# Project Networks and Reallocation Externalities

Vibhuti Dhingra
Schulich School of Business, York University, Toronto, ON M3J 1P3, Canada vibhutid@schulich.yorku.ca

Harish Krishnan
University of British Columbia, Vancouver, BC V6T 1Z4, Canada Harish.Krishnan@sauder.ubc.ca

Juan Camilo Serpa
Desautels Faculty of Management, McGill University, Montreal, QC H3A 1G5, Canada juan.serpa@mcgill.ca

A project involves several participants—including clients, contractors, and subcontractors—that work concurrently on multiple projects and allocate resources among them. This interdependency creates a network of otherwise unrelated projects. We map the network of U.S. government projects involving over 150,000 participants. We show that a seemingly localized disruption, affecting only one project site, eventually causes delays across unrelated projects. This is because participants opportunistically reallocate resources into disrupted projects, at the expense of other projects, triggering a domino effect of further reallocations in the network. Thus, the costs of on-site disruptions end up being shared by multiple participants in the network, rather than being fully absorbed by the affected project. Performance-based incentives, which reward contractors for timeliness, exacerbate these externalities by encouraging self-interested resource reallocation.

*Key words*: Government procurement, Network effects, Empirical operations, Project management
*History*: This version: August 11, 2024

## 1. Introduction

A project involves several participants—including contractors, clients, and subcontractors—all of whom manage multiple projects concurrently and allocate scarce resources (such as workers and machinery) among them. This resource interdependence creates an interconnected network of otherwise unrelated projects.

We demonstrate that, due to this interdependency, the adverse effects of a seemingly isolated disruption (e.g., a weather disruption at a construction site) propagate across the network and, eventually, delay projects and affect participants that had nothing to do with the original disruption. Specifically, when a disruption causes a project to fall behind schedule, participants respond by reallocating resources to speed up the disrupted project at the expense of other projects in the network. This self-interested reallocation may benefit the affected project, but it leaves other projects with less slack time and puts them at risk of future delays by reducing their margins of error.

Moreover, this initial reallocation of resources triggers a domino effect, leading to further reallocation decisions by other participants in the network. Thus, the effects of a seemingly isolated disruption to a single project end up being parceled out across other projects in the network, instead of being fully absorbed by the disrupted project. The project network, in this sense, inadvertently acts as an insurance net for the initially disrupted project. In doing so, however, progress on other projects is hampered as a result of resource-reallocation externalities.

**Figure 1 A network with three highway construction projects (involving two asphalting firms and two excavation companies).**

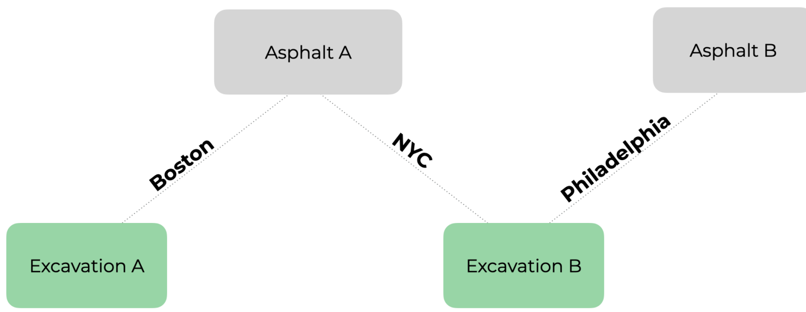


*Note.* In this network, there are three concurrent highway projects: one taking place in Boston (executed by Excavation A and Asphalt A); one in New York City (executed by Asphalt A and Excavation B); and one in Philadelphia (executed by Excavation B and Asphalt B).

We illustrate this idea using a hypothetical project network from Figure 1, which consists of three highway construction projects in Boston, New York City, and Philadelphia. Each project is managed by an excavation company (Excavation A or Excavation B) and an asphalting company (Asphalt A or Asphalt B). Initially, all projects expect to finish ahead of time. But suppose a flood at Boston's project site puts it behind schedule. Asphalt A responds by reallocating machinery and workers, from NYC to Boston, to speed up the Boston project. Now, both the Boston and NYC projects are on track to finish just in time, but without any slack. This reallocation is optimal for Asphalt A, since bringing resources from NYC helps the Boston project avoid a delay. But now, any future disruption at NYC could trigger a delay there (and affect Excavation B for no fault of its own). Furthermore, if such a disruption were to occur, Excavation B may need to reallocate resources from Philadelphia to NYC. Thus, a disruption originating in Boston may trigger a delay in NYC's project and eventually even ripple to Philadelphia, thus impacting participants unrelated to the initial disruption due to decisions made elsewhere in the network.

The main challenge in empirically substantiating this argument is to build a project network and to document how disruptions ripple across it. Doing so requires a centralized source of data on projects, the participants involved in each task, and the disruptions affecting project operations. We overcome this challenge by mapping the network of public projects awarded by the U.S. federal government. We obtain data on 2.4 million projects awarded between 2011 and 2015, involving 124,026 contractors, 27,660 subcontractors, 3,559 awarding offices, and 1,188 types of tasks. The projects span tasks ranging from the construction of a hospital to the installation of streetlights. The data also includes time-stamped records of every modification made to a project, allowing us to create a timeline for each project and track its progress towards completion.

To measure the propagation of localized disruptions in the network, and distinguish such shocks from potential confounders, we focus on disruptions caused by *localized* weather events.[1] For instance, a thunderstorm may flood a construction site or cause electric outages that affect project operations. To this end, we leverage severe weather events affecting a project site using data from the National Oceanic and Atmospheric Administration (NOAA). We then assess the impact of such disruptions on other projects that are linked through the network but are not directly affected by the disruption.

Our identification approach involves grouping the projects into partitions. Each partition contains projects that perform the same task, in nearby regions, and under similar contractual arrangements. Although projects in a partition share many similarities, they differ in three key ways: (i) they are connected to different projects in the network, (ii) only some projects in a partition experienced a localized disruption and (iii) the ones that did experience a disruption did so at different times.

For example, a partition may include five windmill installation projects $(p_1, p_2, p_3, p_4, p_5)$ conducted in Texas in 2011 by similar contractors (who are each linked to different participants in the project network). But only two of these projects experienced a localized disruption, namely $p_1$ and $p_2$, with the disruptions occurring in March 2011 and July 2012, respectively. These key differences enable us to use an identification strategy that exploits two sources of variation within a partition: variation in the cross-section (i.e., only a subset of projects experienced a disruption) and variation in the time series (i.e., the timing of the disruptions varied). If disruptions indeed propagate across the network, we would then expect that the projects linked to $p_1$and $p_2$ are more susceptible to delays than those linked to $p_3$, $p_4$, and $p_5$—but also that this increased likelihood coincides with the timing of the disruptions. Thus, projects connected to $p_1$ should be more vulnerable to delays only after March 2011, while those connected to $p_2$ should be more vulnerable to delays only after July 2012.

We applied this insight using a multi-period difference-in-differences estimator with recurrent and temporary treatment spells, and implementing this design with thousands of partitions in our project network. This analysis provides systematic evidence that disruptions propagate through the network. Specifically, we demonstrate that being connected to a disrupted project increases the likelihood of being delayed by up to 17% in the quarter-year following the disruption, relative to the within-partition control projects. In terms of delay time, being connected to a disrupted project increases

[1] In Appendix C, we also consider reallocation externalities due to other types of disruptions.

the expected completion time by up to 6 days per quarter, relative to the within-partition controls. The analysis also shows that the disruptions propagate, albeit on a smaller scale, to the second tier.

We also analyzed *directly-disrupted* projects and found that projects with fewer network connections were much more likely to experience delays, presumably due to their limited access to resources from the network. Furthermore, network-connected projects suffered almost as much as directly impacted projects, indicating that the costs of disruptions are passed on to other projects in the network, instead of being entirely absorbed by the impacted project.

But are resource reallocation decisions really driving the results? After all, we do not observe which resources are being used by project participants, nor how they are being allocated. With this caveat in mind, we provide evidence pointing to a resource reallocation, by showing that disruptions propagate mainly when participants find it *feasible* to reallocate and when they have *incentives* to do so. Specifically, we show that connected projects are more likely to be delayed when they are geographically closer to the disrupted project and perform the same task. The connected projects also experience greater delays when the disrupted project receives performance incentives (for timely completion) or when the disrupted project has a substantially higher budget. Further, we show that machinery and equipment is more prone to reallocation (following a disruption) as opposed to labor.

Taken together, our results suggest that project delays are driven (at least in part) by a reallocation of resources due to disruptions in the project network. We demonstrate that this spillover effect of resource reallocation externalities is economically significant – our analysis (in Appendix B) provides some conservative estimates of the cost of these delay spillovers. Our work sheds new light on what causes project delays, and what can be done to mitigate their impact.

## 2. Theoretical Foundations

### 2.1. What drives project delays?

According to recent studies, 92% of infrastructure projects are completed behind schedule or are over-budget (Vartabedian 2021). A natural question is, "What causes delays in so many projects?" Much of the work on this point has focused on establishing systematic behaviors that explain why delays have become endemic in project management. Researchers offer four broad reasons:

1. *Poor planning*. Lovallo and Kahneman (2003), Grushka-Cockayne (2020), and Baucells et al. (2024) suggest that delays are driven by planning fallacies and cognitive biases, which leads project participants to be overly optimistic (or pessimistic) and, in turn, to make unrealistic projections.[2] Staats et al. (2012) offer a similar explanation showing that, when the project is large, decision makers underestimate the time required to coordinate a big team. In contrast, Flyvbjerg (2009)

[2] In a related case study, Grushka-Cockayne (2014) demonstrates the existence of planning fallacy using data on construction projects of the New York City Department of Parks and Recreation.

argues that delays are caused due to "strategic misrepresentation" by contractors who purposely set short timelines to secure a contract.

