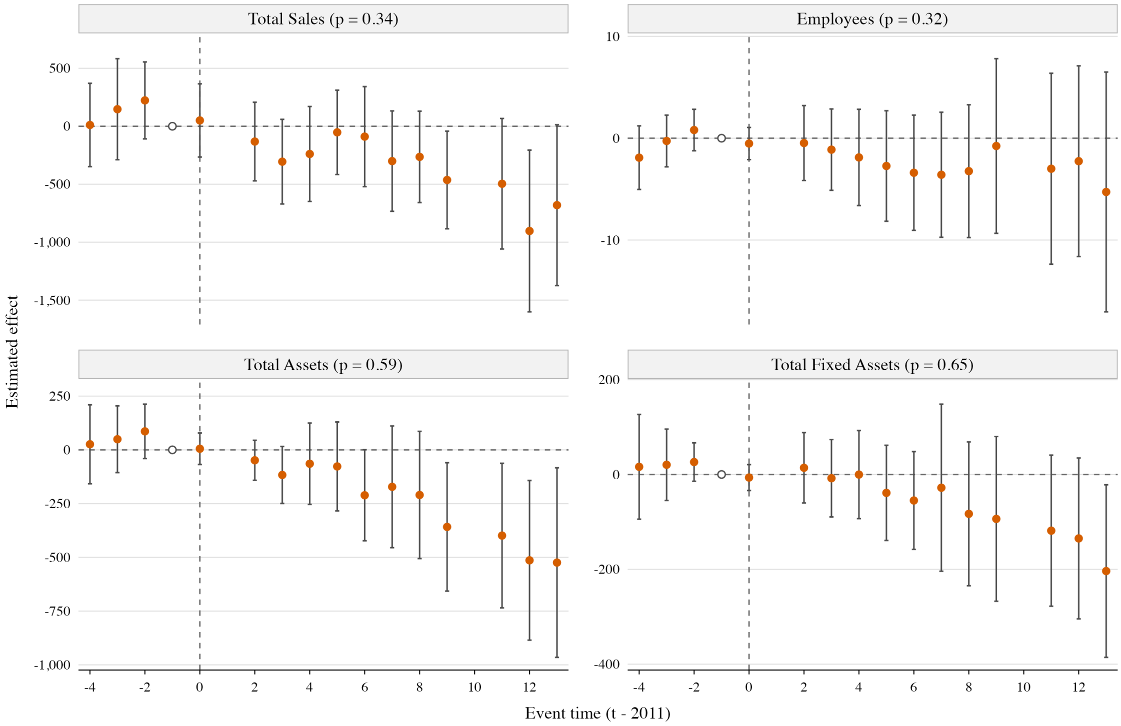

Notes: Event-study estimates from Equation (2) with the construction-by-year fixed effects replaced by manufacturing (2010) × year fixed effects, level outcomes, balanced panel reconstructed excluding the 2012 and 2021 fiscal years (274 firms, of which 25 treated); event years $k = 1$ and $k = 10$ are absent. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Table A.10. Heterogeneous DD Estimates: manufacturing vs. non-manufacturing (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t \times Manufacturing_i$	−495.212	−10.566*	−254.138	−6.741
	(352.586)	(6.341)	(161.922)	(76.164)
$Treat_i \times Post_t \times NonManufacturing_i$	−335.516	−1.288	−180.162	−114.205
	(253.278)	(3.126)	(158.097)	(93.090)
$Post_t \times Manufacturing_i$	−94.764	2.192	−207.487	−189.843**
	(224.609)	(2.723)	(141.784)	(79.008)
p-value: $\beta_M = \beta_{NM}$	0.71	0.19	0.74	0.37
Observations	1854	1854	1854	1854
Firms (manufacturing)	22	22	22	22
Firms (non-manufacturing)	81	81	81	81
Pre-treatment mean (manufacturing)	1,818.69	50.62	1,440.78	570.00
Pre-treatment mean (non-manufacturing)	1,599.83	38.50	1,125.38	347.92
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t \times Manufacturing_i$	−306.150	0.609	−315.273	−36.365
	(224.168)	(6.365)	(192.527)	(87.021)
$Treat_i \times Post_t \times NonManufacturing_i$	−499.907*	−3.912	−227.672*	−109.438
	(268.586)	(2.844)	(134.665)	(68.199)
$Post_t \times Manufacturing_i$	−340.054	−1.712	−35.930	−87.045
	(211.847)	(3.116)	(150.125)	(64.976)
p-value: $\beta_M = \beta_{NM}$	0.58	0.52	0.71	0.51
Observations	4384	4384	4384	4384
Firms (manufacturing)	78	78	78	78
Firms (non-manufacturing)	196	196	196	196
Pre-treatment mean (manufacturing)	2,482.00	71.99	2,251.79	1,043.36
Pre-treatment mean (non-manufacturing)	1,891.80	40.85	1,411.24	566.33
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and year fixed effects are included in all specifications, and the $Post_t \times Manufacturing_i$ main effect is included as a control; the $Treat_i \times Manufacturing_i$ interaction is constant within firm and absorbed by the firm fixed effects. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment means pool treated and control firm-years 2007–2010 within each group. The p-value row tests $\beta_M = \beta_{NM}$, the equality of the manufacturing and non-manufacturing treatment coefficients, from the same specification re-parameterised with non-manufacturing firms as the baseline, so that the manufacturing term is the between-group difference. In Panel B, the 2012 and 2021 fiscal years are excluded, so the post period covers twelve rather than fourteen years.

Figure A.14: Event-study: Manufacturing versus Non-Manufacturing: *Balanced panel*

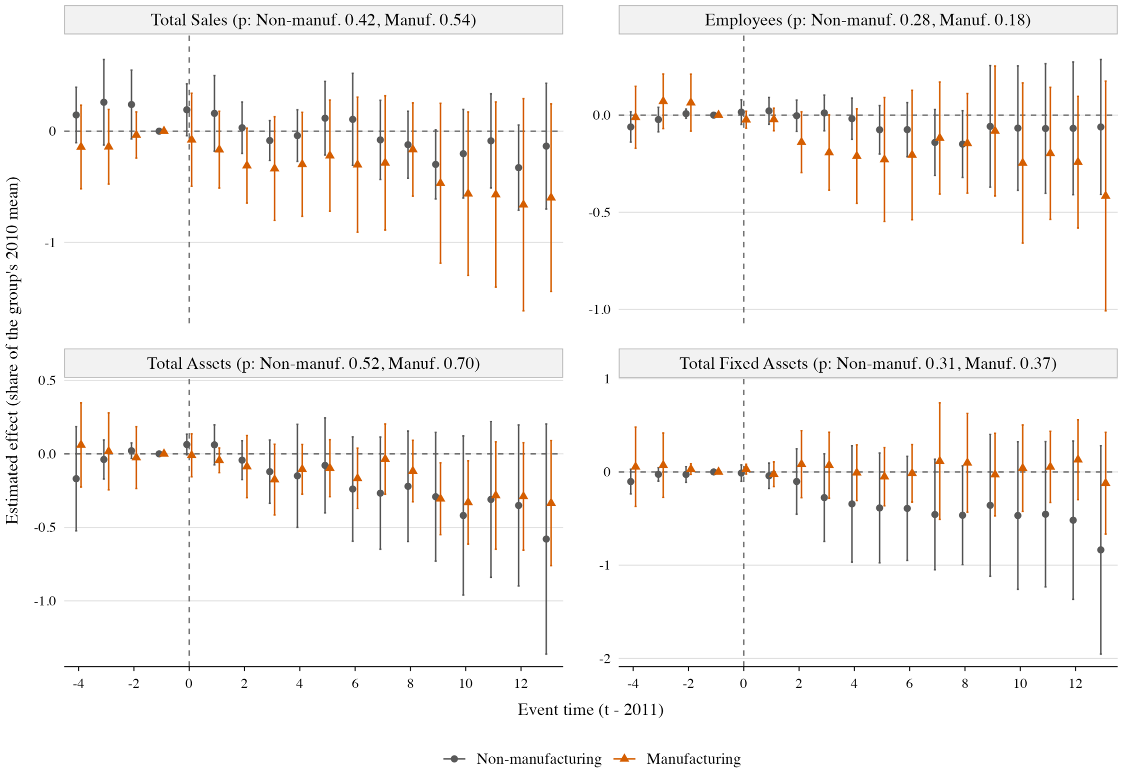

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the manufacturing (22 firms) and non-manufacturing (81 firms) subsamples of the balanced panel with firm and year fixed effects; within a subsample, the year effects are the group's own year effects. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Figure A.15: Event-study: Manufacturing versus Non-Manufacturing: *Balanced panel excluding 2012 and 2021*

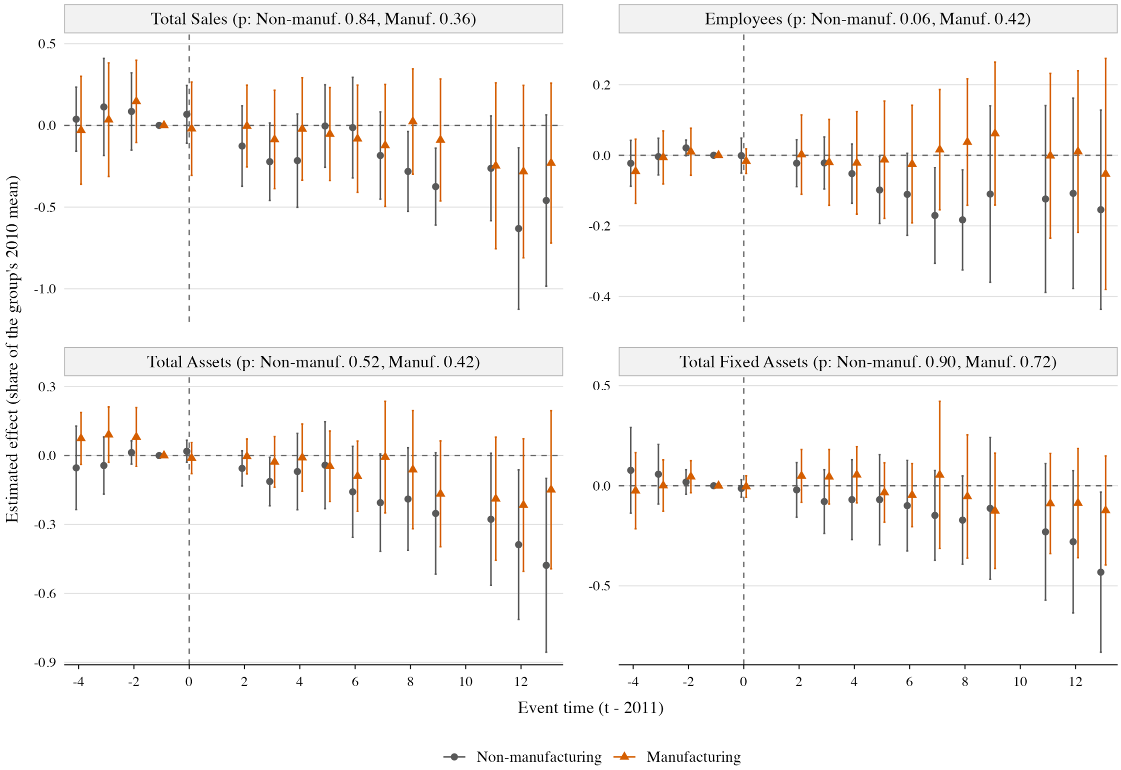

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the manufacturing (78 firms) and non-manufacturing (196 firms) subsamples of the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, with firm and year fixed effects; event years $k = 1$ and $k = 10$ are absent. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Table A.11. Firm Exit after the GEJE: Registry-Based Hazard Estimates

	Coefficient	Std. error	Observations
Panel A: Baseline			
$Treat_i \times Post_t$	−0.306***	(0.044)	19 615 222
$Treat_i \times D_t^{2011\text{–}2012}$	−0.546***	(0.089)	19 615 222
$Treat_i \times D_t^{2013\text{–}2016}$	−0.431***	(0.076)	19 615 222
$Treat_i \times D_t^{2017\text{–}2025}$	−0.172***	(0.043)	19 615 222
***Panel B: Heterogeneity** ($Treat_i \times Post_t \times$ **group**)*			
Firm size: 4 or fewer employees	−0.381***	(0.062)	19 615 222
Firm size: 5–9 employees	−0.352***	(0.061)	19 615 222
Firm size: 10–29 employees	−0.240***	(0.060)	19 615 222
Firm size: 30 or more employees	0.032	(0.052)	19 615 222
Firm size: missing	−0.487***	(0.113)	19 615 222
Industry: construction	−0.314***	(0.055)	19 615 222
Industry: manufacturing	−0.009	(0.074)	19 615 222
Industry: other	−0.360***	(0.056)	19 615 222
***Panel C: Robustness** ($Treat_i \times Post_t$)*			
Excluding the 2012 and 2021 vintages	−0.318***	(0.042)	17 371 142
Censoring exits identified only in 2025	−0.329***	(0.042)	19 587 178
Excluding relocated firms	−0.346***	(0.049)	17 982 123
Firms in ≥ 3 pre-2011 vintages	−0.265***	(0.044)	16 385 720
Failure-only outcome, 2013–2023 ($Treat_i$)	−0.066***	(0.012)	11 674 435

Notes: Registry-based extensive-margin analysis; estimates from the discrete-time linear-probability exit hazard of the TSR registry cohort (1,458,060 firms first observed by 2010 with a geocodable address; about 19.6 million firm-years at risk). Treatment is physical exposure at the firm's latest pre-2011 address: location inside or within 500 m of the GSI tsunami inundation footprint, or interpolated JMA instrumental seismic intensity of 5.5 or more. The outcome is 100 × an indicator that the year is the first registry vintage after the firm's last appearance, a broad exit measure that bundles voluntary closure, merger, and loss of TSR coverage; coefficients are in percentage points of the annual exit hazard. All specifications include pre-2011 industry × year and pre-2011 employment-size-bin × year fixed effects; standard errors are clustered at the municipality. In Panel C, the final row replaces the outcome with exits matched to a TSR bankruptcy record; because bankruptcy coverage begins in July 2012, that row reports a post-period treated–control level difference over 2013–2023 rather than a difference-in-differences.

Figure A.16: Head-Office Relocation after 2011

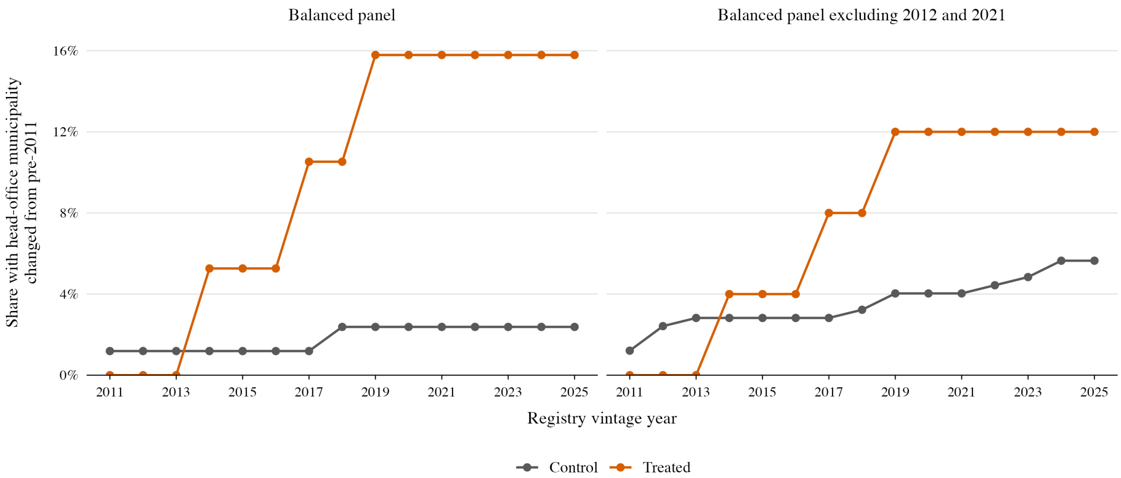

Notes: Share of estimation-sample firms whose head-office municipality (per-vintage JIS address code in the TSR registry) in the indicated vintage year differs from the municipality at the firm's latest pre-2011 vintage, by treatment group. Each year's denominator is the firms present in that vintage with a valid pre-2011 address code. Moves within a municipality are undetectable.