2. *Poor hiring.* Other researchers have shown that projects can suffer when managers do not put in the effort to assemble competent teams. For instance, Coviello and Mariniello (2014) used Italian procurement data to show that delays are more likely when governments fail to publicize project solicitations and, accordingly, to attract capable contractors. And Bordat et al. (2004) found that projects experience delays when they do not foster competition in procurement auctions.

3. *Poor contracting.* Researchers have also linked delays to how contractors and clients set up contractual incentives. For instance, Bajari and Tadelis (2001) find that pegging the risk of delays on the contractor is most effective when the contractor is assigned a complex task, and vice versa. Gopal and Sivaramakrishnan (2008), in addition, use data from offshore software development projects to show that project outcomes improve when contractual incentives are carefully designed. Finally, Warren (2014) show that a surge in contracting officers' workload influences the selection criteria and the contractual parameters on government projects.

4. *Poor oversight.* Studies also examine how the level of supervision, from the client to the contractor, affects project outcomes. For instance, Calvo et al. (2019) show that government oversight hamstrings operational efficiency, especially when contractors are new to public procurement and Coviello et al. (2018) show that red tape in procurement increases expected delays, while Roy et al. (2022) show that projects experience more baseline changes when there is insufficient oversight capacity.

In all cases above, researchers attribute delays to factors that operate *within* project boundaries. This means that project participants have, at least in theory, some degree of control over delays. From this perspective, it is reasonable to penalize or reward participants for completing a project on time. However, in this paper, we show that project outcomes are influenced by decisions made elsewhere in the network. We contribute to the literature by identifying a new factor that influences project delays: one that intertwines the decisions of multiple actors in a so-called "project network."

### 2.2. What is a project network?

A "project" is defined as a one-time process with a clear output and a fixed deadline. This output is achieved when multiple participants contribute resources that are unique and, often, non-substitutable (e.g., workers, machinery, engineers, and inspectors).

To visualize a project network, assume that only two types of participants are involved: *contractors* and *agencies*. Consider a set of $M$ contractors, $\mathbf{C} = \{c_1, ..., c_M\}$, and a set of $N$ agencies, $\mathbf{A} = \{a_1, ..., a_N\}$, which pair up to deploy multiple projects. Because agencies and contractors only pair

**Figure 2 Representation of a project network**

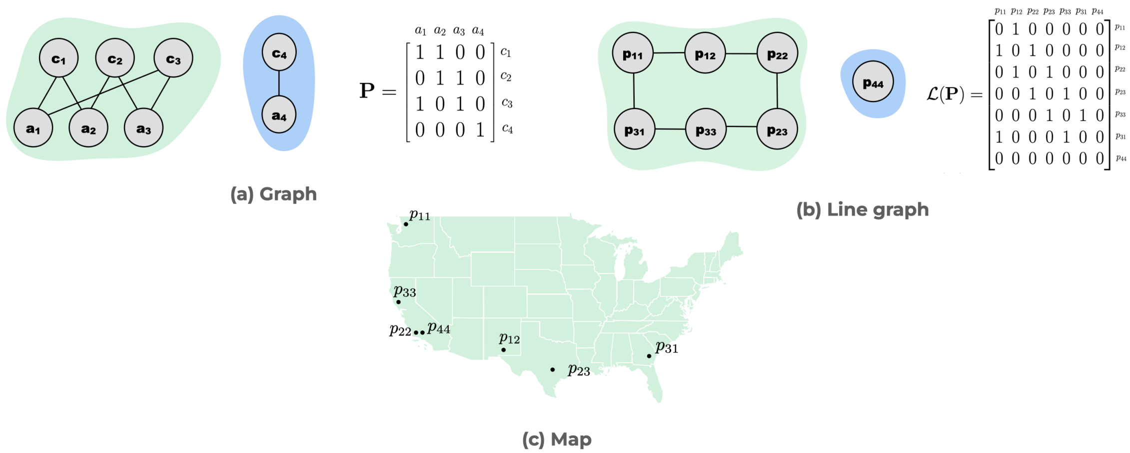


*Note.* This network has seven projects represented by the edges $\{p_{11}, p_{12}, p_{22}, p_{23}, p_{33}, p_{31}, p_{44}\}$. Projects in this network involve four contractors and four agencies, i.e., $\mathbf{C} = \{c_1, c_2, c_3, c_4\}$ and $\mathbf{A} = \{a_1, a_2, a_3, a_4\}$. The network has two connected components $\{p_{11}, p_{12}, p_{22}, p_{23}, p_{33}, p_{31}\}$ and $\{p_{44}\}$. Panel (a) presents the graph of the network, Panel (b) presents the line graph representation, and Panel (c) presents the geographic location of these projects. Matrix $\mathcal{L}(\mathbf{P})$ denotes the incidence matrix of the line graph corresponding to Network $\mathbf{P}$.

with each other and not among themselves, a project network is bipartite.[3] If we let $p_{ij}$ represent a project performed by contractor $c_i \in \mathbf{C}$ and agency $a_j \in \mathbf{A}$, then

$$\mathbf{P} = \begin{bmatrix} p_{11} & \cdots & p_{1N} \\ \vdots & \ddots & \vdots \\ p_{M1} & \cdots & p_{MN} \end{bmatrix}$$

is the $M \times N$ relationship matrix with element $p_{ij} = 1$ if contractor $c_i$ and agency $a_j$ are involved in a project, and is zero otherwise.[4] Figure 2 illustrates a sample network with four contractors and four agencies working simultaneously on seven projects (in different locations, as indicated on the map).

**Distance between projects.** We can represent $\mathbf{P}$ via a line graph, which maps the adjacency between the edges of $\mathbf{P}$. We let $\mathcal{L}(\mathbf{P})$ denote the adjacency matrix of the line graph (see panel (b) in Figure 2). $\mathcal{L}(\mathbf{P})$ is a square matrix of dimension $k$, with $k$ being the number of projects. We can use the line graph to determine the degree of separation between two projects. To this end, we present the following definition.

Definition 1. $\psi^d(p_{ij})$ is the set containing all projects with $d$ degrees of separation from $p_{ij}$.

Applying Definition 1 to the network in Figure 2, we can verify that $\psi^1(p_{11}) = \{p_{12}, p_{31}\}$, i.e., $p_{12}$ and $p_{31}$ have one degree of separation from $p_{11}$. Similarly, we can verify that $\psi^2(p_{12}) = \{p_{23}, p_{31}\}$, i.e., $p_{23}$ and $p_{31}$ have two degrees of separation from $p_{12}$.

[3] In reality, a project could involve multiple types of participants (which could make the network tripartite or multipartite). However, the same concepts highlighted here would apply.

[4] In reality, an agency and a contractor can work on multiple projects at the same time. For notational simplicity, we consider elements of $\mathbf{P}$ to be binary, i.e., they represent a single project.

### 2.3. Disruptions in the project network

The paper presents two arguments and two corresponding hypotheses, which we intend to test empirically. Our first argument is that when a project $p$ experiences a disruption, its impact will propagate across projects that were connected to $p$ but not directly affected by the disruption. To this end, we will empirically examine the following proposition:

***Baseline proposition (disruption spillovers)****. A "localized" disruption in one project will ripple across the same connected component in the project network, by increasing the delay likelihood and the expected future delays of network-connected projects.*

This proposition argues that the costs of a localized disruption end up being borne by other participants in the project network rather than being fully absorbed by the directly disrupted project. To illustrate, consider the network from Figure 2. Suppose that projects $p_{22}$ and $p_{44}$, which are both located in San Diego, CA, as shown in the figure, are identical in all observable respects: they execute the same task, are in nearby locations, have a similar timeline, and operate under the same contractual guidelines. However, $p_{22}$ and $p_{44}$ are in different connected components in the network.

Our baseline proposition argues that if a localized disruption hits $p_{12}$, then $p_{22}$ is more likely to suffer a future delay than $p_{44}$. In other words, we will show that the sole fact that $p_{22}$ was connected to $p_{12}$—whereas $p_{44}$ was not—made it more likely to report a future delay. We also intend to show that these shocks can propagate further in the network, and affect projects that are more than one tier away from the disruption (for instance, $p_{23}$ or $p_{33}$).

But what is driving such spillovers? We have the following proposition.

***Mechanism proposition (resource reallocation)****. The observed disruptions spillovers are driven, at least in part, by a reallocation of resources by the project participants.*

Note that a reallocation will only be necessary if the disruption is significant enough to push the project behind schedule. And if the project does fall behind schedule, participants would typically draw resources from projects with enough slack time, so that both projects remain on schedule after the reallocation. This means that a reallocation of resources by itself is not expected *a priori* to create a delay in any project. It does, however, make other projects more vulnerable to future delays by reducing their margins of error. For example, if a disruption in $p_{12}$ (from Figure 2) prompts the participants to reallocate resources from $p_{11}$ and $p_{22}$, then these two projects would be left with a smaller buffer and, accordingly, be more susceptible to delays. As discussed in Section 1, this reallocation may trigger a domino effect of further re-allocations in the network, and delay projects that are several degrees of separation away from the disrupted project.

## 3. Empirical framework

We test our hypotheses by mapping the network of U.S. federal projects. We first describe the role of contractors and agencies in allocating resources to keep these projects on track.[5]

### 3.1. Role of Agencies

An agency is a government organization that is responsible for overseeing the day-to-day operations of a project. Each agency has a number of "awarding offices" that carry out its responsibilities. These offices verify that the contractor complies with all quality and regulatory guidelines.

**What resources do agencies manage?** A key responsibility of the agencies is to manage and allocate *government-furnished property* (GFP) – these are resources supplied by the government for a contractor to utilize in performing project tasks. Government-furnished property can encompass equipment, tools, materials, vehicles, and facilities. Agency officers can request for equipment to be transferred from storage centers, or from other projects, to any given location for the use of a particular task. Since agency officers are tasked with managing GFP resources across projects, they need to determine when specific equipment is available to the contractor and whether it is concurrently used by other projects. If an agency officer finds it advantageous (or necessary) to reallocate GFP from one project to another, a formal procedure exists for transferring such equipment (see FAR §52.245-1). GFP transfers are quite common—in fact, our data show several records where GFP transfers were made to speed up a project.

**How are agencies held accountable for delays?** Agency officers get evaluated based on their performance through internal assessments. The Federal Acquisition Regulation outlines the authority and responsibilities of agency officers overseeing contracts—including how they are selected, appointed and can be terminated for unsatisfactory performance.[6] Falling behind schedule in a project handicaps agency officers if they wish to advance in their careers. According to a recent report, more than 80% of the officers reported that project outcomes such as timeliness, quality, and cost were factored into their performance ratings (McPhie et al. 2005).

Furthermore, when projects are delayed, the inspector-general may conduct a thorough investigation of the concerned agency and identify the personnel responsible for mismanagement. Punitive actions are then taken against these individuals along with agency-wide recommendations to improve operations—see, for example, Office of Inspector General (2018) and Hull (2019). Finally, the government also compensates the contractor by making equitable adjustments for the costs incurred due to agency-caused delays (see FAR §552.243-71)—we observe several text descriptions in our dataset evidencing this compensation.

[5] For details about project regulations, see the Federal Acquisition Regulation (FAR) book, which is available at https://www.acquisition.gov/sites/default/files/current/far/pdf/FAR.pdf.

[6] See `https://www.acquisition.gov/far/subpart-1.6`

### 3.2. Role of Contractors

A contractor is a firm hired by the government to carry out project tasks. Contractors also own and operate property (referred to as "contractor-acquired property" in FAR §52.245-1); and provide resources such as machinery, tools, equipment, skilled labor, and input materials for the project.

**How are contractors held accountable for delays?** There are two ways in which the government aligns the contractors incentives with its own.

*1. Career concerns (Contractor Performance Assessment Reporting System).* First, the government can impede contractors' ability to bid on future projects if their past performance is not up to the mark. The federal government uses a tier system to rank each contractor's performance in the Contractor Performance Assessment Reporting System (CPARS). This system ranks contractors on various dimensions (one of them being their timeliness or "schedule"), via a five-tier ranking: Exceptional, Very Good, Satisfactory, Marginal, and Unsatisfactory.

*2. Financial penalties (Performance-based incentives).* Second, the government can use financial incentives to influence the contractors' decisions. About 41% of the government contracts include performance-based incentives, which levy fines on (or withhold bonuses from) the contractors when a project is delivered late. When a contract includes a delay-based penalty, this penalty is specified through a liquidated damages clause. Performance-based contracts focus entirely on project outcomes rather than the processes followed to achieve them. In these projects, the contractor typically receives a larger fixed-fee payment but is subject to penalties if the pre-established outcome is not achieved—for instance, if the deadline is not met (see FAR §37.6 for details). This type of contracting aims to reduce oversight and surveillance, allowing the agency to take on a less active role. Without performance-based acquisition, the agency has a more active role in ensuring the contractor finishes on time. In such cases, the agency shares some of the burden for delays, as it is responsible for dictating how the work should be done or developing plans to expedite a delayed project.

### 3.3. Can delays be excusable?

Since we focus on delays that are caused by exogenous disruptions (such as weather) in the network, one may argue that project participants are not responsible for such delays. In 2010, however, the U.S. Federal Court ruled (in Edge Construction Co. v. United States, 95 Fed. Cl. 407, 420) that project participants are contractually liable for all delays, regardless of the root cause, with the exception of delays resulting from government negligence. The Federal Court of Appeals later confirmed that contractual penalties (or rewards) cannot be adjusted due to severe weather. This ruling aimed to address opportunistic behavior from contractors who often exploited weather-related events to justify operational inefficiency and evade penalties (see Appeal of Charles G. Williams Const., Inc., No. 42592, 92-1). Therefore, all project participants have an incentive to avoid being late, even when the delay is due to factors beyond their control.

### 3.4. Data sources

We obtain data on project and modification records from USAspending.gov; and data on weather records from the National Oceanic and Atmospheric Administration (NOAA).

**Project and Modification records.** We obtained 2.4 million project records from the USAspending database, including every infrastructure project with a budget surpassing $10,000. There are 261 variables in each project record. We use the variables describing the characteristics of a project's (1) task, (2) contractual terms, (3) contractor, and (4) agency office.

*1. Task characteristics.* We retrieve data on each project's start date, the expected completion date, and the address where it was carried out. We also observe a variable describing and categorizing the nature of the task. There are over 1,000 task categories, represented via a four-digit code. The first digit provides a broad categorization of the task (e.g., *Y = Construction of Structures/Facilities*), whereas the latter three digits provide a more granular description about each task (*e.g., Y1BD = Construction of airport runways and taxiways*).

*2. Contract characteristics.* We observe several variables on the contractual terms of a project, including the initial budget, payment scheme (e.g., fixed price, time and materials, or cost-plus pricing), and details on the acquisition scheme (performance-based or non-performance-based). We also observe whether the project was competitively awarded and the number of bids made in the solicitation stage and all regulatory bylaws affecting a particular project.

*3. Agency characteristics.* We observe details of the agency offices responsible for overseeing each project, as well as the government agency and sub-agency that it represents (e.g., Air Force is a sub-agency of the Department of Defense).

*4. Contractor characteristics.*We observe the contractor's name, DUNS identification number, headquarters location, annual revenue, number of employees, and industry (NAICS) code for the project in which the contractor is involved.

We winsorize the numerical variables to reduce the influence of reporting errors or outliers in the data. Table 1 illustrates a project record drawn from the database. Table D. 3, in the appendix, provides descriptive statistics about the project records in our sample.

Any disruption or event affecting a project's operations must be logged as a "modification record." We observe around five million time-stamped modification records detailing such changes (e.g., see Table 2). These records have an identification number, a text description, and a category code describing the nature of the modification. Moreover, we observe if the completion date of the project was changed as a result of the event, for instance, if there was a delay caused by a disruption. Note that modifications are reported even if there is no change to a project's deadline. Several modifications are purely administrative in nature and lead to no change on project outcomes.

**Table 1 Sample project record**

| RECORD OF PROJECT | |
|---|---|
| PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER FA446011C0005 | |
| | |
| TASK CHARACTERISTICS | |
| Description | HOSPITAL RENOVATION AND MRI ADDITION, MARINE CORPS BASE, CAMP LEJEUNE, NORTH CAROLINA |
| Task category | Y: Construction of structures/Facilities |
| Task code | Y141: Construct/Hospitals & Infirmaries |
| Start date | 09/27/2012 |
| Expected completion date | 10/29/2015 |
| Actual completion date | 10/29/2015 |
| Place of performance | CAMP LEJEUNE, NC 28542-0089 Congressional District: NC-03 UNITED STATES |
| CONTRACT CHARACTERISTICS | |
| Initial Contract value | $62,159,163 |
| Type of Contract Pricing | J: Firm Fixed Price |
| Performance-Based Acquisition | No |
| Solicitation procedure | A: Full and open competition |
| Solicitation procedure | NP: Negotiated Proposal/Quote |
| Davis Bacon Act | X: Not Applicable |
| Clinger Cohen Act | X: Not Applicable |
| CONTRACTOR CHARACTERISTICS | |
| Name | W.M.JORDAN COMPANY, INCORPORATED |
| DUNS | 005810890 |
| Address (Branch) | 11010 JEFFERSON AVE NEWPORT NEWS, VA 23601-2717 |
| Parent company | W.M.JORDAN COMPANY, INCORPORATED |
| Parent DUNS | 005810890 |
| Annual revenue | $170.25 million |
| Employees | 340 |
| AGENCY CHARACTERISTICS | |
| Agency: Office | NAVAL FAC ENGINEERING CMD MID LANT |
| Office address | 1322 Patterson Ave. SE, Suite 1000, D.C. 20374-5065 |
| Agency: Sub-department | Department of the Navy |
| Agency: Department | Department of Defense |

*Note.* The complete record (that includes over 260 fields) can be retrieved from USAspending.gov.

**Table 2 Sample modification record**

| MODIFICATION RECORD | |
|---|---|
| PROJECT INSTRUMENT ID: | FA466114P0083 |
| | |
| MODIFICATION NO.: P00001 | EFFECTIVE DATE: 09/05/2014 |
| Major (Parent) Agency | 9700: DEPT OF DEFENSE |
| Contracting Office Agency ID | 5700: DEPT OF THE AIR FORCE |
| CONTRACTOR INFORMATION | |
| GENECO TECHNOLOGIES, LLC | |
| 649 SCOTT ST, 79563-2225 TYE, TEXAS | |
| DESCRIPTION OF MODIFICATION | |
| Completion date change (calendar days) | 14 |
| Total cost price change | $0 |
| TYPE OF MODIFICATION | |
| Reason for Modification | M: OTHER ADMINISTRATIVE ACTION |
| EXTEND PERIOD OF PERFORMANCE BY 14 DAYS FROM 8 SEPTEMBER 14 TO 22 SEPTEMBER 14 DUE TO DAMAGES IN THE VENTILATION SYSTEM | |

*Note.* This table represents an example of a modification record, where the project's scheduled end date was extended by 14 days.

**Weather records.** To obtain causal estimates, our identification strategy will rely on disruptions caused by *localized* weather events. To this end, we use the Storm Events Database, available through NOAA, which includes records on over 50 types of weather events. Each record includes the date when the weather event occurred, the counties affected by it, and the dollar value of damaged property and crops. Table 3 provides a sample record drawn from this database, In Appendix D, Table D. 4 shows descriptive statistics about all weather events, and Figure D.2 maps these weather events by county and event type.