Table A.12. Excluding relocated firms: average DD Estimates (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t$	−354.560	−4.896	−251.151	−133.518
	(222.337)	(3.758)	(153.556)	(90.314)
Observations	1764	1764	1764	1764
Firms	98	98	98	98
Pre-treatment mean	1,630.58	41.55	1,192.96	399.33
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t$	−370.269*	−1.339	−210.634*	−89.546
	(201.768)	(3.491)	(127.029)	(63.724)
Observations	4112	4112	4112	4112
Firms	257	257	257	257
Pre-treatment mean	1,999.94	48.31	1,599.12	696.92
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Firms whose head-office municipality (per-vintage JIS address code in the TSR registry) in any post-2011 vintage differs from the municipality at their latest pre-2011 vintage are excluded. Within-municipality moves are undetectable, and the registry address refers to the head office only. Firms without a valid pre-2011 address code are retained. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. In Panel B, the 2012 and 2021 fiscal years are excluded.

Figure A.17: Dynamic Effects Excluding Relocated Firms: *Balanced panel*

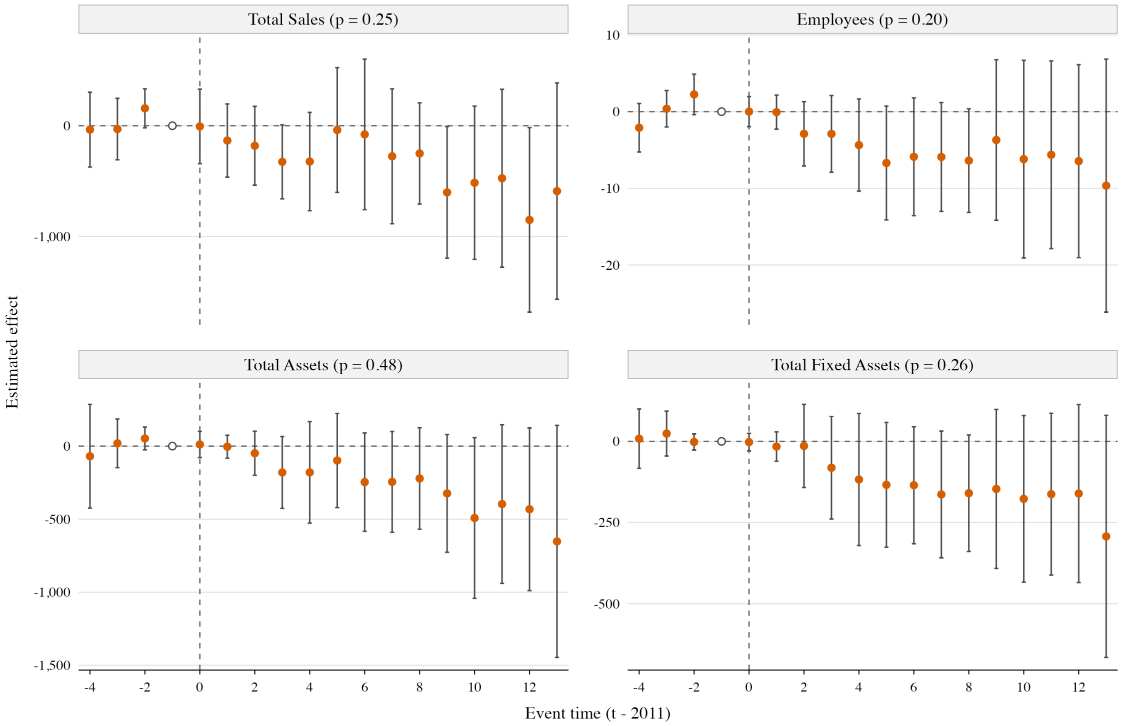

Notes: Event-study estimates from Equation (2), level outcomes, balanced panel excluding firms whose head-office municipality in any post-2011 registry vintage differs from its latest pre-2011 municipality (98 firms remain, of which 16 treated). The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure A.18: Dynamic Effects Excluding Relocated Firms: *Balanced panel excluding 2012 and 2021*

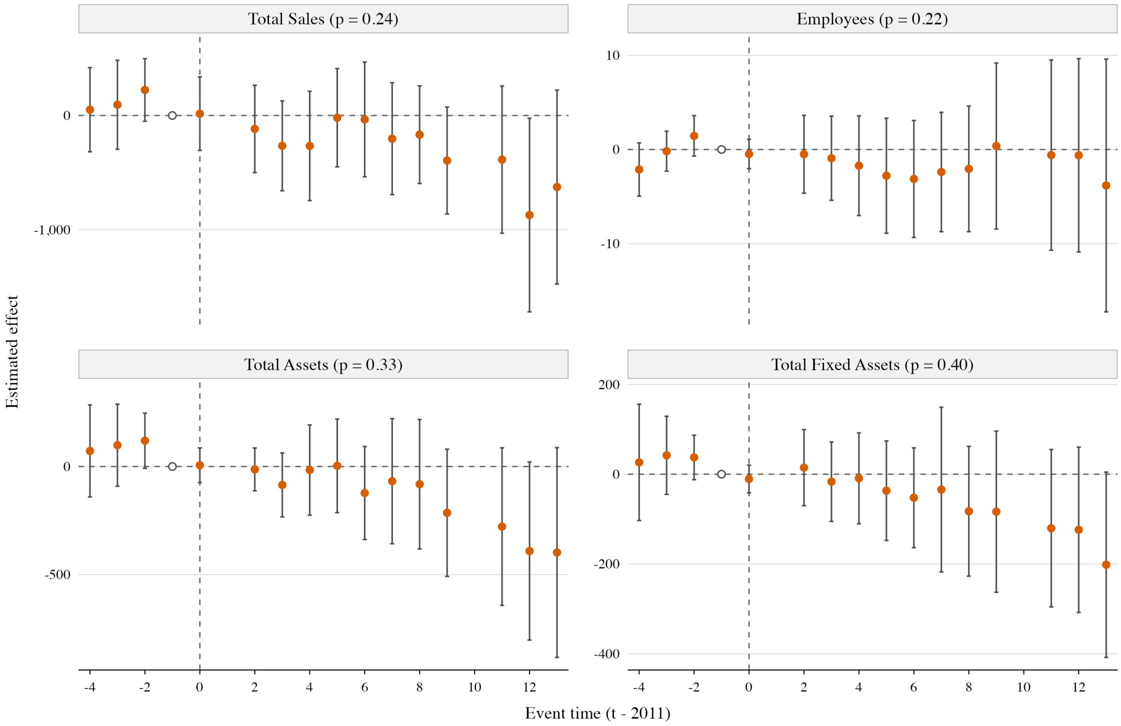

Notes: Event-study estimates from Equation (2), level outcomes, balanced panel reconstructed excluding the 2012 and 2021 fiscal years and excluding firms whose head-office municipality in any post-2011 registry vintage differs from its latest pre-2011 municipality (257 firms remain, of which 22 treated); event years $k = 1$ and $k = 10$ are absent. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

A.6. Physical (GIS-Based) Exposure

This appendix collects the supporting material for the physical-exposure treatment of Section 4.1. A firm is treated if its geocoded pre-disaster address (registry address pinned to 2010, with the earliest available registry address as a fallback) lies in or within 500 meters of the tsunami-inundation footprint mapped by the Geospatial Information Authority of Japan, or in the zone of instrumental seismic intensity of at least 5.5 interpolated from Japan Meteorological Agency stations. The intensity measure follows Hanaoka et al. (2018): for each address we take the weighted average of the instrumental seismic intensities (*keisoku shindo*, the continuous scale underlying the reported Japan Meteorological Agency intensity classes) recorded for the March 11, 2011 main shock at the three closest stations, weighting each station by the inverse of its distance, with no distance cutoff and distances computed in projected metres (JGD2011 Japan Plane Rectangular CS IX, EPSG:6677). The one departure from Hanaoka et al. (2018) is the reference point: they measure distances from the city hall of the respondent's municipality, whereas we measure them from the firm's own geocoded address. Two further choices are ours rather than theirs: they use the intensity continuously, while we dichotomize it at 5.5, and we combine the shaking zone with the tsunami footprint by a logical or. The samples, outcomes, and specification are identical to the main design; only the treatment changes. Self-reported exposure plays no role: survey-treated firms that are not physically exposed serve as controls, and firms without a geocodable address are dropped (none on either panel). Only nine firms of the 103-firm balanced panel (ten of the 274-firm panel excluding 2012 and 2021) are physically treated, so the estimates pin down direction rather than magnitude. The average level estimates are reported in the main text (Panel A of Table 7 in Section 4.1) and the log estimates in Table A.13 below; Figures A.21–A.24 plot the corresponding event studies. Figures A.19 and A.20 map the two ingredients of the treatment.

Figure A.25 traces the relationship between the two exposure instruments on the full set of geocoded respondents; Figures A.26 and A.27 repeat the comparison within the balanced and ex-2012/2021 estimation samples.

Table A.13. Physical (GIS-based) exposure: average DD Estimates (log)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t$	0.199	0.247	0.380**	0.277
	(0.139)	(0.212)	(0.181)	(0.198)
Observations	1854	1854	1854	1854
Firms	103	103	103	103
Treated firms	9	9	9	9
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Treatment is physical exposure at the firm's geocoded pre-disaster address (registry address pinned to 2010, with the earliest available registry address as a fallback): a firm is treated if the address lies within 500 m of the tsunami-inundation footprint mapped by the Geospatial Information Authority of Japan or in the zone of instrumental seismic intensity of at least 5.5 interpolated from Japan Meteorological Agency station records. Self-reported exposure plays no role: survey-treated firms that are not physically exposed serve as controls, and firms without a geocodable address are excluded. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Outcomes are $\log(1+Y)$, so coefficients can be approximately interpreted as proportional changes. The estimation sample is the balanced panel. Appendix Table B.5 reports the same specification on the balanced panel excluding 2012 and 2021.

Figure A.19: The 2011 tsunami-inundation footprint

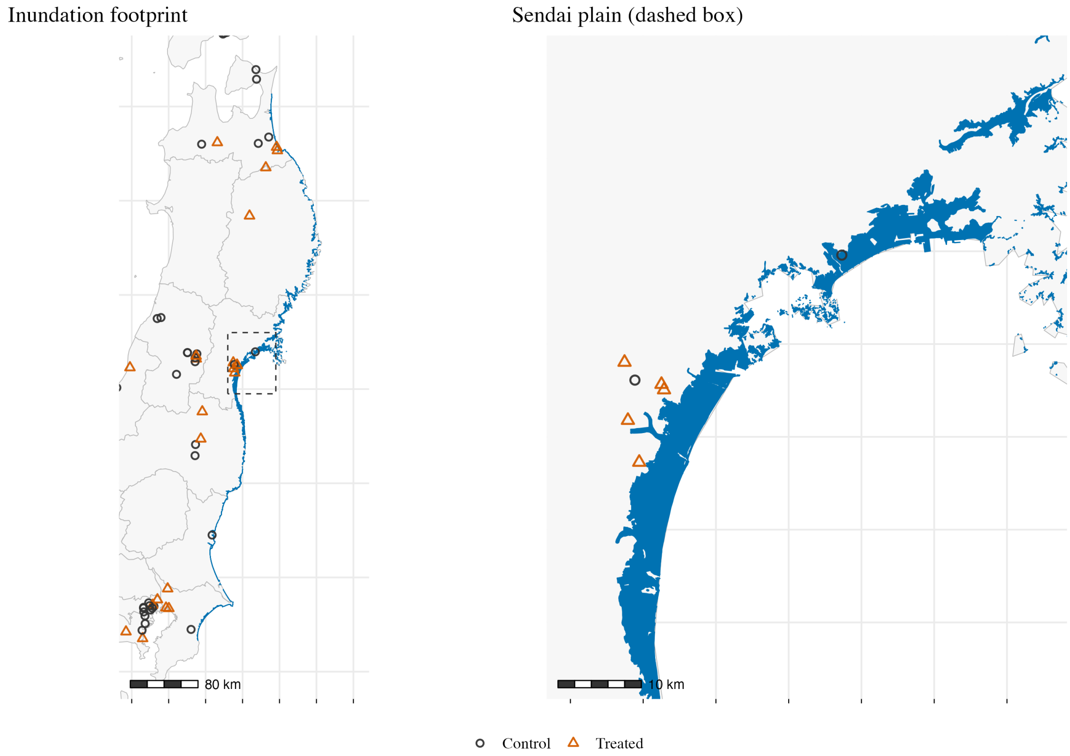

Notes: The inundation extent of the March 2011 tsunami mapped by the Geospatial Information Authority of Japan (D1-No.691, 1:100,000 inundation overview map), in blue. The left panel shows the full footprint along the Pacific coast; the right panel enlarges the dashed box around the Sendai plain, where the inundated band is widest. Open triangles (treated) and open circles (control) are the sample firms of the panel excluding 2012 and 2021 (274 firms, a superset of the 103-firm balanced panel; only firms within each map window are shown) at their geocoded pre-disaster addresses; treatment refers to the self-reported survey exposure. Grey lines are prefecture boundaries; the scale bars give distance in kilometers. A firm is tsunami-exposed when its address lies inside or within 500 meters of the footprint.

Figure A.20: Instrumental seismic intensity of the March 2011 main shock

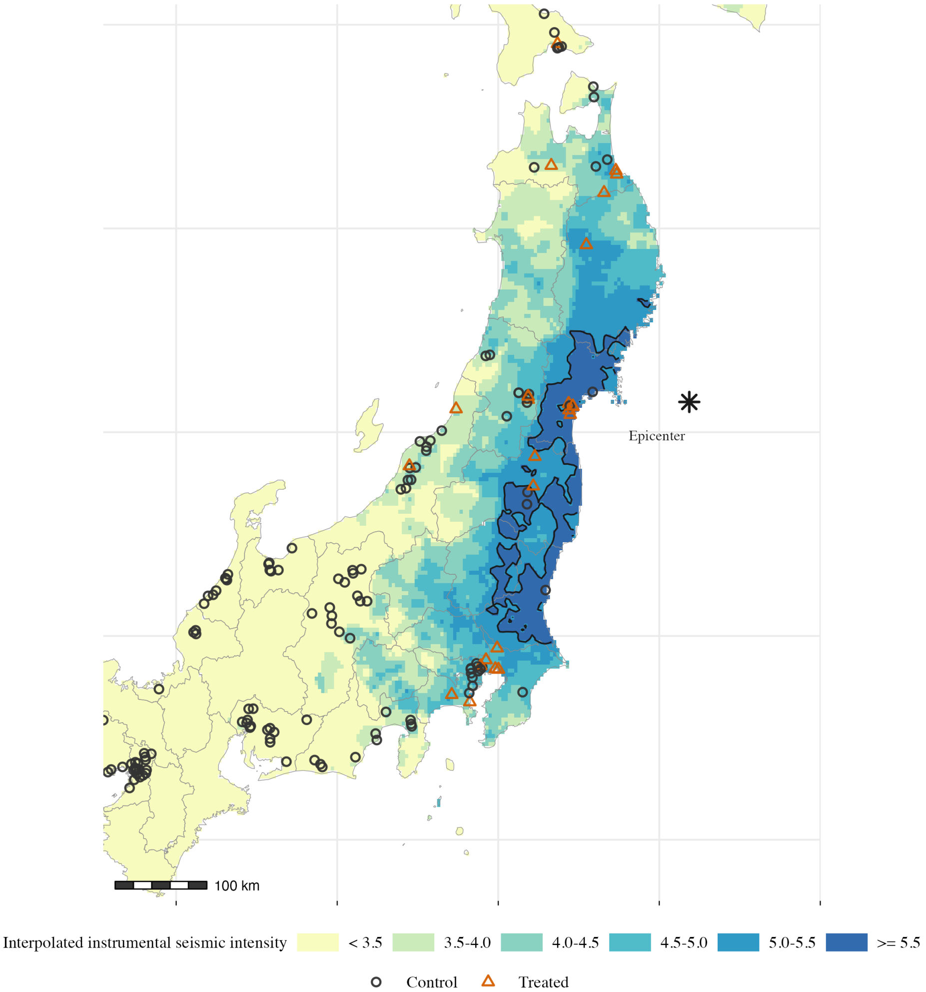

Notes: The shaded surface interpolates the instrumental seismic intensity recorded by Japan Meteorological Agency stations during the March 11, 2011 mainshock onto a grid, using the same estimator as the physical-exposure treatment (inverse-distance weighting over the three nearest stations, following Hanaoka et al., 2018); the black contour marks interpolated intensity 5.5, the shaking cutoff of the treatment, and the star marks the epicenter. Open triangles (treated) and open circles (control) are the sample firms of the panel excluding 2012 and 2021 (274 firms, a superset of the 103-firm balanced panel; only firms within the map window are shown) at their geocoded pre-disaster addresses; treatment refers to the self-reported survey exposure. Grey lines are prefecture boundaries; the scale bar gives distance in kilometers.