We note that every weather disruption that occurs at a project's site is recorded in the data, irrespective of whether the disruption causes a delay on the project or not.

**Table 3 Sample weather record**

| WEATHER RECORD | |
|---|---|
| EVENT TYPE | FLASH FLOOD |
| Cause | Heavy rain |
| State | Georgia |
| County/Area | Floyd |
| Begin Date | 01/22/2017, 20:20 |
| Begin Location | 1SW COOSA |
| End Date | 01/23/2017, 03:30 |
| End Location | 1SW COOSA |
| | REPORTED DAMAGE |
| Deaths | 0 |
| Injuries | 0 |
| Property damage | $40,000 |
| Crop damage | $0 |
| Episode Narrative | The atmosphere over north and central Georgia was extremely moist and unstable. |

*Note.* The complete record can be retrieved from NOAA's Storm Events Database.

**Figure 3 Illustration of the network.**

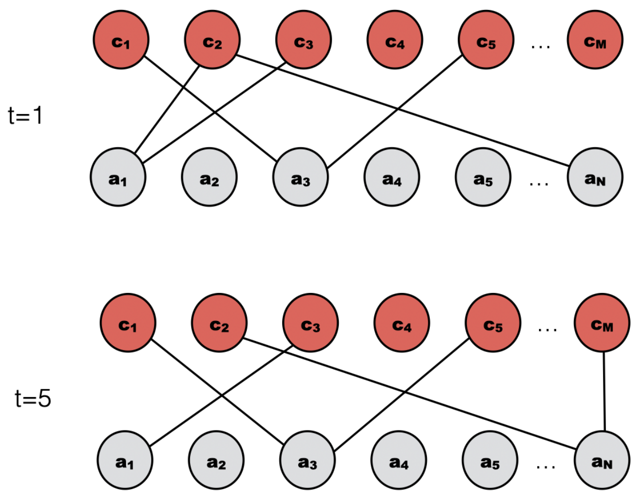


$$\mathbf{P_1} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & \dots & 0 \\ 1 & 0 & 0 & 0 & 0 & \dots & 1 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 \end{bmatrix}$$

$$\mathbf{P_5} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 1 \\ 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \dots & 1 \end{bmatrix}$$

*Note.* Based on our Definition from Section 2, the $i, j-th$ element of matrix $\mathbf{P_t}$ equals one if contractor $i$ and agency-office $j$ were executing a project at time $t$.

## 3.5. Mapping the network

From the project records, we retrieve each project's start and actual completion date to create a network that includes $M = 124,026$ contractors and $N = 3,559$ awarding offices, and that spans the years 2011 through 2015. We use quarter-years as our time unit of analysis.[7] The structure of the network changes in each quarter, as does its adjacency matrix—depending on which projects are currently being executed, and which ones have ended. This means that we have a unique matrix $\mathbf{P_t}$, and line graph $\mathcal{L}(\mathbf{P_t})$, for each quarter, as illustrated in Figure 3. For instance, if project $p_{ij}$ began at $t = 1$ and concluded at $t = 5$, then the $(i,j)^{th}$ element would equal one in Matrices $\mathbf{P_1}$ to $\mathbf{P_5}$, but zero in Matrices $\mathbf{P_j}$ for $j \geq 6$.

[7] Given that disruptions are relatively rare occurrences and project timelines are long, this time window is adequate and consistent with what is used in the literature (e.g., Barrot and Sauvagnat 2016).

**Figure 4** **Illustration of a project timeline.**

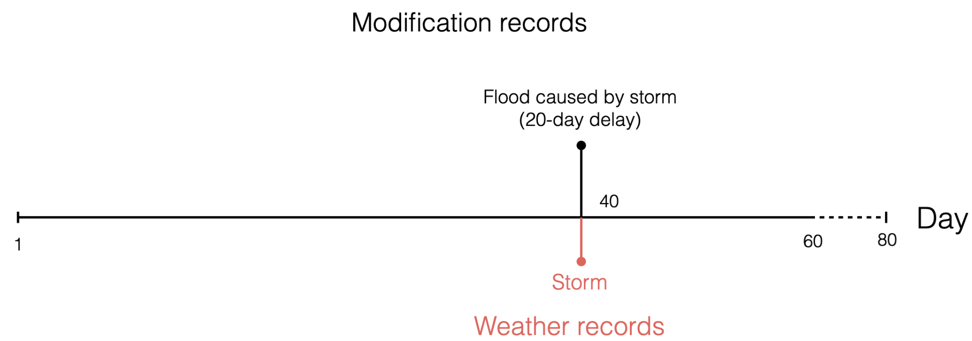


### 3.6. Mapping the disruptions

Using severe weather events at a project site allows us to obtain clean and unbiased estimates of the causal effect of disruptions on the project network. This is due to the following reasons.

*1. Weather-induced disruptions are random and exogenous.* The occurrence of weather disruptions is random, exogenous, and independent of the nature of the project task. This is unlike other disruptions such as a labor strike or a machine breakdown which may be affected by the actions of the contractor or the agency.

*2. All weather disruptions are recorded without omission.* Since the NOAA data is independent of USAspending, we observe severe localized weather events at a project site even if they are not documented as disruptions in the project records. All other types of disruptions compel us to rely solely on project modification records, which could introduce biases. For instance, we will not be able to ascertain if there were unrecorded machine breakdowns or paperwork issues. In theory, these disruptions should be recorded by the contracting officers but there we cannot determine if any recording biases or omissions occur.

*3. The severity of the event can be objectively measured.* Using data from the NOAA's weather records, we can ascertain the degree of severity of disruptive events such as hurricanes, storms, or tornadoes. We also know the precise areas that were affected by the event, and the overall property and material damage resulting from it. Unlike the disruptions reported in the modification records, the severity recorded for weather events are quantifiable and not subjective.

*4. We observe the exact disruption time.* Another benefit of using NOAA's records is that we observe the exact time when a weather disruption occurred. In contrast, there is likely some lag between the event and the date in which the modification record was logged in project records.

Collectively, these reasons will allow us to measure spillovers consistently and without questioning the nature, endogeneity, or accuracy of the event.[8]

[8] While weather disruptions lend themselves to cleaner estimation, in Appendix C, we re-estimate our results by including other types of disruptions.

### 3.7. Mapping the timelines

By combining information from the two data sources, we created a timeline for every project in our network. In particular, we use (i) the project and modification records to obtain each project's start date, initial deadline, and changes to the completion date; and (ii) weather records to obtain disruptions affecting a given project.

To calculate delays, we obtain information on changes to a project's completion time. Consider, for example, a project with the timeline illustrated in Figure 4. Using project records, we could ascertain that this project started at $Day = 1$, had an expected completion date at $Day = 60$, but was ultimately completed at $Day = 80$. Moreover, using weather records, we could observe that this project was affected by a disruption at $Day = 40$ arising from a flood caused by a storm.

## 4. Econometric model

Here we present the identification strategy used to test the baseline hypothesis, which states that a localized disruption in one project will propagate through the network by indirectly increasing the likelihood and magnitude of delays at connected projects.

### 4.1. Designing an identification strategy

In designing an identification strategy, it is helpful to imagine an "ideal" controlled experiment to estimate disruption spillovers in a project network. Our experiment would begin by picking a large set of projects (say, 10,000 pavement repair projects), all of which are disconnected from each other. We would then create a localized disruption at half of these projects (at different times), while the other 5,000 projects remain undisturbed. After each (simulated) disruption event, we would compare (i) the pre- and post-event performance of projects connected to the disrupted project, to (ii) the pre- and post-event performance of projects not connected to the disrupted project. We can confidently infer network spillovers if the first group of projects consistently falls behind schedule and the timing coincides with each disruption event.

While this "ideal" experiment is clearly infeasible in practice, we can design an identification strategy that resembles it. To do so, we will need to formally define a disruption and find a way to consistently measure a disruption spillover. We formally define a disruption event as follows.

Definition 2. *Let $\delta_{p,t} \in \{0, 1\}$ define a disruption event, such that $\delta_{p,t} = 1$ if project $p$ experienced a disruption in period $t$, and $\delta_{p,t} = 0$ otherwise.*

To measure how disruptions spill over in the project network, our identification strategy will compare hundreds of thousands of projects that are similar across dozens of characteristics, including task, location, timeline, and contractual terms. However, these projects will differ in one important way, namely, that they will have different network connections. Consider, for example, two comparable projects $p'$ and $p''$. These projects will fall into one of three cases in each period (as shown in Figure 5):

**Figure 5** **Three cases for a pair of duplicate projects** $(p', p'')$

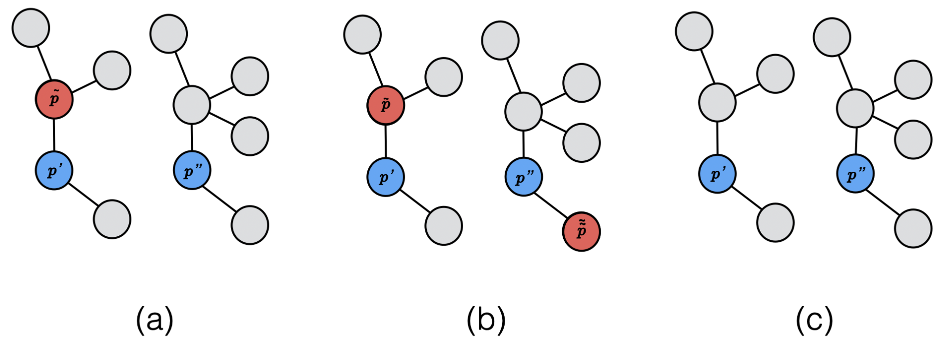


*Note.* Two comparable projects, $(p', p'')$ may fall into three cases. In Case A (depicted in the leftmost figure), only $p'$ is connected to a disrupted project ($\tilde{p}$). In Case B (depicted in the center), both projects were connected to a disrupted project ($\tilde{p}$ and $\tilde{\tilde{p}}$). In Case C (depicted in the rightmost figure), none of the projects were connected to a disrupted project.

**Figure 6** **Evidence of disruption spillovers – Illustration**

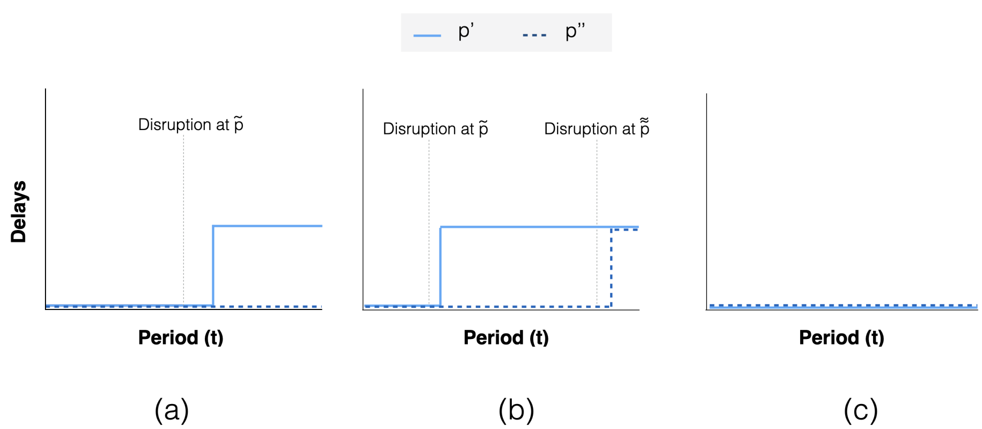


*Note.* In the three figures above, we show how delays could occur at two duplicate projects—for Cases A, B, and C—in the presence of disruption spillovers.

**Case A.** Project $p'$ was directly connected to a disrupted project (say $\tilde{p}$) but $p''$ was not.

**Case B.** Both $p'$ and $p''$ were connected to different disrupted projects (say $\tilde{p}$ and $\tilde{\tilde{p}}$).

**Case C.** Neither $p'$ nor $p''$ were connected to a disrupted project.

Comparing the outcomes of $p'$ and $p''$ in all three cases will allow us to assess the existence of disruption spillovers. In particular, if our hypothesis is correct (and localized disruptions indeed propagate) then the following will be true:

**Case A.** Project $p'$ is more likely to experience a delay after the disruption at $\tilde{p}$ (see Figure 6 (a));

**Case B.** Both $p'$ and $p''$ are more likely to fall behind schedule, following the disruptions at $\tilde{p}$ and $\tilde{\tilde{p}}$ (see Figure 6 (b)); and

**Case C.** Neither project is more likely to experience a delay (see Figure 6 (c)).

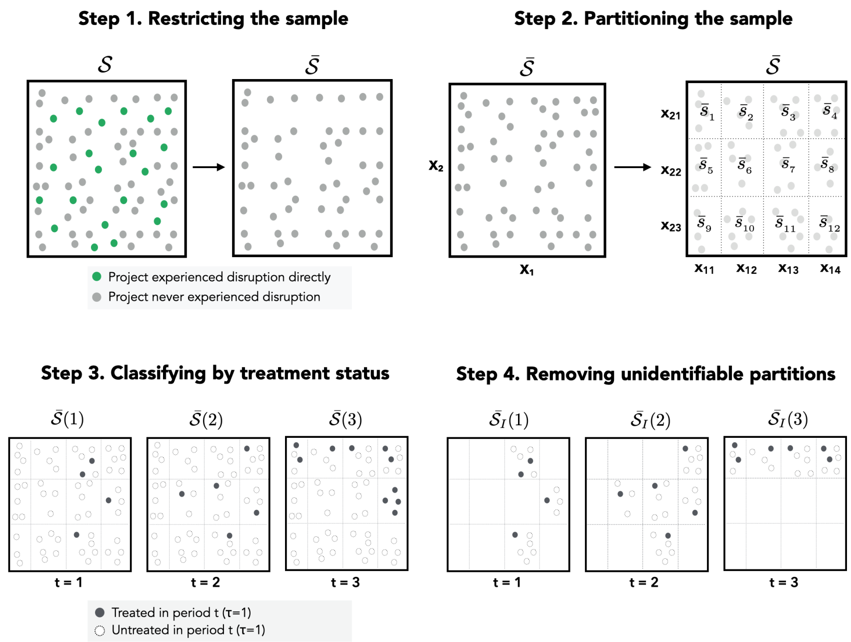


**Figure 7**

Illustration of project matching strategy

*Note.* The figure above illustrates the process of finding a grouping of duplicate projects.

If this phenomenon is systematically replicated across hundreds of thousands of comparable projects, then we will be able to confidently conclude that disruptions spill over in the project network. In the next sections, we will implement this identification strategy.

## 4.2. Estimation procedure

We use a five-step estimation procedure, which borrows elements from Ho et al. (2007) and Barrot and Sauvagnat (2016)). Figure 7 illustrates this approach.

### Step 1 (Restricting the sample)

We start with a sample of 2.4 million projects, $\mathcal{S}$, and want to determine how projects are affected by localized disruptions arising *elsewhere* in the network. This means that both our treatment and control observations will need to be restricted to the subsample of projects that did not experience a disruption directly on their site, that is, to subsample $\bar{\mathcal{S}} = \{p \in \mathcal{S} : \delta_{p,t} = 0 \text{ for all } t\}$, as shown in the upper-left pane of Figure 7. In this step, we dropped 287,102 projects or approximately 11% of the sample.

### Step 2 (Partitioning the sample)

Because our dataset is large, we can find thousands of groups of projects that are similar across multiple characteristics. We partition $\bar{\mathcal{S}}$ into $K$ subsets $\{\bar{s}_1, \bar{s}_2, \ldots, \bar{s}_K\}$ of comparable projects using Coarsened Exact Matching. This methodology follows four sub-steps:

1. *Define a vector of categorical variables* ($\mathbf{X}'$): In our analysis, $\mathbf{X}'$ comprises four categorical variables, which describe: (i) the task type; (ii) the pricing scheme, (iii) the project's NAICS code; and (iv) the geographic location in which the project took place.
2. *Define a vector of numerical variables* ($\mathbf{X}''$): In our analysis, $\mathbf{X}''$ comprises five numerical variables, which describe (i) the contractor's annual revenue in the preceding year, (ii) the number of employees that the contractor employs, (iii) the number of bids the project received, (iv) the project's initial duration, and (v) the project's initial budget.
3. *Coarsen variables in* $\mathbf{X}''$: Once the two vectors are defined, the variables in $\mathbf{X}''$ are coarsened into discrete categories. For instance, the annual revenue variable may be discretized into four or five bins.
4. *Partition projects.* After coarsening the numerical variables, we partition the projects. Two projects belong to the same partition if they match exactly across all variables in $\mathbf{X}'$ and they also fall within the same bin for every variable in $\mathbf{X}''$. Figure 7 shows a partition with two matching covariates, $\mathbf{X_1}$ and $\mathbf{X_2}$, where $\mathbf{X_1}$ could be a discrete variable with four categories, and $\mathbf{X_2}$ could be a numerical variable that is coarsened into three bins.

**Partition granularity.** Our dataset allows us to create partitions with varying degrees of granularity — from very fine to very coarse. We could partition the projects using stringent criteria, for example, by using the *four-digit* task code (with $\geq 1,000$ categories), the *six-digit* NAICS code (1,176 categories), and the *county* where the project is performed ($\approx 3,000$ categories). Similarly, we could create hundreds of discrete bins for each numerical variable using an automated technique such as the Sturges rule.[9] This approach can allow us to get highly accurate matches, but may also sacrifice valuable data.

To get a more representative sample, we could also match projects using higher level discrete variables (e.g., the two-digit task code or the state where the project is performed), and manually specify the number of bins to be created for the numerical variables. To balance the trade-off between sample size and project similarity, we use several different partitioning approaches and show that our results are robust (in both sign and significance) to the partitioning approaches. Specifically, in Table 4, we match projects on the two-digit task code and two-digit NAICS code, and vary the number of bins (i.e., granularity) for numerical variables. In Table A.2, we progressively alter the partition granularity by varying the levels for discrete variables (e.g., matching at the county level and 5-digit, 4-digit, or 3-digit NAICS codes). Further, in Figures A.1 and A.2, we also ran a series of regressions using other matching methods such as optimal full matching, nearest neighbor matching

[9] Sturges algorithm is used to determine the bin size in a data-driven way. This algorithm creates evenly spaced bins based on the scale of each numerical variable. This technique not only prevents researcher manipulation to "enhance" the results, but also ensures a good fit.

with different number of neighbors, propensity score matching with varying caliper sizes, Mahalanobis distance matching, and Lasso and tree-based approaches. The results are insensitive to the choice of matching methodology.

**Correlation between categories.** Finally, we note that there is a correlation among the categories, but this correlation is both natural and expected. For example, the task of *Construction of Marine Facilities* would be confined to coastal counties, and the task of *Construction of Air Traffic Control Towers* would only occur in counties with airport facilities. Similarly, it would be quite rare to find roofing contractors involved in aircraft manufacturing projects. In other words, a given project task would not require contractors from all industry (NAICS) codes.

This correlation also implies that projects are not uniformly distributed across all possible partitions. Although the number of potential partitions is exceedingly large, reaching into the millions, a significant portion of these partitions are empty. The majority of observations in our sample are concentrated within a select subset of partitions. For example, there are over a thousand task categories, but 81.37% of the projects fall within the top 200 tasks. Conversely, the bottom 200 categories comprise less than 2% of the sample. The concentration of project tasks simply reflects the fact that some tasks are inherently rare (e.g., construction of dams), while others are quite common (e.g., repair of office buildings).