Figure A.21: Dynamic Effects under Physical (GIS-Based) Exposure: *Balanced panel, level outcomes*

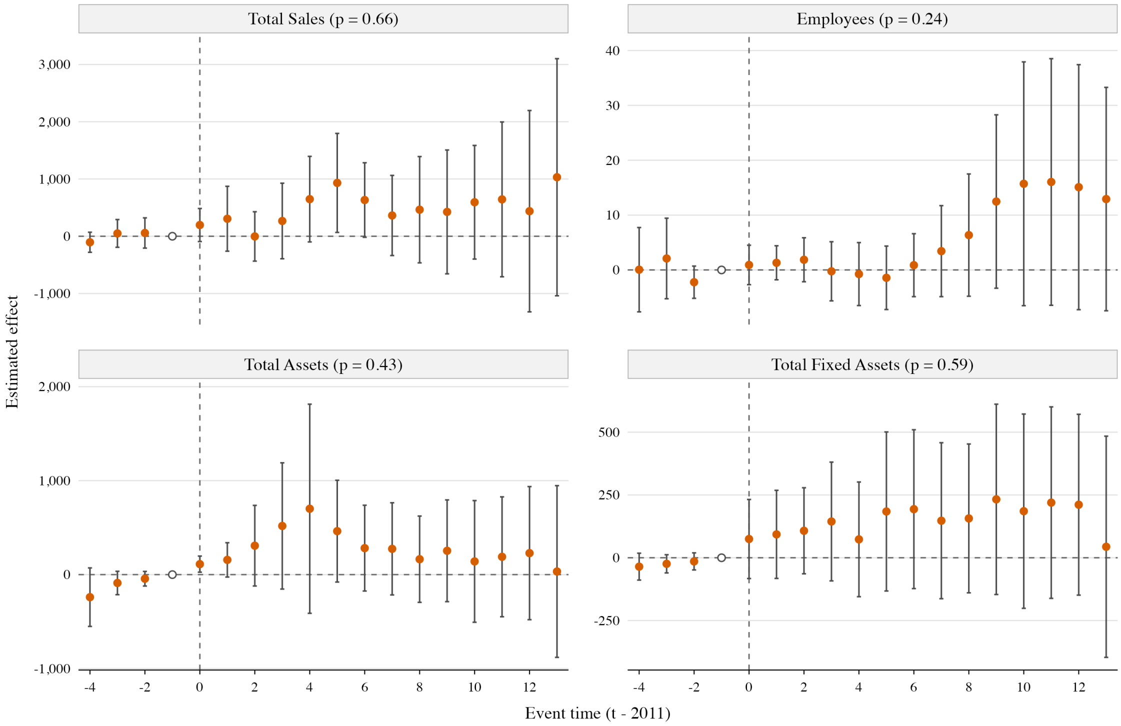

Notes: Event-study estimates from Equation (2), level outcomes, balanced panel (103 firms, of which 9 physically treated), under the physical-exposure treatment of this appendix: a firm is treated if its geocoded pre-disaster address lies within 500 meters of the GSI tsunami-inundation footprint or in the zone of interpolated instrumental seismic intensity of at least 5.5, and survey-treated firms that are not physically exposed serve as controls. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure A.22: Dynamic Effects under Physical (GIS-Based) Exposure: *Balanced panel excluding 2012 and 2021, level outcomes*

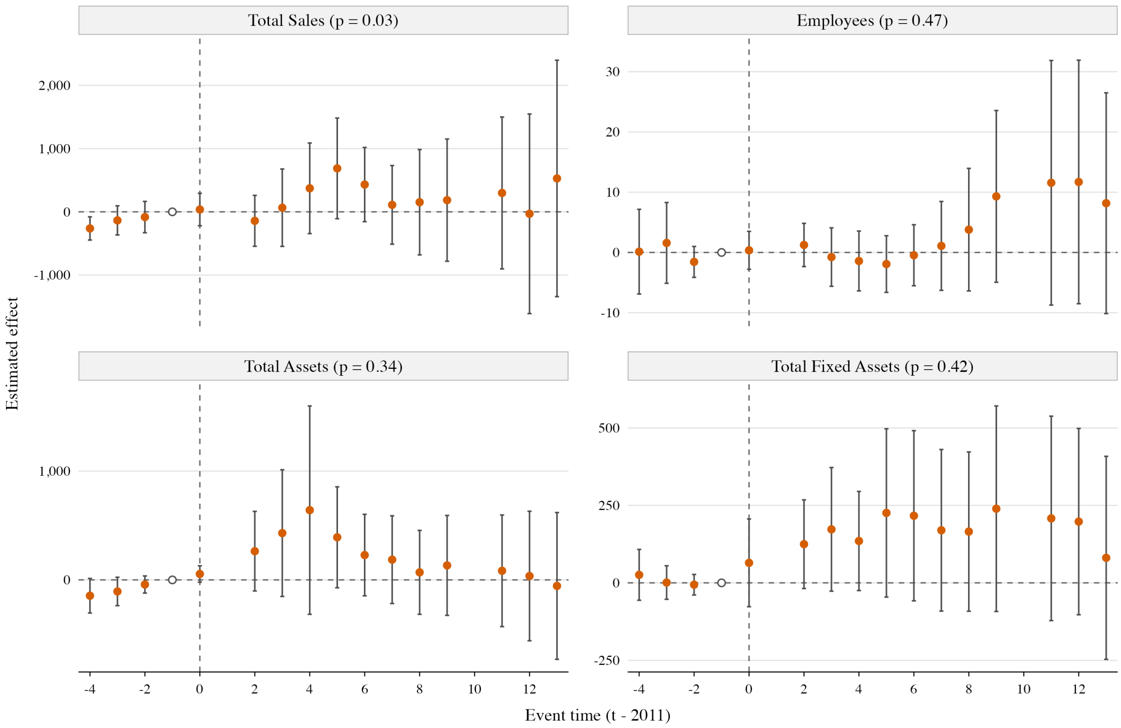

Notes: Event-study estimates from Equation (2), level outcomes, balanced panel reconstructed excluding the 2012 and 2021 fiscal years (274 firms, of which 10 physically treated), under the physical-exposure treatment of this appendix; event years $k = 1$ and $k = 10$ are therefore absent. A firm is treated if its geocoded pre-disaster address lies within 500 meters of the GSI tsunami-inundation footprint or in the zone of interpolated instrumental seismic intensity of at least 5.5, and survey-treated firms that are not physically exposed serve as controls. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure A.23: Dynamic Effects under Physical (GIS-Based) Exposure: *Balanced panel, log outcomes*

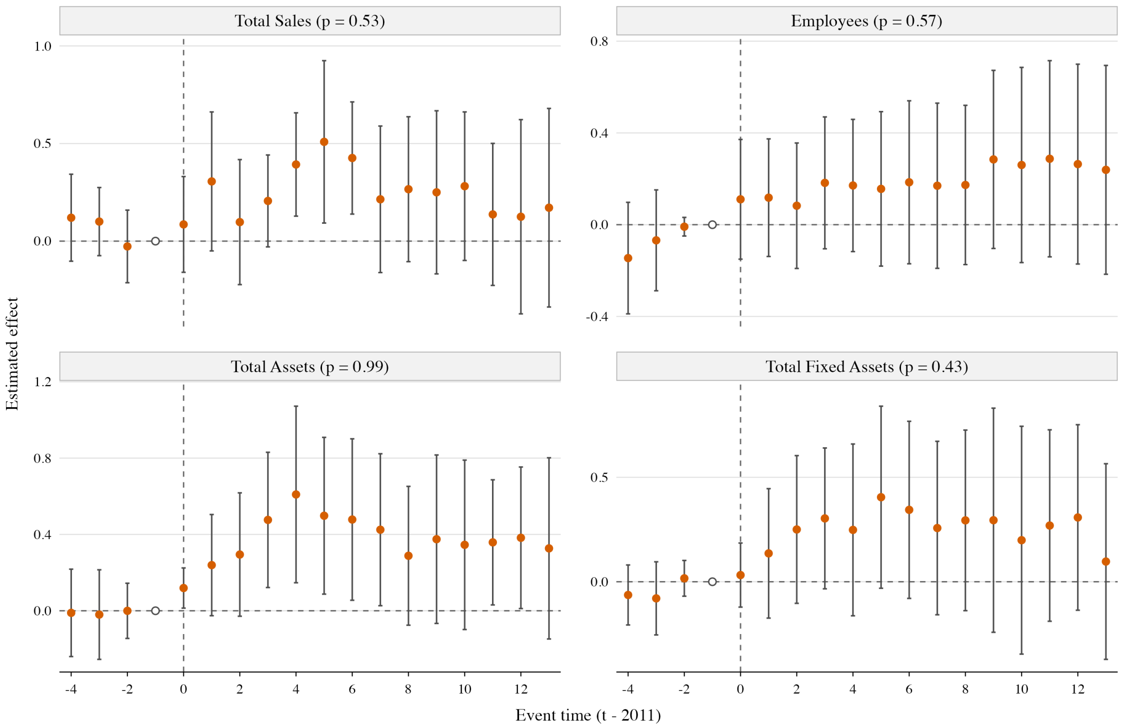

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, balanced panel (103 firms, of which 9 physically treated), under the physical-exposure treatment of this appendix: a firm is treated if its geocoded pre-disaster address lies within 500 meters of the GSI tsunami-inundation footprint or in the zone of interpolated instrumental seismic intensity of at least 5.5, and survey-treated firms that are not physically exposed serve as controls. The reference year is 2010 ($k=-1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Coefficients can be approximately interpreted as proportional changes.

Figure A.24: Dynamic Effects under Physical (GIS-Based) Exposure: *Balanced panel excluding 2012 and 2021, log outcomes*

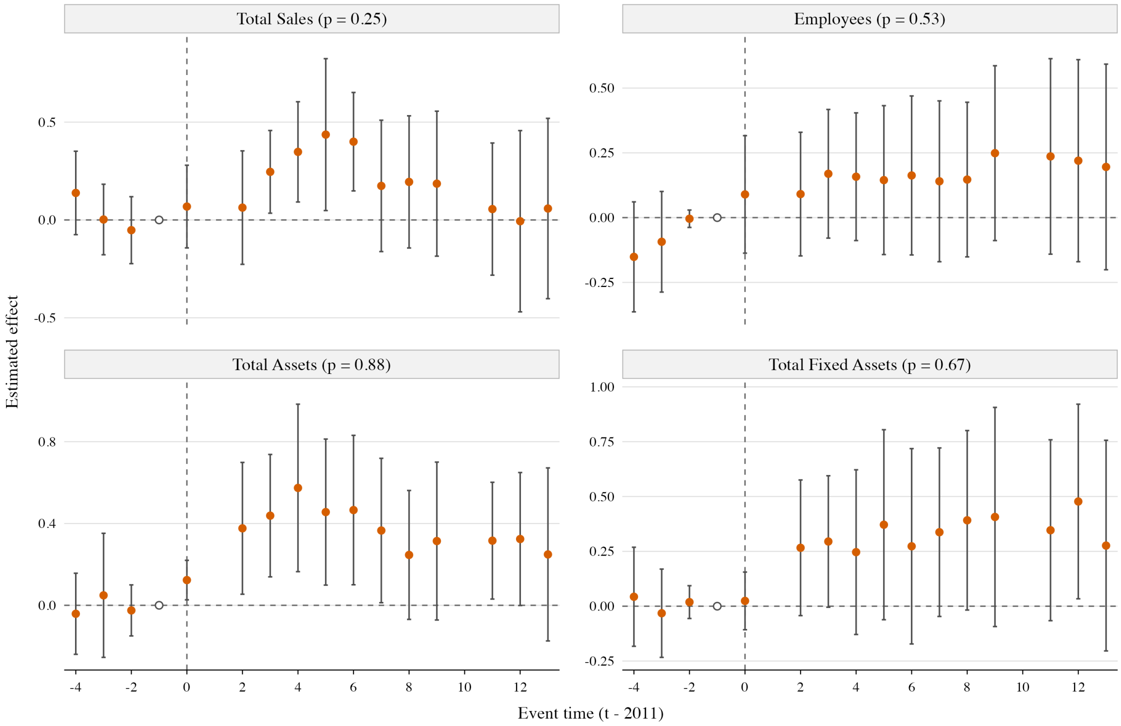

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, balanced panel reconstructed excluding the 2012 and 2021 fiscal years (274 firms, of which 10 physically treated), under the physical-exposure treatment of this appendix; event years $k=1$ and $k=10$ are therefore absent. A firm is treated if its geocoded pre-disaster address lies within 500 meters of the GSI tsunami-inundation footprint or in the zone of interpolated instrumental seismic intensity of at least 5.5, and survey-treated firms that are not physically exposed serve as controls. The reference year is 2010 ($k=-1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Coefficients can be approximately interpreted as proportional changes.

Figure A.25: Self-reported versus physical exposure

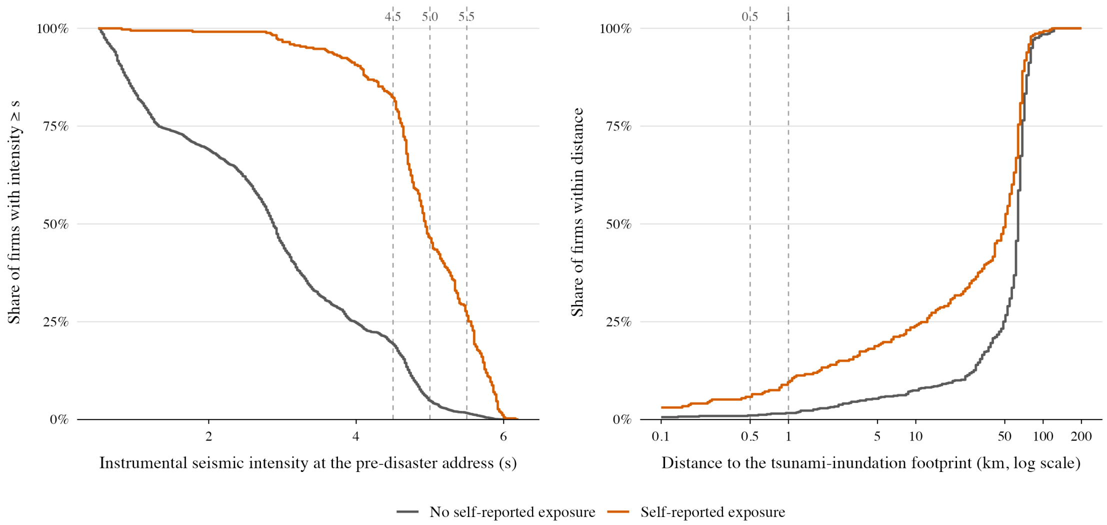

Notes: Survivor curves of physical exposure at the firm's geocoded pre-disaster address (registry address pinned to 2010, with the earliest available vintage as a fallback), for the 343 respondents reporting GEJE exposure and the 3,167 respondents reporting none. Left: share of firms whose instrumental seismic intensity (inverse-distance-weighted interpolation from the three nearest JMA stations, following Hanaoka et al., 2018) is at least s; dashed lines mark intensities 4.5, 5.0, and 5.5, the last being the severe-shaking cutoff of the physical-exposure treatment. Right: share of firms whose address lies within a given distance of the GSI tsunami-inundation footprint (log scale); dashed lines mark 0.5 and 1 km, the footprint buffer of the physical-exposure treatment. Self-reported exposure is strongly related to physical exposure — 82 percent of self-reporting firms experienced intensity of at least 4.5, against 19 percent of non-reporting firms — but only 27 percent of self-reporting firms cross the severe 5.5 threshold and only 10 percent lie within 1 km of the inundation footprint, so the self-reported treatment pools severely and mildly exposed firms.

Figure A.26: Self-reported versus physical exposure: *Balanced panel*

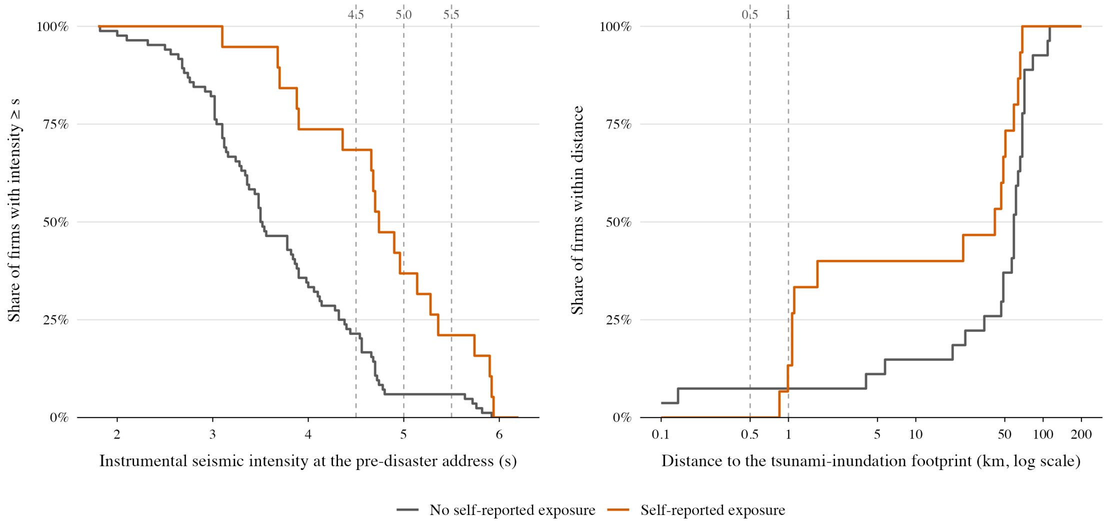

Notes: Same design as Figure A.25, restricted to the balanced panel (19 firms reporting GEJE exposure, 84 reporting none). Within this sample, 68 percent of self-reporting firms experienced instrumental seismic intensity of at least 4.5, against 21 percent of non-reporting firms, and 21 percent (6 percent) cross the severe 5.5 threshold that defines physical treatment. With 19 treated firms the curves are correspondingly coarse.

Figure A.27: Self-reported versus physical exposure: *Balanced panel excluding 2012 and 2021*

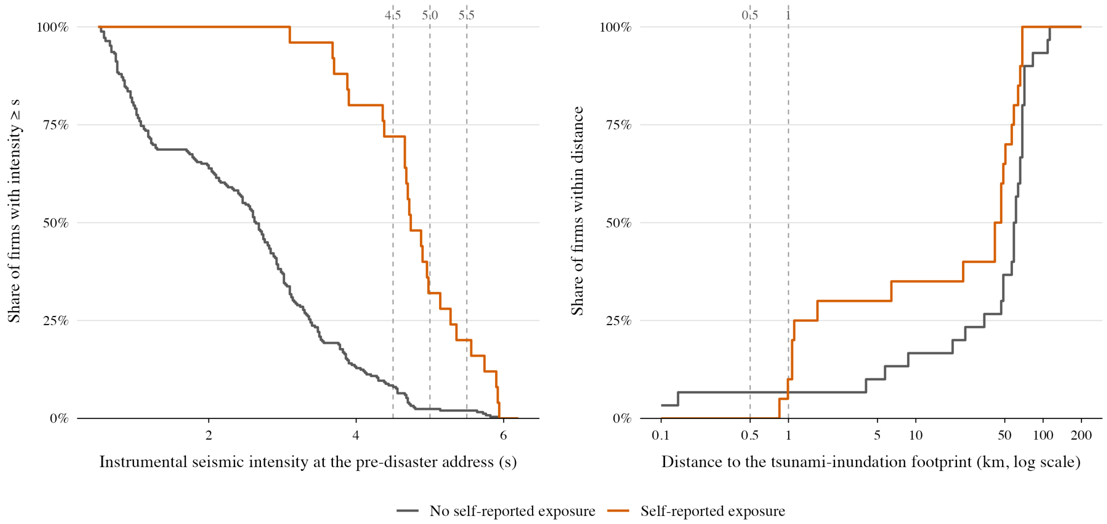

Notes: Same design as Figure A.25, restricted to the balanced panel reconstructed excluding the 2012 and 2021 fiscal years (25 firms reporting GEJE exposure, 249 reporting none). Within this sample, 72 percent of self-reporting firms experienced instrumental seismic intensity of at least 4.5, against 8 percent of non-reporting firms, and 20 percent (2 percent) cross the severe 5.5 threshold that defines physical treatment.

B. Online Appendix

Contents of the Appendix

B.1. The Balanced Panel Excluding 2012 and 2021

The 2012 and 2021 TSR vintages carry coverage discontinuities (Section 2.2), and dropping those two fiscal years from the balance requirement rebuilds the panel on a much larger cross-section: 274 firms, of which 25 treated, against the 103 firms (19 treated) of the balanced panel used throughout the main text. Because the balanced panel is a subset of this one, and because both the treated and the control group change between them, the two are not independent replications of each other. They are, however, differently composed in ways that matter for the reading of the estimates: the two samples differ in the geography of the control group (Figure A.2), in industry composition (Figure 2), and in the treated share (Figure A.3). This appendix re-estimates the main-text exhibits on the larger sample. Throughout, the point estimates are close to their balanced-panel counterparts and more precisely estimated, which suggests that the losses documented in Section 3 are not generated by the particular set of control firms that the stricter balance requirement selects.

The average DD (Table B.1) reproduces the headline result with tighter standard errors: total sales fall by 476.8 million yen (23 percent of the pre-treatment mean) and total assets by 275.1 million yen (17 percent), both now significant at the 5 percent level, against the 10 percent level on the balanced panel, while the fixed-asset estimate (−94.2 million yen, 13 percent) remains insignificant. The construction split (Table B.2) confirms the sectoral divergence: total sales among treated non-construction firms fall by 922.3 million yen (34 percent, 1 percent level) and total assets by 533.7 million yen (24 percent, 1 percent level), though the employment effect loses significance. The size split (Table B.3) sharpens the concentration of the losses among larger treated firms: 597.8 million yen in total sales and 327.9 million yen in total assets, both at the 5 percent level, with the small-firm coefficients still indistinguishable from zero. The two location-based treatment definitions of Section 4 behave on this sample as they do on the balanced panel (Tables B.4, B.5, and B.6).

Table B.1. Average DD Estimates (level): *balanced panel excluding 2012 and 2021*

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t$	−476.833**	−2.260	−275.062**	−94.189
	(213.075)	(3.192)	(120.334)	(61.219)
Observations	4384	4384	4384	4384
Firms	274	274	274	274
Pre-treatment mean	2,059.81	49.71	1,650.52	702.12
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. The estimation sample is the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, so the post period covers twelve rather than fourteen years. Table 2 reports the same specification on the balanced panel.

Table B.2. Heterogeneous DD Estimates: construction vs. non-construction (level): *balanced panel excluding 2012 and 2021*

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t \times Construction_i$	88.071	0.404	52.927	63.022
	(133.539)	(3.247)	(129.734)	(65.347)
$Treat_i \times Post_t \times NonConstruction_i$	−922.293***	−4.361	−533.701***	−218.159**
	(327.456)	(5.044)	(163.647)	(85.141)
p-value: $\beta_C = \beta_N$	< 0.01	0.43	< 0.01	< 0.01
Observations	4384	4384	4384	4384
Firms (construction)	123	123	123	123
Firms (non-construction)	151	151	151	151
Pre-treatment mean (construction)	1,216.17	28.25	958.79	260.10
Pre-treatment mean (non-construction)	2,747.02	67.20	2,213.98	1,062.19
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications; the construction × year cells absorb the $Post_t \times Construction_i$ main effect. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment means pool treated and control firm-years 2007–2010 within each group. The p-value row tests $\beta_C = \beta_N$, the equality of the two group coefficients, from the same specification re-parameterised with construction firms as the baseline, so that the non-construction term is the between-group difference. The estimation sample is the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, so the post period covers twelve rather than fourteen years. Table 3 reports the same specification on the balanced panel.

Table B.3. Heterogeneous DD Estimates: small vs. large firms (level): *balanced panel excluding 2012 and 2021*

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t \times Small_i$	−63.860	0.373	9.055	0.032
	(149.016)	(2.808)	(94.379)	(62.789)
$Treat_i \times Post_t \times Large_i$	−597.823**	−2.582	−327.919**	−110.996
	(274.319)	(4.376)	(155.944)	(77.462)
$Post_t \times Large_i$	374.405***	4.677***	412.536***	140.490***
	(108.997)	(1.597)	(65.532)	(31.288)
p-value: $\beta_L = \beta_S$	0.06	0.57	0.06	0.24
Observations	4384	4384	4384	4384
Firms (small)	55	55	55	55
Firms (large)	219	219	219	219
Pre-treatment mean (small)	165.46	9.98	131.91	53.75
Pre-treatment mean (large)	2,535.56	59.69	2,031.90	864.96
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. Firm and construction (2010) × year fixed effects are included; the $Post_t \times Large_i$ interaction controls for differential trends of large firms, and its coefficient is reported. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment means pool treated and control firm-years 2007–2010 within each group. The p-value row tests $\beta_L = \beta_S$ (Equation (4)), the equality of the large- and small-firm treatment coefficients, from the same specification re-parameterised with small firms as the baseline, so that the large-firm term is the between-group difference. The estimation sample is the balanced panel reconstructed excluding the 2012 and 2021 fiscal years. Table 4 reports the same specification on the balanced panel.

Table B.4. Physical (GIS-based) exposure: average DD Estimates (level): *balanced panel excluding 2012 and 2021*

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t$	344.728	3.517	278.557	161.507
	(337.918)	(3.862)	(213.403)	(116.875)
Observations	4384	4384	4384	4384
Firms	274	274	274	274
Treated firms	10	10	10	10
Pre-treatment mean	2,059.81	49.71	1,650.52	702.12
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Treatment is physical exposure at the firm's geocoded pre-disaster address (registry address pinned to 2010, with the earliest available registry address as a fallback): a firm is treated if the address lies within 500 m of the tsunami-inundation footprint mapped by the Geospatial Information Authority of Japan or in the zone of instrumental seismic intensity of at least 5.5 interpolated from Japan Meteorological Agency station records. Self-reported exposure plays no role: survey-treated firms that are not physically exposed serve as controls, and firms without a geocodable address are excluded. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. The estimation sample is the balanced panel reconstructed excluding the 2012 and 2021 fiscal years. Panel A of Table 7 reports the same specification on the balanced panel.

Table B.5. Physical (GIS-based) exposure: average DD Estimates (log): *balanced panel excluding 2012 and 2021*

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t$	0.163	0.229	0.358**	0.302*
	(0.129)	(0.186)	(0.166)	(0.182)
Observations	4384	4384	4384	4384
Firms	274	274	274	274
Treated firms	10	10	10	10
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Treatment is physical exposure at the firm's geocoded pre-disaster address (registry address pinned to 2010, with the earliest available registry address as a fallback): a firm is treated if the address lies within 500 m of the tsunami-inundation footprint mapped by the Geospatial Information Authority of Japan or in the zone of instrumental seismic intensity of at least 5.5 interpolated from Japan Meteorological Agency station records. Self-reported exposure plays no role: survey-treated firms that are not physically exposed serve as controls, and firms without a geocodable address are excluded. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Outcomes are $\log(1 + Y)$, so coefficients can be approximately interpreted as proportional changes. The estimation sample is the balanced panel reconstructed excluding the 2012 and 2021 fiscal years. Table A.13 reports the same specification on the balanced panel.

Table B.6. Disaster-area (municipality-based) treatment definition: average DD Estimates (level): *balanced panel excluding 2012 and 2021*

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
$Treat_i \times Post_t$	−20.522	−4.131	24.263	43.872
	(355.341)	(4.804)	(206.851)	(99.946)
Observations	4368	4368	4368	4368
Firms	273	273	273	273
Treated firms	13	13	13	13
Pre-treatment mean	2,056.09	49.53	1,640.27	697.69
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Treatment is assignment to the officially designated disaster area of Carvalho et al. (2021): a firm is treated if its geocoded pre-disaster address (registry address pinned to 2010, with the earliest available registry address as a fallback) lies in one of the 41 disaster-area municipalities designated by the MLIT decree of 28 April 2011 and the government evacuation decrees of 21–22 April 2011. Self-reported exposure plays no role: survey-treated firms outside the disaster area serve as controls, and firms without a geocodable address are excluded. Treatment is assigned at the municipality level; standard errors are clustered at the firm level as in all other tables. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. The estimation sample is the balanced panel reconstructed excluding the 2012 and 2021 fiscal years. Panel B of Table 7 reports the same specification on the balanced panel.

Table B.7. Alternative treatment definitions: treated and total firm counts, *balanced panel excluding 2012 and 2021*

Treatment definition	Treated	Firms
Baseline: self-reported exposure	25	274
Physical (GIS-based) exposure	10	274
Designated disaster area (41 municipalities)	13	273

Notes: As Table 5, on the balanced panel reconstructed excluding the 2012 and 2021 fiscal years. One firm lacks a municipality assignment and so drops from the disaster-area definition. Definitions are those of Table 5.

Table B.8. Alternative treatment definitions: cross-tabulation of the treated groups, *balanced panel excluding 2012 and 2021*

	Self-reported	Physical	Disaster area
Self-reported exposure	25	5	9
Physical (GIS-based) exposure	5	10	9
Designated disaster area	9	9	13

Notes: As Table 6, on the balanced panel reconstructed excluding the 2012 and 2021 fiscal years. A firm enters the cross-tabulation only if all three definitions are defined for it, which drops the one firm without a municipality assignment and leaves 273 of the panel's 274 firms.