Another potential concern with the data is that the NAICS code and task category could be highly correlated. While there is some correlation, it is not a perfect match. For instance, a firm with the NAICS code 238160 (*Roofing Contractors*) or 238110 (*Poured concrete foundation and structure contractors*) are involved in various types of repair tasks.

**Step 3 (Classifying observations by treatment status).**

We want to compare the outcome of "treated" and "untreated" projects in a given time period. To this end, we expand each partition $\bar{s}_k$ by period, generating partitions $\bar{s}_k(t)$ for each $t$, where $t$ is measured at the quarter-year level. This means that while partition $\bar{s}_k$ represents projects with similar characteristics across the sample, $\bar{s}_k(t)$ stands for projects within the partition that were active at time $t$. Figure 7 illustrates the sample expansion across three periods.

We say a project $p$ is treated at time $t$ if an adjacent project was disrupted at time $t$; otherwise, we say that the project was untreated. Formally, we have the following definition.

DEFINITION 3. Let $\tau_{p,t} \in \{0,1\}$ define the treatment status of project $p$ at time $t$, such that $\tau_{p,t} = 1$ if there exists a $\tilde{p} \in \psi^1(p)$ such that $\delta_{\tilde{p},t} = 1$.

In Appendix D, we show the distribution of partitions in terms of cardinality under the most stringent matching criteria (i.e., using four-digit task code, six-digit NAICS code, and county). The sample includes a total of 677,558 project-quarter partitions and, in a given period, about 19.5% of the projects were treated.

**Step 4 (Assigning weights to the sample partitions).**

To determine how a localized disruption propagates, we compare treated and untreated projects within the same partition. To this end, we use Ho et al. (2007)'s weighting technique to obtain a weight vector, $\mathcal{W}$, which assigns weights $w_{p,t}$ to each project $p$ in partition $s_k(t)$. The weighting technique assigns a weight of one to each matched treated project. The matched untreated projects are assigned weights that balance the variation in the number of treated and untreated observations both within and across partitions. Unmatched observations are assigned a weight of zero. For further details, please see Ho et al. (2007).

We end up with up to 50,000 identifiable partitions along agency and contractor nodes under different matching methodologies (see Table 4). Altogether, approximately three million project-quarter observations are used to estimate the effect.

**Step 5 (Running a weighted diff-in-diff estimator).**

As in a prototypical diff-in-diff model, the key idea is to compare the incidence of delays (i) pre-disruption and post-disruption (the time-level difference) and (ii) across the treated and untreated units (the group-level difference). But our model differs from a canonical diff-in-diff model in four aspects:

1. We observe units across multiple periods, instead of two periods.
2. The treatment affects units asynchronously, instead of being a synchronized treatment (because a localized weather event can occur at any time).
3. Observations can be treated multiple times, instead of only once (because a project can be affected by a localized network disruption several times within the timeline).
4. The effect of a shock dissipates across time (because a storm is a one-time event with a temporary effect rather than a persistent policy).

Collectively, these four conditions make our identification stronger. For instance, Condition 2 means that our shock is less likely to be confounded by time-specific unobservable shocks that are correlated with the treatment date (which is one of the biggest drawbacks of a prototypical diff-in-diff model). And Conditions 3 and 4 allow us to observe the impact of a shock on the same unit of analysis, but at different times.

To properly model this empirical setting, we use a multi-period diff-in-diff model with recurrent and temporary treatment spells, developing a methodology that borrows elements from Barrot and Sauvagnat (2016) and Callaway and Sant'Anna (2021)'s doubly-robust estimators.

We let $Y_{p,t}$ denote the completion time of project $p$ at time $t$. Then, $Delay_{p,t} = Y_{p,t} - Y_{p,t-1}$ measures the reported delay for project $p$ in time $t$. Using the weights vector, $\mathcal{W}$, we estimate the weighted average delay of the treatment and control units as follows. With a slight abuse of notation, we let $\tau = 1$ represent all treated projects, and $\tau = 0$ represent all control projects.

$$Delay^{\tau}_{\mathcal{W}} = \begin{cases} \frac{\sum w_{p,t}\tau_{p,t}Delay_{p,t}}{\sum w_{p,t}\tau_{p,t}} & \text{if } \tau_{p,t} = 1 \\ \frac{\sum w_{p,t}(1-\tau_{p,t})Delay_{p,t}}{\sum w_{p,t}(1-\tau_{p,t})} & \text{if } \tau_{p,t} = 0 \end{cases}$$

Using this weighted average, we can obtain a diff-in-diff estimate, which will compare the delays in the treated group against the delays in the control group, as follows.

$$\beta_{DD} = Delay^{1}_{\mathcal{W}} - Delay^{0}_{\mathcal{W}} = \underbrace{(Y^1_{p,t} - Y^1_{p,t-1})_{\mathcal{W}}}_{\substack{\text{Weighted delay for} \\ \text{treated projects in time } t}} - \underbrace{(Y^0_{p,t} - Y^0_{p,t-1})_{\mathcal{W}}}_{\substack{\text{Weighted delay for} \\ \text{control projects in time } t}}$$

To estimate the diff-in-diff effect ($\beta_{DD}$), we set $Delay_{p,t}$ as the model's dependent variable, and regress this differenced variable against a time-specific treatment indicator ($\tau_{p,t}$), a vector of fixed effects, and project controls, i.e.,

$$Delay_{p,t} = f(\underbrace{\tau_{p,t}}_{\text{Treatment variable}}, \underbrace{\gamma_{contractor}}_{\text{Contractor FE}}, \underbrace{\gamma_{agency}}_{\text{Agency FE}}, \underbrace{\gamma_{task}}_{\text{Task FE}}, \underbrace{\gamma_{price}}_{\text{Price FE}}, \underbrace{X_p}_{\text{Project controls}}, \underbrace{\epsilon_{p,t}}_{\text{Error term}} \mid \underbrace{\mathcal{W}}_{\text{Weights vector}})$$

The control variables are specified in each regression table's caption, and all errors are robust and clustered at the project level. We also consider other levels of clustering and alternate standard error computations (in Appendix A.6) to ensure our estimates are robust. Since we use matching along with an outcome regression, our estimation is doubly robust and will yield statistically consistent estimates (as only one of the two models needs to be correctly specified)—see Funk et al. (2011).

## 5. Results

Table 4 presents estimates of the extent to which a disruption in one project spills over to concurrent projects by respectively looking at the spillover effect across (i) agency-connected nodes (Columns I-VI) and (ii) contractor-connected nodes (Columns VII-XII). For instance, if a disruption were to affect a project jointly executed by Agency $a$ and Contractor $c$, then: (i) Columns I-VI's estimates would capture the spillover effect across projects directly connected to $a$, whereas (ii) Columns VII-XII's estimates would show this same effect across projects connected to $c$.[10]

This table includes diff-in-diff estimates for 24 different specifications (two regressions per column). The top estimate is from a logistic model in which the outcome variable, $Delay_{p,t}$, is a binary variable that equals one if there was a positive delay and zero otherwise.[11] The bottom estimate in each

[10] For illustration purposes, Table 4 only display the diff-in-diff treatment effect and omit the single difference estimates and the coefficient estimates of the control variables.

[11] Binary estimates are interpreted as probabilities for ease of exposition.

**Table 4 Spillover effect**

| | Agency nodes | | | | | | Contractor nodes | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI | VII | VIII | IX | X | X1 | XII |
| *Dep. variable:* | | | | | | | | | | | | |
| Delay probability | 0.10*** | 0.10*** | 0.11*** | 0.11*** | 0.12*** | 0.17*** | 0.07*** | 0.08*** | 0.07*** | 0.08*** | 0.05*** | 0.13*** |
| | (0.01) | (0.01) | (0.01) | (0.02) | (0.02) | (0.03) | (0.01) | (0.01) | (0.01) | (0.02) | (0.02) | (0.03) |
| Delay days | 3.25*** | 3.28*** | 3.41*** | 3.74*** | 3.73*** | 4.89*** | 3.24*** | 4.12*** | 4.00*** | 4.32*** | 3.19*** | 6.01*** |
| | (0.47) | (0.48) | (0.50) | (0.53) | (0.56) | (0.89) | (0.62) | (0.65) | (0.69) | (0.78) | (0.86) | (1.47) |
| Agency FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Contractor FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Task FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| *Matching specification:* | | | | | | | | | | | | |
| Number of bins | 2 | 2 | 3 | 4 | 5 | Sturges | 2 | 2 | 3 | 4 | 5 | Sturges |
| Location | No | State | State | State | State | State | No | State | State | State | State | State |
| Identifiable partitions | 22,219 | 49,299 | 52,508 | 52,461 | 51,412 | 29,884 | 19,158 | 45,183 | 46,756 | 45,944 | 44,674 | 24,006 |
| Observations | 1,118,720 | 706,319 | 583,401 | 491,846 | 432,645 | 139,532 | 1,087,065 | 666,372 | 539,474 | 445,101 | 391,094 | 117,751 |
| $R^2$ | 0.23 | 0.25 | 0.25 | 0.26 | 0.26 | 0.31 | 0.08 | 0.10 | 0.10 | 0.11 | 0.11 | 0.15 |

*Note.* This table presents the estimated coefficients for the spillover effect of a network disruption on delay days and delay probability. Columns I-VI show the effect for projects connected through the agency nodes, and Columns VII-XII show the effect for projects connected through the contractor nodes. Treated and control projects are always matched on two-digit task code, two-digit industry code, and four price categories, and numerical variables (number of bids, initial budget, initial duration, annual revenue, and number of employees) using different levels of coarsening. In Columns I and VII, numerical variables are coarsened into two bins. Columns II-VI and VIII-XII gradually increase the number of bins yielding finer partitions. All specifications include fixed effects for the agency, contractor, and project task. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