Figure B.1: Trends in mean outcomes for treatment and control groups: *Balanced panel excluding 2012 and 2021*

Notes: Yearly means by treatment group, balanced panel reconstructed excluding the 2012 and 2021 fiscal years (274 firms, of which 25 treated); the two excluded years are absent from the series. The dashed vertical line marks 2011, the year of the GEJE. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure B.2: Dynamic Effects on Firm Outcomes: *Balanced panel excluding 2012 and 2021*

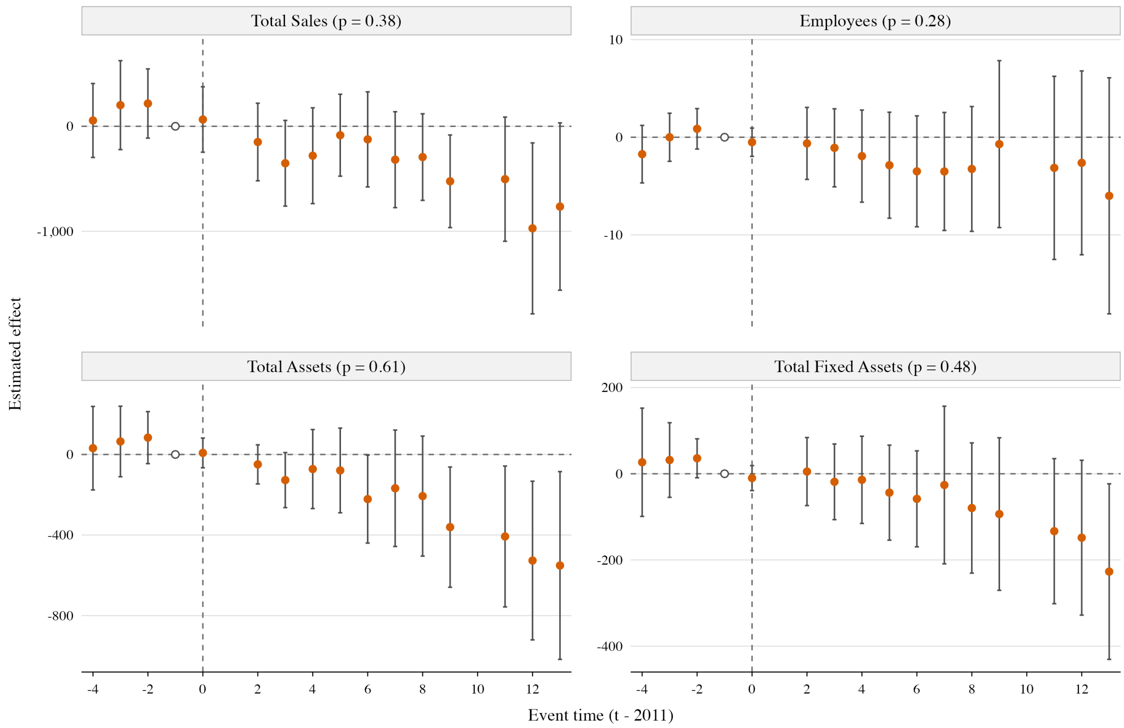

Notes: Event-study estimates from Equation (2), level outcomes, balanced panel reconstructed excluding the 2012 and 2021 fiscal years (274 firms, of which 25 treated); event years $k = 1$ and $k = 10$ are therefore absent. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

B.2. Fiscal-Year Timing: Closing-Month Subsamples

Statements are assigned to the fiscal year in which the accounting period ends, so for a firm closing its accounts in January–March, fiscal 2011 ends essentially before the disaster (Section 2.2). The absence of an immediate effect in the event studies of Section 3 could therefore in principle reflect this timing rather than a genuine absence of an early effect. This appendix splits the estimation panels by the closing month of the fiscal-2011 accounting period and re-estimates the design on each group. Figure B.3 shows the distribution of closing months: March closings are modal (33 percent of the balanced panel and 29 percent of the panel excluding 2012 and 2021), and the January–March group comprises 41 percent of the balanced panel (42 of 103 firms) and 37 percent of the panel excluding 2012 and 2021 (101 of 274 firms).

The direct test is the April–December group, the firms for which fiscal 2011 does contain between roughly one and ten post-disaster months. Its event study (Figure B.4) shows no immediate effect: the coefficient at $k = 0$ is statistically insignificant for every outcome, so the lagged onset documented in Section 3 is not an artifact of partially treated fiscal years. Two caveats apply. The restriction removes about four in ten firms, including every March closer, and leaves only 13 treated firms on the balanced panel, so the subsample estimates are imprecise. And the average effects attenuate relative to the full panel (Table B.9, Panel A against Table 2): the full-panel losses draw importantly on the January–March closers, whose estimates Panel B reports directly.

Among the January–March closers the losses are large despite the small sample (42 firms, of which 6 treated, on the balanced panel), and their event study (Figure B.5) traces the same delayed and deepening trajectory as the headline figures. The two closing-month groups are not exchangeable, however: January–March closers are markedly larger (pre-treatment mean total sales of 2,268.9 million yen against 1,218.1 among April–December closers in the balanced panel), so the concentration of the yen losses among them lines up with the size gradient of Section 3.3.2, larger treated firms bear larger absolute losses, rather than pointing to fiscal-year timing itself as the driver. Figure B.6 shows both associations descriptively: the industry mix of the two closing-month groups differs in the balanced panel (Fisher exact $p = 0.013$), and the January–March closers sit visibly to the right of the April–December closers in the pre-2011 size distribution, significantly so in the panel excluding 2012 and 2021 (Wilcoxon $p = 0.003$ for pre-2011 mean sales, $p = 0.001$ for employment).

Figure B.3: Distribution of Fiscal-2011 Closing Months

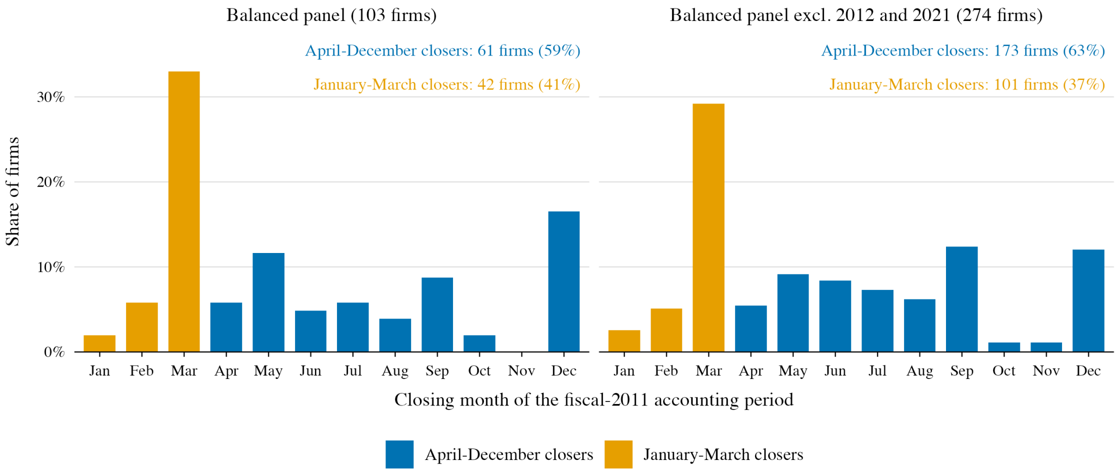

Notes: Share of firms by the closing month of their fiscal-2011 accounting period, within each estimation panel. Colors mark the two closing-month groups used in Table B.9 and Figures B.4 and B.5: April–December closers, whose fiscal 2011 contains post-disaster months, and January–March closers, whose fiscal 2011 ends essentially before the disaster. Annotations report each group's firm count and share of the panel.

Figure B.4: Dynamic Effects among April–December Closers: *Balanced panel*

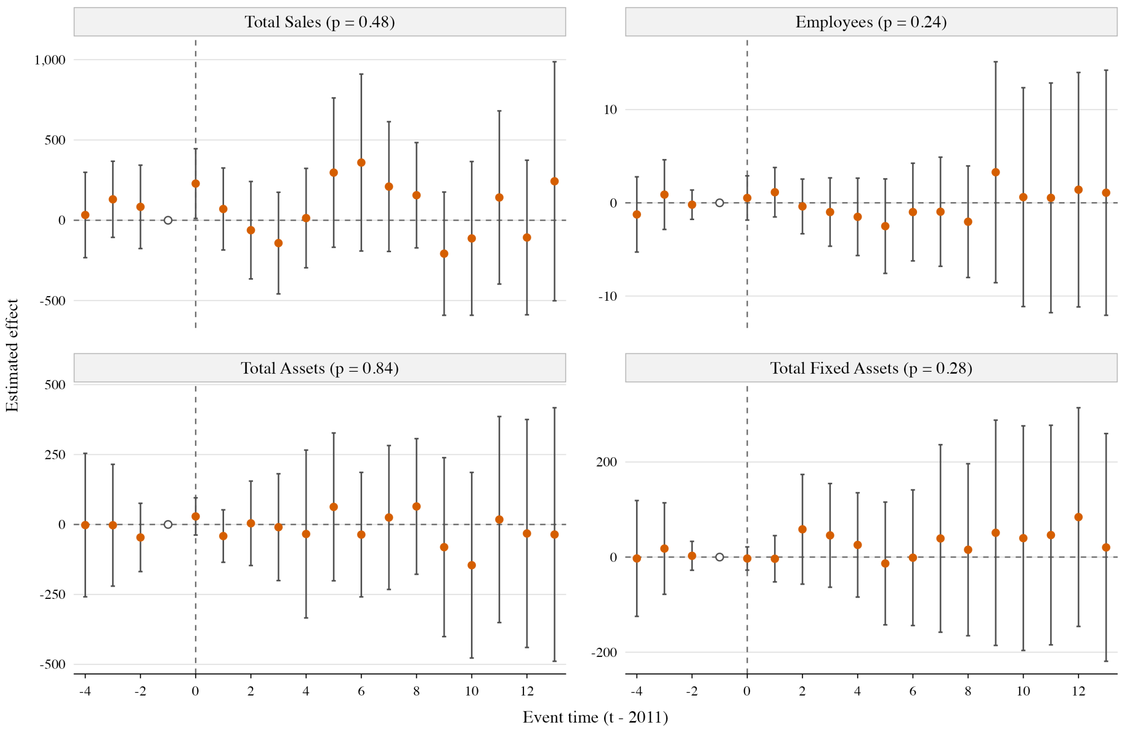

Notes: Event-study estimates from Equation (2), level outcomes, estimated on the firms of the balanced panel whose fiscal-2011 accounting period closes in April–December, so that fiscal 2011 contains post-disaster months. The coefficient at $k = 0$ is therefore a direct test of an immediate effect. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Table B.9. Fiscal-2011 timing: average DD by closing-month group (level)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel, April–December closers				
$Treat_i \times Post_t$	15.898	0.086	−2.428	24.478
	(136.367)	(3.149)	(121.479)	(71.557)
Observations	1098	1098	1098	1098
Firms	61	61	61	61
Pre-treatment mean	1,218.06	30.33	894.78	309.04
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel B: Balanced panel, January–March closers				
$Treat_i \times Post_t$	−1549.254**	−16.749**	−984.499**	−511.582**
	(629.203)	(6.656)	(373.225)	(248.043)
Observations	756	756	756	756
Firms	42	42	42	42
Pre-treatment mean	2,268.94	56.73	1,625.51	520.71
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel C: Balanced panel excluding 2012 and 2021, April–December closers				
$Treat_i \times Post_t$	−14.817	−0.835	−129.914	−33.475
	(108.794)	(2.611)	(116.934)	(66.345)
Observations	2768	2768	2768	2768
Firms	173	173	173	173
Pre-treatment mean	1,525.38	39.46	1,268.69	505.80
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel D: Balanced panel excluding 2012 and 2021, January–March closers				
$Treat_i \times Post_t$	−1562.054**	−4.013	−704.467**	−294.274**
	(640.738)	(7.597)	(273.453)	(140.289)
Observations	1616	1616	1616	1616
Firms	101	101	101	101
Pre-treatment mean	2,975.22	67.28	2,304.54	1,038.41
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Firms are split by the closing month of their fiscal-2011 accounting period. Statements are assigned to the fiscal year in which the accounting period ends, so fiscal 2011 contains post-disaster months for an April–December closer but ends essentially before the disaster for a January–March closer. Of the 103 balanced-panel firms (19 treated), 61 (13 treated) close in April–December and 42 (6 treated) in January–March; of the 274 firms on the panel excluding 2012 and 2021 (25 treated), 173 (16 treated) and 101 (9 treated) respectively. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. Figures B.4 and B.5 report the corresponding event studies on the balanced panel.

Figure B.5: Dynamic Effects among January–March Closers: *Balanced panel*

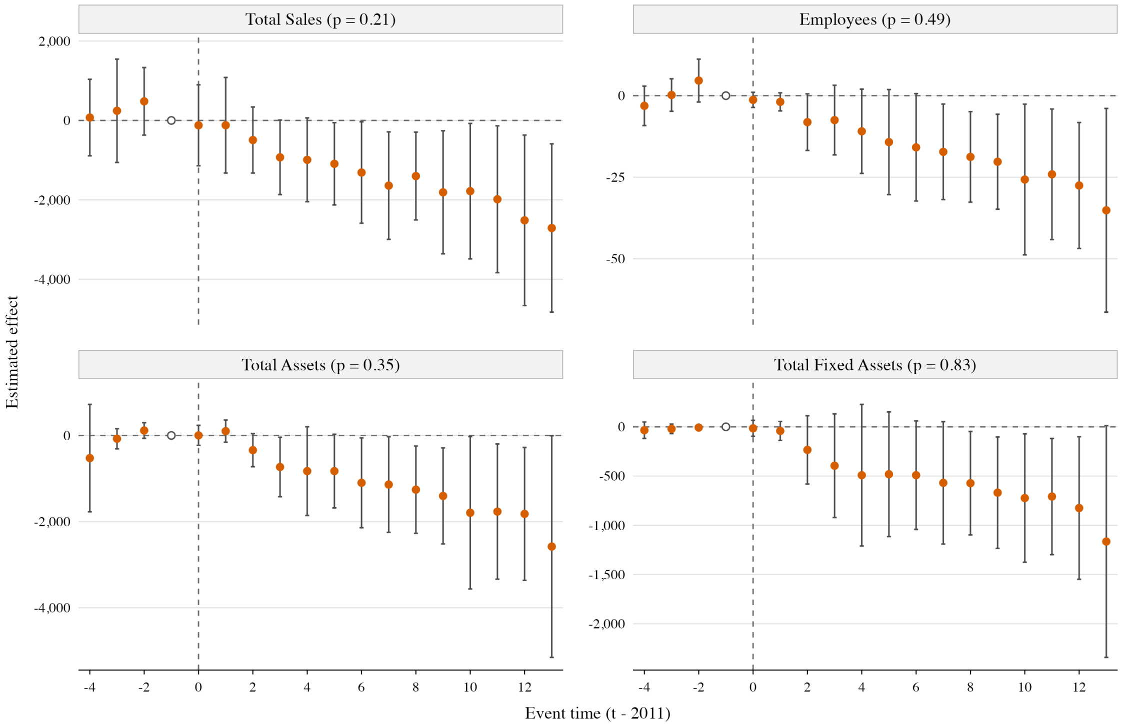

Notes: Event-study estimates from Equation (2), level outcomes, estimated on the firms of the balanced panel whose fiscal-2011 accounting period closes in January–March, so that fiscal 2011 ends essentially before the disaster and the coefficient at $k = 0$ is close to a placebo. The reference year is 2010 ($k = -1$, open circle); bars are 95% confidence intervals from standard errors clustered at the firm level. Firm and construction (2010) $\times$ year fixed effects are included. The pre-trend p-value in each panel title is from a joint Wald test that all pre-2011 coefficients equal zero. Monetary variables are measured in 1 million yen; the number of employees is in persons.

Figure B.6: Closing-Month Groups: Industry Composition and Pre-Disaster Firm Size

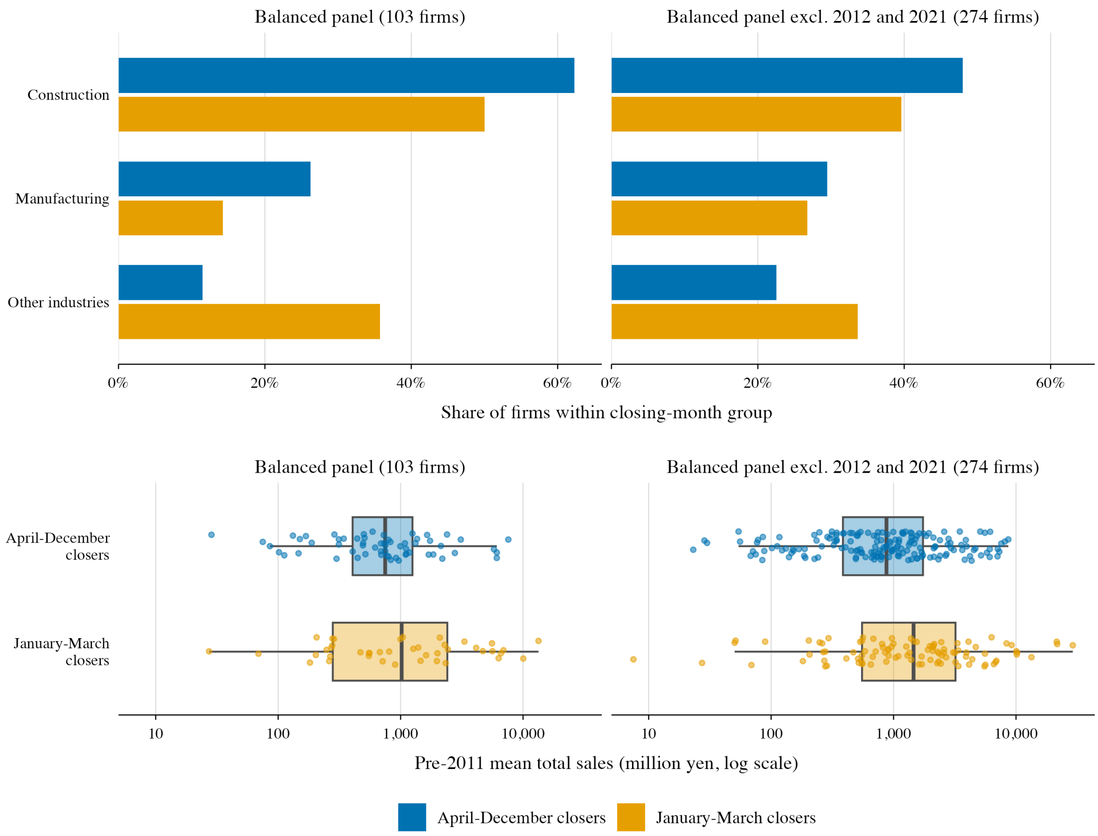

Notes: Firms grouped by the closing month of their fiscal-2011 accounting period, within each estimation panel. Top: industry composition of each closing-month group (2010-pinned industry division, grouped into construction, manufacturing, and other). Bottom: distribution of each firm's 2007–2010 mean total sales (million yen, log scale); boxes span the interquartile range with the median line, and dots are individual firms. Fisher exact tests for equal industry composition across closing-month groups: $p = 0.013$ (balanced panel), $p = 0.131$ (excl. 2012 and 2021). Wilcoxon tests for equal pre-2011 size distributions: mean total sales $p = 0.166$ (balanced), $p = 0.003$ (excl. 2012 and 2021); mean employment $p = 0.187$ (balanced), $p = 0.001$ (excl. 2012 and 2021).

B.3. Firm Size: Additional Results

The event studies (Appendix Figures B.7 and B.8) show that the large-firm losses follow the same lagged trajectory as the pooled results, emerging several years after the disaster and deepening thereafter, while the small-firm coefficients remain near zero throughout. In logs (Appendix Table B.10 and Figures B.9 and B.10), the only significant large-firm loss is in fixed assets (-0.28, 5 percent level), against a positive small-firm coefficient (0.53, 10 percent level); the two differ at $p = 0.01$. For the other three outcomes the log coefficients are small and imprecisely estimated. The level results therefore likely reflect the economic scale of large treated firms rather than a uniformly larger proportional decline.[15]

Firm size is defined as in Section 3.3.2: small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster accounting data. The tables use the main-specification fixed effects and add reported interaction controls for differential trends of large firms — $Post_t \times Large_i$ in the average tables, window $\times$ $Large_i$ in the period tables; the event-study figures estimate each group on its own subsample with the main-specification fixed effects. Because the estimation panels select toward larger firms, this respondent-wide cutoff produces uneven within-panel splits (26 small versus 77 large firms on the balanced panel; 55 versus 219 excluding 2012 and 2021); the final table and event-study figures therefore re-estimate the split with the median recomputed within each estimation panel, which makes the split even by construction (51/52 and 137/137).

[15] The concentration among large firms is not an artifact of the uneven split that the respondent-wide median induces (26 versus 77 firms). Recomputing the median within each estimation panel, which makes the split even by construction (51 versus 52), yields the same pattern, if anything sharper: the large-firm sales loss is 873.7 million yen and the small-firm coefficients remain indistinguishable from zero (Appendix Table B.11). The between-group test is also more informative under the even split: on the panel excluding 2012 and 2021 it rejects equality for total sales ($p = 0.02$) and is marginal for total assets ($p = 0.07$), though on the balanced panel it still does not reject for any outcome.

Figure B.7: Event-study estimates by firm size: *Balanced panel, level outcomes*

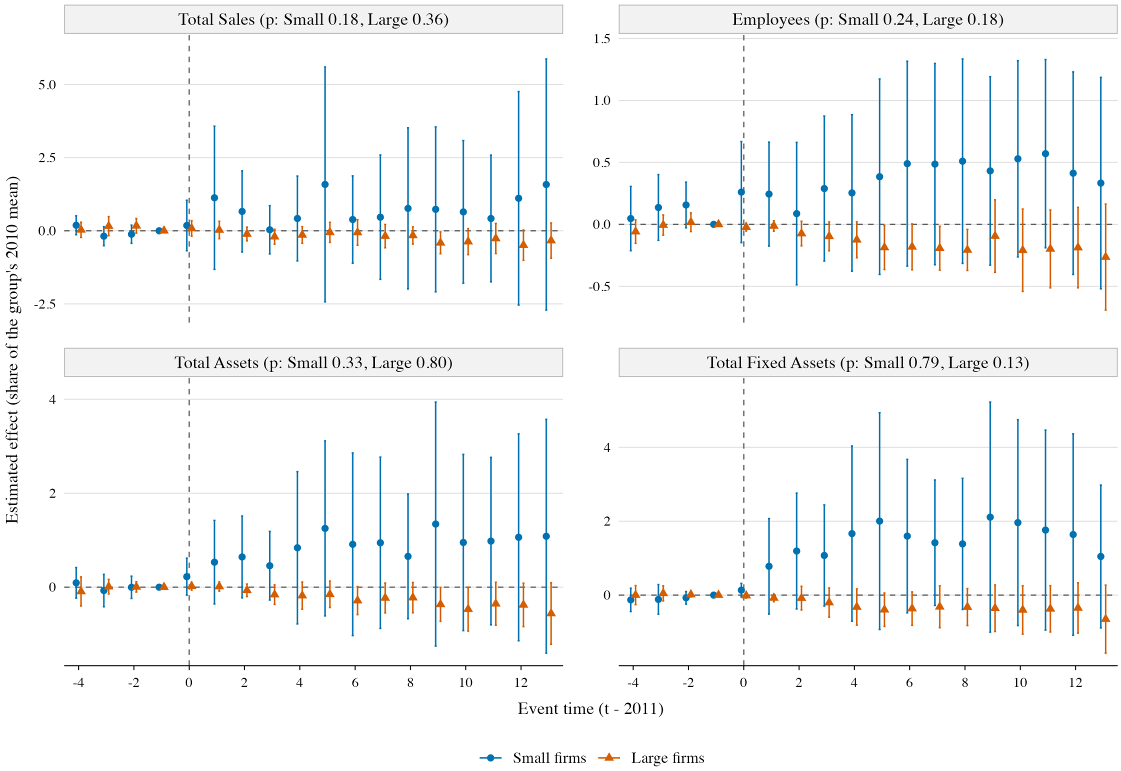

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the small-firm (26 firms) and large-firm (77 firms) subsamples of the balanced panel with the main-specification fixed effects. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Figure B.8: Event-study estimates by firm size: *Balanced panel excluding 2012 and 2021, level outcomes*

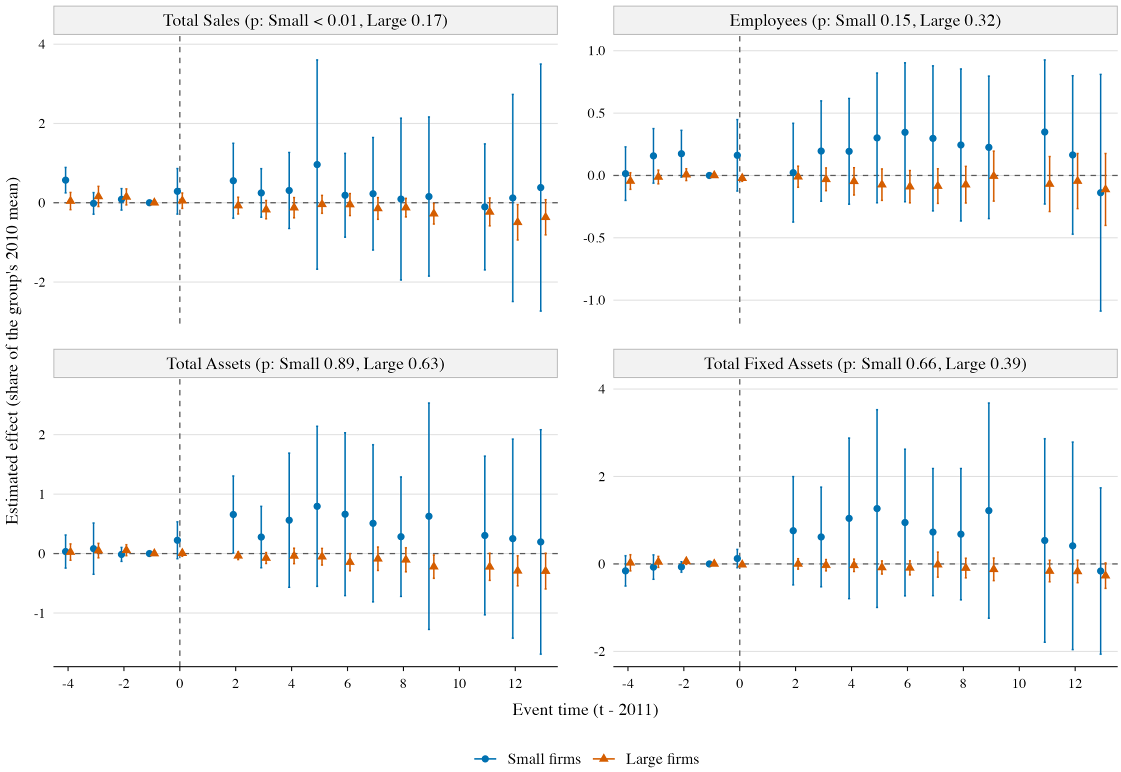

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the small-firm (55 firms) and large-firm (219 firms) subsamples of the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, with the main-specification fixed effects; event years $k = 1$ and $k = 10$ are therefore absent. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Table B.10. Heterogeneous DD Estimates: small vs. large firms (log)

Outcome:	(1) Total Sales	(2) Number of Employees	(3) Total Assets	(4) Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t \times Small_i$	0.005	0.280	0.344	0.528*
	(0.223)	(0.291)	(0.289)	(0.293)
$Treat_i \times Post_t \times Large_i$	−0.139	−0.089	−0.125	−0.284**
	(0.095)	(0.088)	(0.095)	(0.115)
$Post_t \times Large_i$	0.071	0.124*	0.220	0.362**
	(0.096)	(0.070)	(0.140)	(0.158)
p-value: $\beta_L = \beta_S$	0.56	0.23	0.13	0.01
Observations	1854	1854	1854	1854
Firms (small)	26	26	26	26
Firms (large)	77	77	77	77
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t \times Small_i$	−0.059	0.230	0.178	0.333
	(0.180)	(0.238)	(0.239)	(0.256)
$Treat_i \times Post_t \times Large_i$	−0.050	−0.037	−0.077	−0.173**
	(0.068)	(0.065)	(0.071)	(0.085)
$Post_t \times Large_i$	−0.017	0.066	0.024	0.103
	(0.066)	(0.047)	(0.099)	(0.130)
p-value: $\beta_L = \beta_S$	0.97	0.28	0.31	0.06
Observations	4384	4384	4384	4384
Firms (small)	55	55	55	55
Firms (large)	219	219	219	219
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. Outcomes are $\log(1 + Y)$, so coefficients can be approximately interpreted as proportional changes. Firm and construction (2010) × year fixed effects are included; the $Post_t \times Large_i$ interaction controls for differential trends of large firms, and its coefficient is reported. Standard errors clustered at the firm level are reported in parentheses. The p-value row tests $\beta_L = \beta_S$ (Equation (4)), the equality of the large- and small-firm treatment coefficients, from the same specification re-parameterised with small firms as the baseline, so that the large-firm term is the between-group difference. In Panel B, the 2012 and 2021 fiscal years are excluded.

Figure B.9: Event-study estimates by firm size: *Balanced panel, log outcomes*

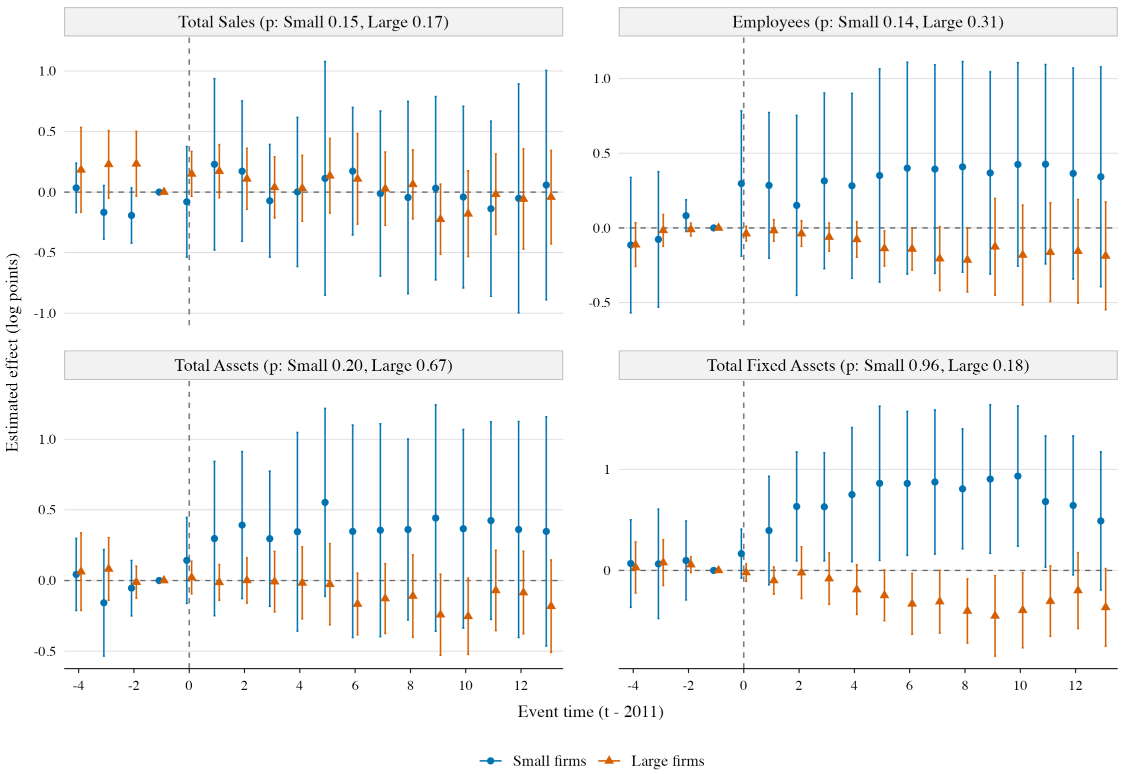

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, estimated separately on the small-firm (26 firms) and large-firm (77 firms) subsamples of the balanced panel with the main-specification fixed effects. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero. Since outcomes are in logs, coefficients can be approximately interpreted as proportional changes relative to the control group.