**Figure 8 Delay days treatment effect conditional distribution (on projects reporting a delay)**

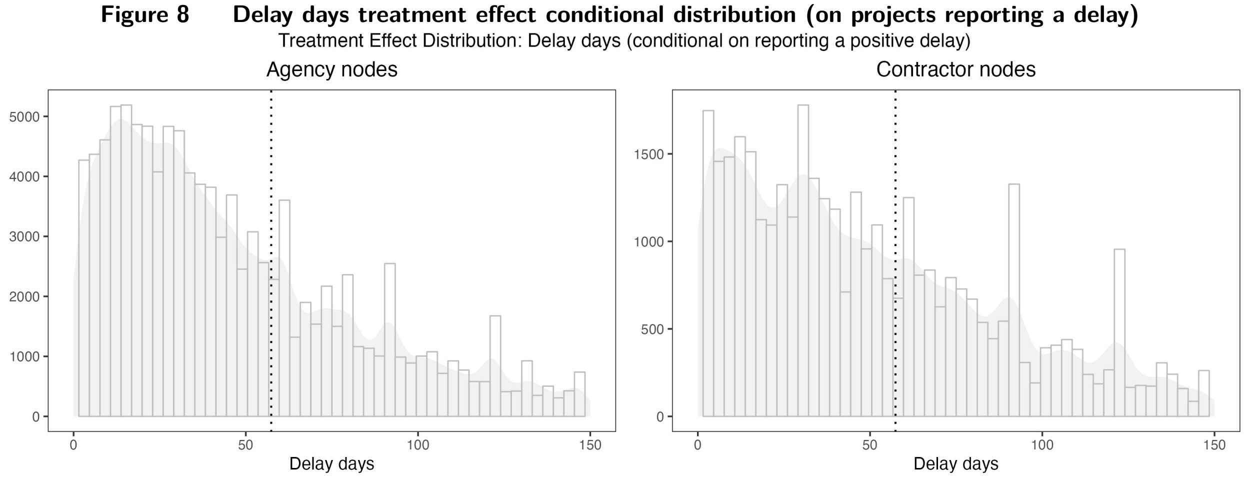


*Note.* These plots show the distribution of the treatment effect across each control-match pairs, for the estimates obtained in Table 4 (conditional on experiencing a positive delay). The conditional average treatment effect is (i.e., for all delayed treated projects) is depicted by the vertical dashed lines.

column is from a linear model in which the outcome variable, $Delay_{p,t}$, is a continuous variable that measures the reported days of delay. Each column displays results from a unique combination of matching specifications, and includes contractor, task, and agency fixed effects.

A look at all 24 regression estimates show evidence of a large, positive, and statistically significant spillover effect (both across agency and contractor nodes). For instance, the logistic regression estimates from Column XII should be read as follows: suppose we take two identical projects, $p_1$ and $p_0$, which are executing the same task, in the same location, in the same period, and under the same contractual terms. Moreover, suppose that neither $p_1$ nor $p_0$ have experienced a disruption at their site. In fact, the only observable difference between these two projects is that $p_1$ was connected to a disrupted project (via the contractor), whereas $p_0$ was not. This difference alone tells us that $p_1$ is 13% more likely to experience a delay following the disruption in a given year-quarter. And, as Figure 8 shows, the average delay in $p_1$ (if one were to happen) would have an average magnitude of 52 days. Thus, the expected spillover effect amounts to about $0.13 \times 52 = 6.7$ delay days.[12] Naturally, this calculation yields somewhat similar results to the estimates from the linear model in Column XII of Table 4, which uses *Delay Days* as the dependent variable and finds an average effect equivalent to 6.01 days.[13] When it comes to agency-connected nodes, Columns I-VI show that the spillover effect is also positive and significant at the 1% level.

In summary, we find that when an external localized disruption hits one project, its impact propagates to adjacent projects. Below, we study whether this effect spills over to multiple tiers, and also examine the impact on the project which was directly affected by the disruption.

**Multi-tier spillovers.** We extend our analysis to see how spillovers propagate beyond the first tier. We begin with the second-tier spillover effect, by using Definition 2 and redefining a treated unit as

$$\tau_{p,t} = \begin{cases} 1 & \textit{if there exists } \tilde{p} \in \psi^2(p) \ \textit{s.t } \delta_{\tilde{p},t} = 1 \\ 0 & \textit{otherwise} \end{cases}.$$

Estimating the second-tier spillover poses complications, both in terms of identification and computation. The main issue arises in terms of identification—the density of the network makes it harder to match treated/untreated pairs, given that the number of untreated observations reduces substantially.[14] A secondary issue is that it is computationally intensive to obtain the second order of the project's adjacency matrix, which is necessary to identify such spillovers. Both issues are common across network studies that attempt to record multi-tier spillover effects (see Goldsmith-Pinkham and Imbens 2013).

We were able to identify the treatment effect on a sample of 59,347 treated and untreated projects. Table 5 shows that the spillover effect of the disruption propagates to the second tier. The second-tier spillover is robust and positive but the magnitude of the effect decreases relative to the first-tier spillover. In other words, we find evidence that localized disruptions propagate to the second tier though the effect reduces.

We also find some evidence of a spillover propagation to the third-tier, albeit the effect is much smaller. As mentioned before, analyzing the spillover effect beyond first tier is computationally challenging and the number of untreated units reduces significantly. This is even more true for the third

[12] Expected Delay Spillover = Delay likelihood (Column XII estimate) $\times$ Avg delay magnitude (Figure 8) = $0.13 \times 52 = 6.7$ delay days

[13] The slight difference in these two figures lies in the fact that the first calculation (of 6.7 delay days) results from the product of a logistic fit and the distributional average, whereas the second estimate (of 6.01 delay days) results from a regression fit.

[14] Note that, by definition, an untreated project is one that had no disrupted connections.

tier analysis. We, therefore, refrain from making conclusive claims about spillover effects at the third tier (or beyond).

**Table 5 Second-tier disruptions**

| | I | II | III | IV | V |
|---|---|---|---|---|---|
| Delay probability | 0.15*** | 0.14*** | 0.14*** | 0.14*** | 0.14*** |
| | (0.03) | (0.04) | (0.04) | (0.04) | (0.04) |
| Delay days | 3.28** | 3.11** | 3.28** | 3.23** | 2.99** |
| | (1.30) | (1.30) | (1.31) | (1.31) | (1.33) |
| Contractor FE | Yes | Yes | Yes | Yes | Yes |
| Agency FE | No | Yes | Yes | Yes | Yes |
| Task FE | No | No | Yes | Yes | Yes |
| Price FE | No | No | No | Yes | Yes |
| County FE | No | No | No | No | Yes |
| Observations | 145,232 | 145,232 | 145,232 | 145,232 | 145,232 |
| $R^2$ | 0.21 | 0.21 | 0.22 | 0.22 | 0.23 |

*Note.* This table presents the estimated coefficients for the delay days and delay probability due to a disruption originating in the second-tier of the project network. Column I presents the specification that includes contractor fixed effects. Columns II, III, IV, and V add to this specification by including fixed effects for agency, task, price, and county, respectively. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

**Directly disrupted projects.** We have shown how localized disruptions propagate in the project network. But what happened to the project that was *directly* disrupted by the storm? And how did the disruption impact its operations? Was this impact moderated by the network?

We paired all projects that were directly disrupted by a given localized event—for instance, all projects that were disrupted by a given hail storm in Wichita County, Kansas. In our sample, the median weather event disrupted four projects. We then regressed the impact of this event as a function of (i) the project's characteristics, (ii) the characteristics of the contractor and the agency, and (iii) the number of network connections each project had. We run the following statistical model.

$$Delay_{p,t} = f(\underbrace{\alpha_{p,t}}_{\text{fixed effects}}, \underbrace{X^p_{p,t}}_{\text{project traits}}, \underbrace{X^a_{p,t}}_{\text{agency traits}}, \underbrace{X^c_{p,t}}_{\text{contractor traits}}, \underbrace{Connections_{p,t}}_{\text{number of network connections}})$$

By analyzing this sample of projects, via a linear-in-means model, we obtain the estimates found in Table 6. Collectively, these findings lead us to conclude that when multiple projects are *directly* disrupted by the same event, the ones with fewer network connections have a statistically significantly higher delay likelihood and magnitude (after controlling for potential confounders). Simply put, the *same* disruption event seems to have a more damaging effect when the project participants are executing fewer projects. Thus, a disruption event is fully absorbed by the disrupted project when it is operating in isolation. When the project has more network connections, however, the impact of the disruption is smoothed out across the network, and partly absorbed by connected projects.

This means that a localized disruption event is cushioned by the network peers, which are in some way 'taxed' by the event. Put another way, the project network inadvertently ends up acting as an insurance net for localized disruptions.

**Table 6 Directly disrupted projects**

| | Delay days | | Delay likelihood | |
|---|---|---|---|---|
| | I | II | III | IV |
| Constant | 21.40*** | | 0.09*** | |
| | (0.89) | | (0.003) | |
| Number of connections | -0.18* | -0.25*** | -0.0008** | -0.0010*** |
| | (0.09) | (0.09) | (0.0003) | (0.0003) |
| Fixed effects | No | Yes | No | Yes |
| Controls | No | Yes | Yes | Yes |
| Observations | 48,880 | 48,880 | 48,880 | 48,880 |

*Note.* This table presents the estimated coefficients for the delay days and delay probability at directly disrupted projects as a function of their network connections. The standard errors (reported in parentheses) are robust and clustered by the disruption event. Significance levels: 10% (*), 5% (**), and 1% (***)

**Robustness checks.** We conduct several robustness tests to validate our findings (see Appendix A). For instance, we employ two alternative metrics to measure delays, namely, delay relative to project's initial duration and delay relative to a project's last reported duration. We also re-estimate our results using fifteen different matching methods, and by employing alternative partition granularity. Finally, we perform cross-sectional and inter-temporal placebo tests, and also consider alternative specifications of the standard errors.