Figure B.10: Event-study estimates by firm size: *Balanced panel excluding 2012 and 2021, log outcomes*

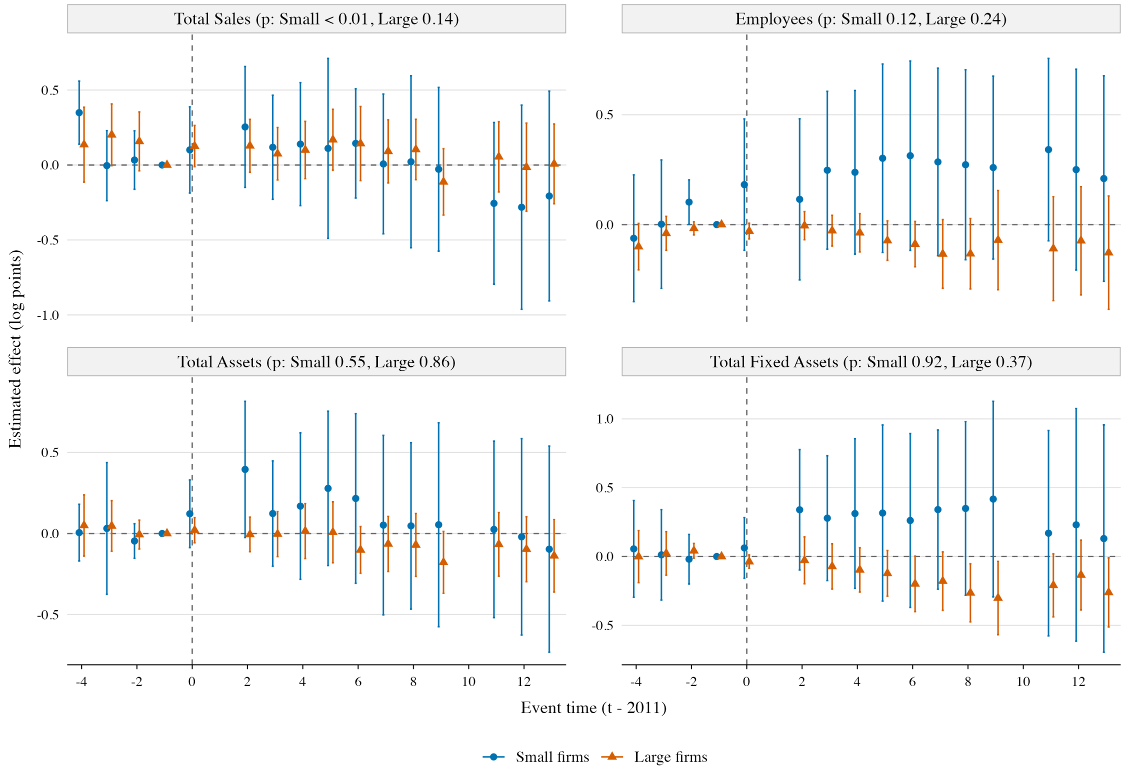

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, estimated separately on the small-firm (55 firms) and large-firm (219 firms) subsamples of the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, with the main-specification fixed effects; event years $k = 1$ and $k = 10$ are therefore absent. Small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median across all respondent firms with pre-disaster data. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero. Since outcomes are in logs, coefficients can be approximately interpreted as proportional changes relative to the control group.

Table B.11. Heterogeneous DD Estimates: small vs. large firms (level, within-panel median)

	(1)	(2)	(3)	(4)
Outcome:	Total Sales	Number of Employees	Total Assets	Total Fixed Assets
Panel A: Balanced panel				
$Treat_i \times Post_t \times Small_i$	−56.710	−0.260	−55.737	−63.607
	(125.817)	(2.604)	(96.789)	(64.397)
$Treat_i \times Post_t \times Large_i$	−873.678*	−9.767	−393.550	−139.049
	(508.684)	(7.418)	(279.951)	(147.662)
$Post_t \times Large_i$	441.238***	5.098***	479.448***	229.313***
	(162.579)	(1.773)	(122.947)	(84.544)
p-value: $\beta_L = \beta_S$	0.12	0.24	0.23	0.60
Observations	1854	1854	1854	1854
Firms (small)	51	51	51	51
Firms (large)	52	52	52	52
Pre-treatment mean (small)	385.56	17.51	372.23	195.43
Pre-treatment mean (large)	2,883.35	64.22	1,997.48	591.43
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021				
$Treat_i \times Post_t \times Small_i$	−3.715	−1.228	−31.851	−32.425
	(88.603)	(2.013)	(64.545)	(39.440)
$Treat_i \times Post_t \times Large_i$	−1009.321**	−2.238	−477.566**	−127.293
	(432.796)	(7.020)	(242.553)	(118.929)
$Post_t \times Large_i$	627.374***	5.545**	573.850***	210.388***
	(221.522)	(2.432)	(111.060)	(57.492)
p-value: $\beta_L = \beta_S$	0.02	0.89	0.07	0.44
Observations	4384	4384	4384	4384
Firms (small)	137	137	137	137
Firms (large)	137	137	137	137
Pre-treatment mean (small)	463.90	20.07	423.03	206.56
Pre-treatment mean (large)	3,655.72	79.35	2,878.01	1,197.69
Firm FE	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓

Notes: Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Unlike the headline size classification (median across all respondent firms with pre-disaster data), here small (large) firms are those whose 2007–2010 mean Total Sales lies below (at or above) the median computed within each estimation panel separately, so the split is roughly even by construction and a firm's size class may differ between panels. Firm and construction (2010) × year fixed effects are included; the $Post_t \times Large_i$ interaction controls for differential trends of large firms, and its coefficient is reported. Standard errors clustered at the firm level are reported in parentheses. Monetary variables are measured in 1 million yen. The pre-treatment means pool treated and control firm-years 2007–2010 within each group. In Panel B, the 2012 and 2021 fiscal years are excluded.

Figure B.11: Event-study estimates by firm size, within-panel median: *Balanced panel, level outcomes*

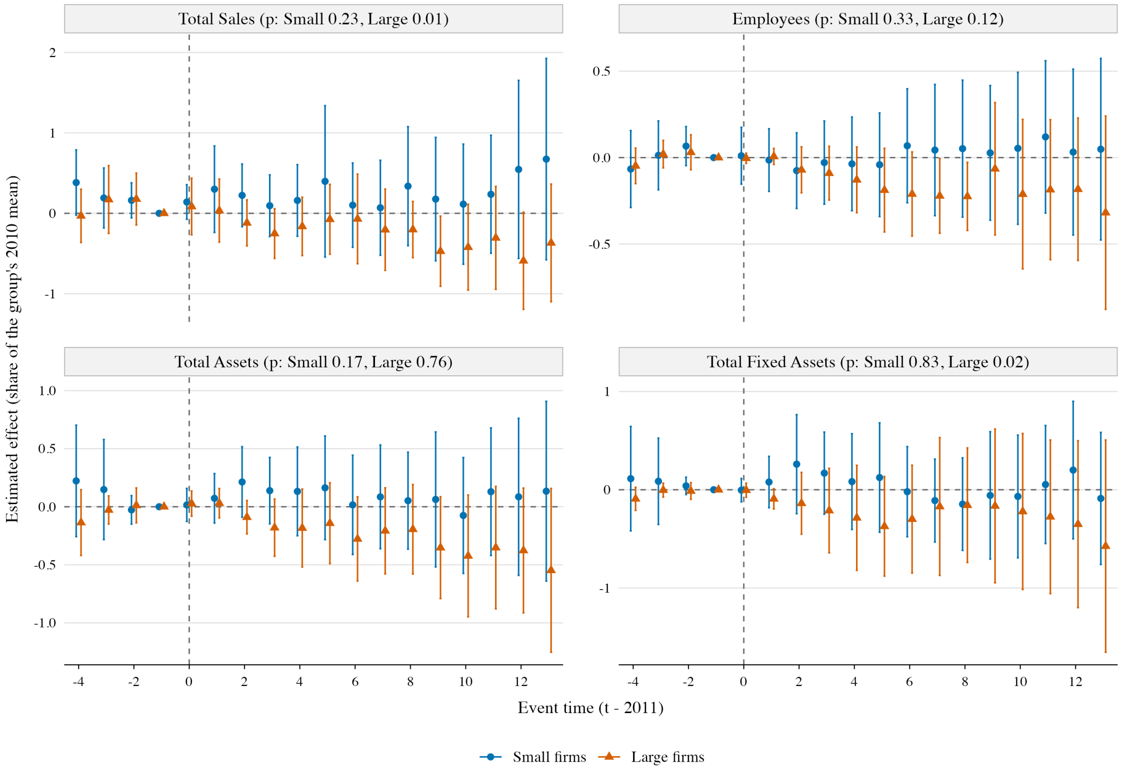

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the small-firm (51 firms) and large-firm (52 firms) subsamples of the balanced panel with the main-specification fixed effects. Unlike the headline classification, small (large) firms are here those whose 2007–2010 mean Total Sales lies below (at or above) the median recomputed within the estimation panel, so the split is even by construction. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Figure B.12: Event-study estimates by firm size, within-panel median: *Balanced panel excluding 2012 and 2021, level outcomes*

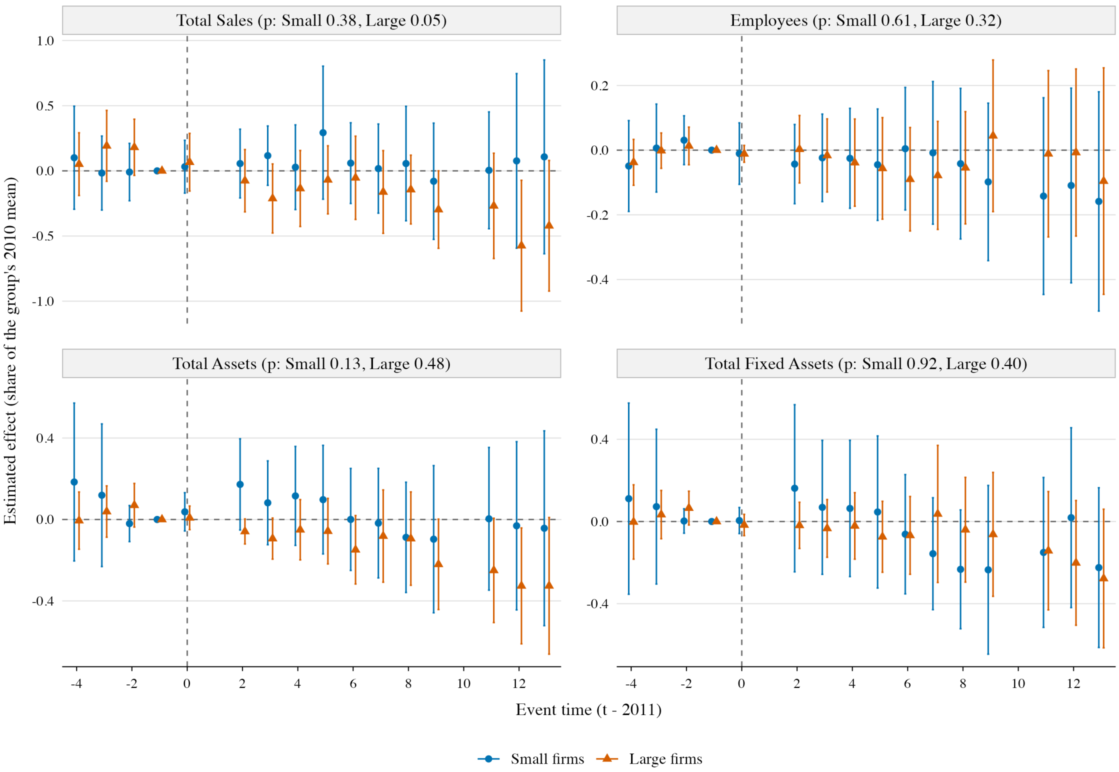

Notes: Event-study estimates from Equation (2), level outcomes, estimated separately on the small-firm (137 firms) and large-firm (137 firms) subsamples of the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, with the main-specification fixed effects; event years $k = 1$ and $k = 10$ are therefore absent. Unlike the headline classification, small (large) firms are here those whose 2007–2010 mean Total Sales lies below (at or above) the median recomputed within the estimation panel, so the split is even by construction. Each group's coefficients are divided by the group's pooled (treated and control) 2010 mean of the outcome, so the vertical axis reads as a share of the group's 2010 level. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.

Figure B.13: Event-study estimates by firm size, within-panel median: *Balanced panel, log outcomes*

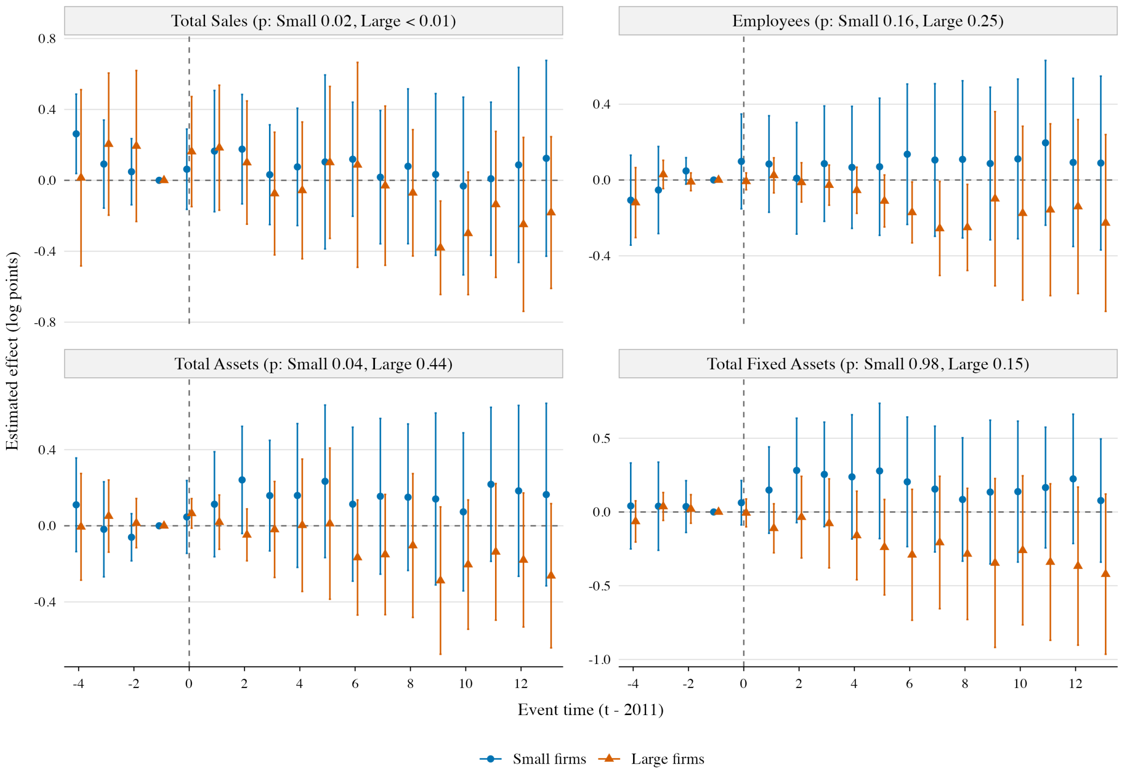

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, estimated separately on the small-firm (51 firms) and large-firm (52 firms) subsamples of the balanced panel with the main-specification fixed effects. Unlike the headline classification, small (large) firms are here those whose 2007–2010 mean Total Sales lies below (at or above) the median recomputed within the estimation panel, so the split is even by construction. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero. Since outcomes are in logs, coefficients can be approximately interpreted as proportional changes relative to the control group.

Figure B.14: Event-study estimates by firm size, within-panel median: *Balanced panel excluding 2012 and 2021, log outcomes*

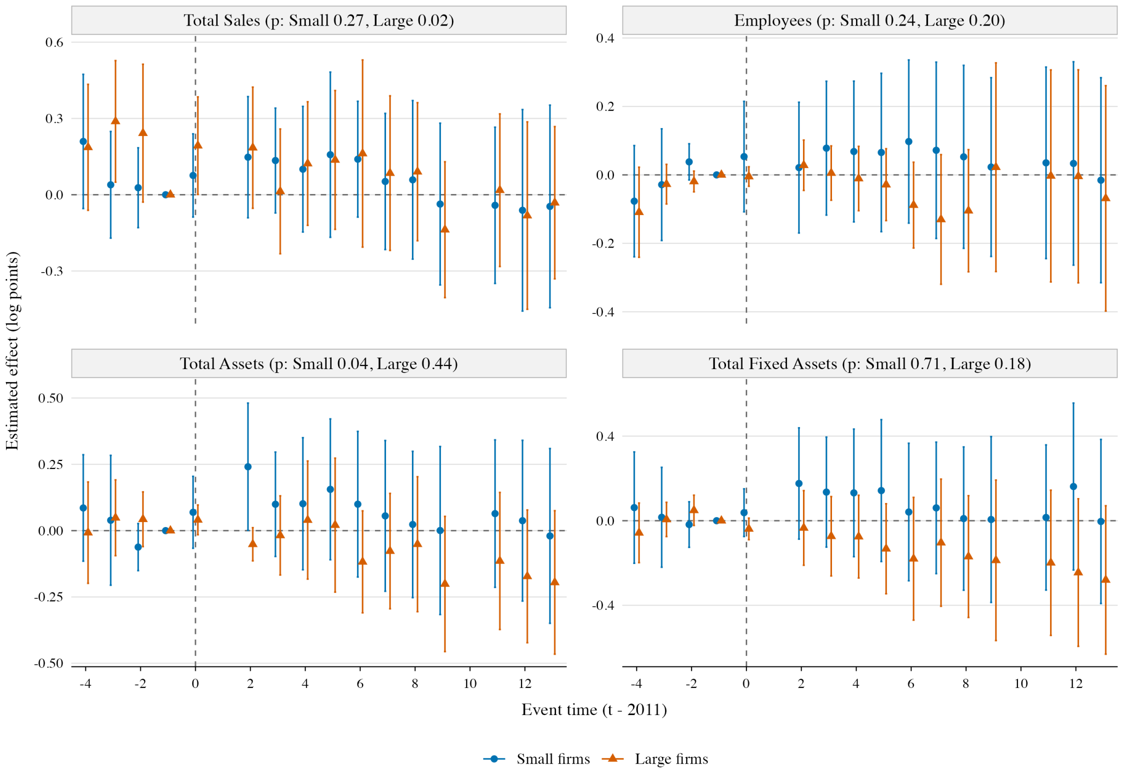

Notes: Event-study estimates from Equation (2), $\log(1+Y)$ outcomes, estimated separately on the small-firm (137 firms) and large-firm (137 firms) subsamples of the balanced panel reconstructed excluding the 2012 and 2021 fiscal years, with the main-specification fixed effects; event years $k = 1$ and $k = 10$ are therefore absent. Unlike the headline classification, small (large) firms are here those whose 2007–2010 mean Total Sales lies below (at or above) the median recomputed within the estimation panel, so the split is even by construction. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero. Since outcomes are in logs, coefficients can be approximately interpreted as proportional changes relative to the control group.

B.4. Additional Outcomes: Profits and Balance-Sheet Positions

The headline analysis studies four outcomes: total sales, employees, total assets, and total fixed assets. This appendix re-estimates the main specification, on the same samples and with the same fixed effects and firm-clustered standard errors, for five further outcomes drawn from the same registry financial statements, all in untransformed levels (million yen): operating profit, gross profit, total liabilities, net worth, and long-term debt. The two profit measures ask whether the persistent contraction in scale was accompanied by a deterioration in earnings, and are read with the reporting caveat of Section 2.2 in mind; the three balance-sheet positions trace how exposed firms' financing absorbed the shock. Table B.12 reports the average DD estimates and Figure B.15 the corresponding event studies.

Because the headline heterogeneity analysis shows that reconstruction demand shields construction firms, Table B.13 and Figure B.16 re-estimate the same five outcomes separately for construction and non-construction firms, in the format of the main construction split (Section 3.3.1): both group coefficients are reported, and the p-value row tests their equality from the same specification re-parameterised with construction firms as the baseline.

Table B.12. Average DD Estimates: additional outcomes in levels

Outcome:	(1) Operating Profit	(2) Gross Profit	(3) Total Liabilities	(4) Net Worth	(5) Long-Term Debt
Panel A: Balanced panel					
$Treat_i \times Post_t$	−42.888**	−92.633*	−233.783	−41.987	−134.945*
	(19.066)	(48.479)	(160.361)	(145.933)	(73.941)
Observations	1854	1854	1854	1854	1854
Firms	103	103	103	103	103
Pre-treatment mean	21.468	248.670	706.675	486.070	160.809
Firm FE	✓	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021					
$Treat_i \times Post_t$	−26.383*	−78.616*	−272.093*	−2.972	−187.384**
	(15.205)	(41.380)	(143.102)	(134.727)	(75.049)
Observations	4384	4384	4384	4384	4384
Firms	274	274	274	274	274
Pre-treatment mean	40.055	301.829	1047.468	603.049	308.563
Firm FE	✓	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓	✓

Notes: All outcomes in million yen, without transformation or winsorization. Long-term debt is long-term borrowings plus corporate bonds. Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. In Panel B, the 2012 and 2021 fiscal years are excluded.

Figure B.15: Event-study estimates: *additional outcomes in levels, balanced panel*

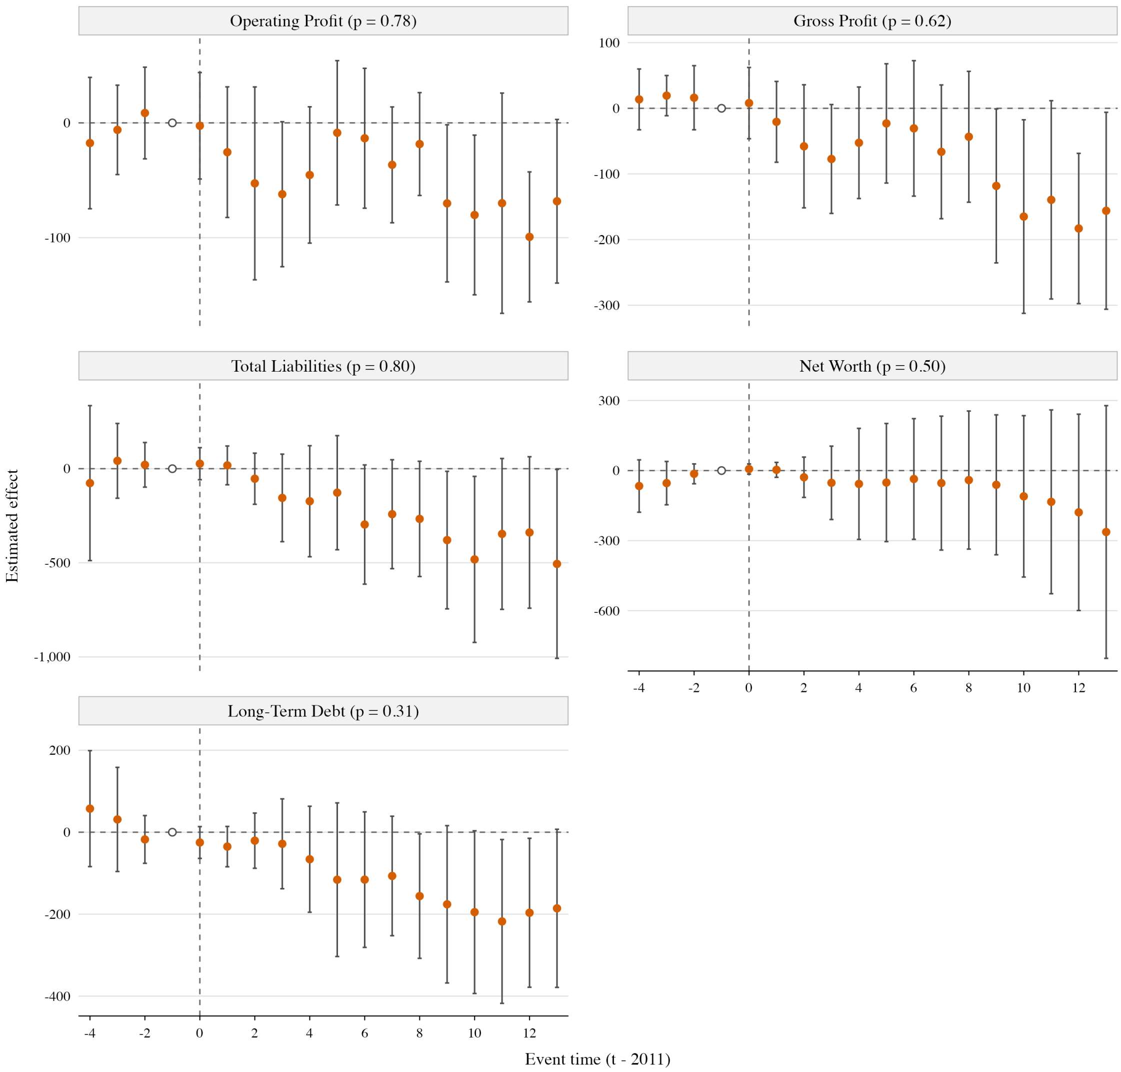

Notes: Event-study estimates from Equation (2) on the balanced panel with the main-specification fixed effects. All outcomes in million yen, without transformation or winsorization; long-term debt is long-term borrowings plus corporate bonds. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level. The pre-trend p-values in each panel title are from joint Wald tests that all pre-2011 coefficients equal zero.

Table B.13. Average DD Estimates: additional outcomes in levels, construction vs. non-construction

	(1)	(2)	(3)	(4)	(5)
Outcome:	Operating Profit	Gross Profit	Total Liabilities	Net Worth	Long-Term Debt
Panel A: Balanced panel					
$Treat_i \times Post_t \times Construction_i$	2.664	34.968	20.291	103.865	16.278
	(16.783)	(50.892)	(76.128)	(98.607)	(43.967)
$Treat_i \times Post_t \times NonConstruction_i$	−87.850***	−218.581***	−484.563	−185.948	−284.208**
	(29.444)	(66.697)	(294.574)	(267.651)	(128.324)
p-value: $\beta_C = \beta_N$	< 0.01	< 0.01	0.10	0.31	0.03
Observations	1854	1854	1854	1854	1854
Firms (construction)	59	59	59	59	59
Firms (non-construction)	44	44	44	44	44
Pre-treatment mean (construction)	10.21	102.53	319.26	372.73	60.37
Pre-treatment mean (non-construction)	36.56	444.63	1,226.17	638.05	295.48
Firm FE	✓	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓	✓
Panel B: Balanced panel excluding 2012 and 2021					
$Treat_i \times Post_t \times Construction_i$	−9.810	9.491	54.128	−1.201	28.956
	(17.763)	(44.814)	(82.883)	(93.433)	(38.491)
$Treat_i \times Post_t \times NonConstruction_i$	−39.453*	−148.094**	−529.338**	−4.368	−357.980***
	(22.811)	(59.786)	(228.613)	(229.457)	(114.663)
p-value: $\beta_C = \beta_N$	0.31	0.04	0.02	0.99	< 0.01
Observations	4384	4384	4384	4384	4384
Firms (construction)	123	123	123	123	123
Firms (non-construction)	151	151	151	151	151
Pre-treatment mean (construction)	23.40	141.34	565.68	393.11	125.20
Pre-treatment mean (non-construction)	53.62	432.56	1,439.92	774.06	457.93
Firm FE	✓	✓	✓	✓	✓
Construction × Year FE	✓	✓	✓	✓	✓

Notes: All outcomes in million yen, without transformation or winsorization. Long-term debt is long-term borrowings plus corporate bonds. The construction × year cells absorb the $Post_t \times Construction_i$ main effect. The p-value row tests $\beta_C = \beta_N$, the equality of the two group coefficients, from the same specification re-parameterised with construction firms as the baseline, so that the non-construction term is the between-group difference. Survey panel; treatment is self-reported exposure to the March 2011 earthquake or tsunami at either establishment. Firm and construction (industry classification as of 2010) × year fixed effects are included in all specifications. Standard errors clustered at the firm level are reported in parentheses. The pre-treatment mean pools treated and control firm-years 2007–2010 in the column's estimation sample. In Panel B, the 2012 and 2021 fiscal years are excluded.

Figure B.16: Event-study estimates by construction status: *additional outcomes in levels, balanced panel*

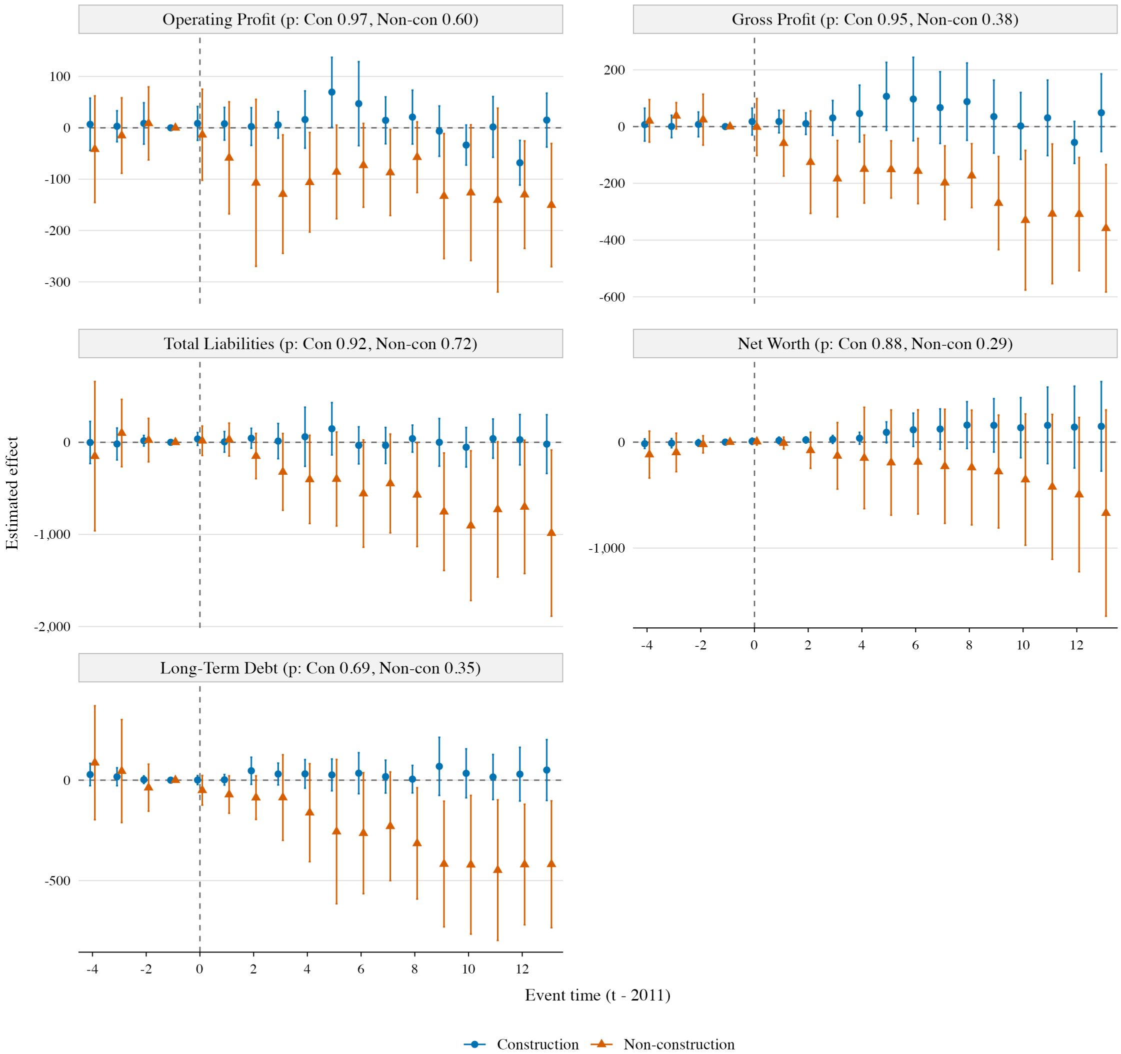

Notes: Event-study estimates from Equation (2), estimated separately on the construction and non-construction subsamples of the balanced panel with the main-specification fixed effects; effects are unscaled. Outcome definitions as in Figure B.15. The reference year is 2010 ($k = -1$); bars are 95 percent confidence intervals from standard errors clustered at the firm level; points are horizontally dodged for readability. The pre-trend p-values in each panel title are from joint Wald tests, one per group, that all of the group's pre-2011 coefficients equal zero.