## 6. Are resources reallocated after a disruption?

We argue that the observed spillovers are driven, at least in part, by a reallocation of resources. In other words, the disruption spills over because participants draw resources from the connected (treated) projects to put into the disrupted projects. Text records from our dataset provide anecdotal evidence for this mechanism:

- *Modification to add 16 additional work days due to weather conditions and delays caused by another contractor outside of the control of the prime contractor.*
- *The purpose of this modification is to accept the transfer of materials from another contract to this contract.*
- *Modification to extend the performance period due to weather delays, having to pull sub-contractors off project in order to work the government's #1 priority project.*

Although these anecdotal accounts are consistent with our hypothesis, we cannot directly observe how the project participants utilize their physical and human resources. And even if such data could be gathered, some resources are inherently unobservable—we cannot, for example, measure the managerial time spent on troubleshooting a project. However, we can still provide empirical evidence pointing towards a resource reallocation. We do so by showing that spillovers are more likely whenever participants can *easily* reallocate resources, or if they have *incentives* to do so. Specifically, we consider the following four channels:

*I. Geographic distance.* Consider a hospital construction project in San Bernardino, CA, that gets disrupted by a localized weather event. Now, suppose that the contractor is concurrently managing two other projects: one in San Diego, CA and one in Boston, MA. In this scenario, it is easy to move slack equipment and workers from San Diego to bring the disrupted project up to speed, but it is expensive, if not infeasible, to reallocate such resources from Boston. Accordingly, if the resource reallocation hypothesis is true, we should observe a greater spillover effect for projects that are geographically closer to the disrupted project, and a lesser effect for those that are further away.[15]

*II. Same task category.* Suppose a disruption occurs in an asphalting project, and the contractor is concurrently managing two other projects: (i) another asphalting project and (ii) a plumbing project. The contractor can reallocate task-specific resources (e.g., asphalting machines) from the first project. It is far-fetched, however, to argue that plumbing tools can be useful at all. Thus, if our hypothesis is correct, we should observe that a disruption spills over more strongly to concurrent projects that carry out the same task as the disrupted one.

*III. Performance-based incentives* In Section 3.2, we discussed that federal projects are often acquired under a performance-based acquisition scheme. On these projects, contractors receive explicit contractual incentives for project outcomes and are, therefore, fully accountable for any delays or failures. Consequently, performance-based acquisition lessens agency oversight, shifting the burden of project performance onto contractors. This means a contractor will be more motivated to shift resources to prevent a delay under performance-based acquisition, whereas the agency will have fewer incentives to do so. Therefore, when the disrupted project includes performance incentives, we expect the spillover effect to be *greater* on concurrent projects held by a contractor, and *lesser* effect on concurrent projects that are held by the agency.

*IV. Relative-project importance.* Suppose we observe a disruption affecting two projects: Project X, a high-profile project with a large budget, and Project Y, a small project involving a modest budget. Participants in Project X will have a greater incentive to keep it on track and, as a result, to reallocate resources to keep it on schedule. Therefore, we expect the spillover effect to be larger when the budget of the disrupted project is higher than that of the connected projects.

To test the above hypotheses, we estimate the treatment effect for each treated project (Ho et al. 2007) and examine how it varies with respect to the four channels identified above. For instance, we compare the average treatment effect on projects that are performing the same task as the disrupted project, relative to the average treatment effect on projects that are performing a task that is different from that of the disrupted project.[16] This is analogous to the approach in Callaway and

[15] Recall that the control projects are located at a very similar location from the treated project (e.g., the matched control project would also be located in San Diego or Boston, just like the treated one). Thus, this indirect effect would also influence the control project and, accordingly, would not be confounded.

[16] To do so, we restrict our sample to those treated projects that are connected to only one disrupted project.

Sant'Anna (2021) that estimates treatment effect heterogeneity by studying how "group-time average treatment effect" varies with the parameters of interest. Table 7 shows evidence of the aforementioned

**Table 7 Diff-in-diff estimates as a function of incentives to reallocate resources**

| | Treatment effect (in days) | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Agency nodes | | | | Contractor nodes | | | |
| | I | II | III | IV | V | VI | VII | VIII |
| Geographic distance (100s of miles) | -0.91*** (0.34) | | | | -0.97*** (0.17) | | | |
| Same task category | | 3.81** (1.86) | | | | 6.47** (2.58) | | |
| Performance incentives | | | -4.15*** (1.43) | | | | 6.71*** (1.82) | |
| Relative project value | | | | 0.07** (0.04) | | | | 0.36*** (0.13) |
| Task FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Num treated | 11,394 | 31,710 | 31,710 | 31,409 | 19,413 | 25,875 | 24,517 | 21,828 |
| $R^2$ | 0.09 | 0.01 | 0.05 | 0.05 | 0.16 | 0.14 | 0.16 | 0.14 |

*Note.* This table presents the estimated coefficients for the treatment effect as a function of participants' incentives to reallocate resources. Each specification controls for the project's initial budget and duration, and number of connections; and includes or excludes task fixed effects. Columns I-IV presents the results for spillover effect along agency nodes, and Columns V-VIII shows the results for contractor nodes. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

moderating effects. Estimates in Columns I and V tell us that, for example, the spillover effect reduces by approximately 1 day for every 100 mile increase in the distance between the disrupted project and the treated project. Thus, distance is a significant negative moderator of the spillover effect. Columns II and VI show that the expected size of the spillover significantly increases when the connected project has the same task category code as the disrupted project. Columns III and VII show that when the disrupted project includes performance-based incentives, the spillover effect is *lower* across agency-connected nodes, but significantly *higher* across contractor-connected nodes. As discussed earlier, this is because performance-based projects impose opposing incentives on contractors and agencies. Finally, Columns IV and VIII show that the spillover effect is larger when the budget of the disrupted project is considerably higher than the budget of the concurrent projects.[17]

## 7. Which resources are reallocated after a disruption?

Having provided evidence in support of resource allocation, we now turn to the question of which types of resources are being reallocated between projects. We conjecture that machinery and equipment is easier to reallocate between projects as opposed to labor or managerial effort. We test this hypothesis

[17] Our estimation assumes that network structure is orthogonal to the treatment, in that there is no treatment-induced endogenous change in the network. This is plausible because the treatment is exogenous and network links between agencies and contractors are established according to policies outlined in the federal acquisition regulations.

in two ways: (1) examining the extent of capital-intensive inputs used in a project, and (2) examining the role of government furnished property and equipment that agencies manage.

### 7.1. Capital-intensive resources

Projects primarily use two types of inputs: capital and labor. We compare how reallocation externalities propagate in the network when disruptions affect capital-intensive versus labor-intensive projects. To this end, we calculate the capital-to-labor ratio (CLR) for each task following the approach in Calvo et al. (2019) and Serpa and Krishnan (2018). The CLR represents the average value of equipment capital expenditure (in dollars) per employee on a given task. For instance, the task of constructing highways and roads is more machine intensive (with a CLR of 3.81), whereas the task of constructing water supply facilities is relatively more labor-intensive (and has a CLR of 0.02). The CLR varies from 0.01 (lowest percentile) to 3.81 (highest percentile) with a median of 0.10. Given that the distribution is highly skewed, we normalize it by taking a log-transformation and, then, regress the conditional treatment effect as a function of the logged CLR. Table 9 shows that

**Table 8 Sample task categories and their average capital-to-labor ratio**

| Code | Task description | CLR |
|---|---|---|
| Y1LZ | Construction of Parking Facilities | 0.013 |
| Y1NE | Construction of Water Supply Facilities | 0.02 |
| N073 | Installation of Equipment- Food Preparation and Serving Equipment | 0.12 |
| Y1ND | Construction of Sewage and Waste Facilities | 0.22 |
| Y222 | Construction of Highways, Roads, Streets, Bridges, and Railways | 3.81 |

**Table 9 Diff-in-diff estimates as a function of resources used**

| | Treatment effect (in days) | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | Agency nodes | | | | Contractor nodes | | | |
| | I | II | III | IV | V | VI | VII | VIII |
| Intercept | 8.30*** | | 15.37*** | | 9.65*** | | 22.73*** | |
| | (1.21) | | (1.73) | | (1.57) | | (2.47) | |
| Government-furnished property | 11.16*** | 8.08*** | | | -10.67*** | -2.64 | | |
| | (2.36) | (2.44) | | | (3.25) | (3.31) | | |
| Capital-to-labor ratio | | | 1.97*** | 2.01*** | | | 5.10*** | 3.09*** |
| | | | (0.39) | (0.41) | | | (0.76) | (0.91) |
| Task FE | No | Yes | No | Yes | No | Yes | No | Yes |
| Controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Num treated | 31,710 | 31,710 | 31,707 | 31,707 | 25,875 | 25,875 | 25,873 | 25,873 |
| $R^2$ | 0.005 | 0.05 | 0.005 | 0.05 | 0.08 | 0.13 | 0.08 | 0.13 |

*Note.* This table presents the estimated coefficients for the treatment effect as a function of resources used on a project. Each specification controls for the project's initial budget and duration, and number of connections; and includes or excludes task fixed effects. Columns I-IV presents the results for spillover effect along agency nodes, and Columns V-VIII shows the results for contractor nodes. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

disruption externalities propagate more when the disrupted project is capital-intensive i.e., it uses more capital equipment than labor. For a 1% increase in a task's CLR, the treatment effect increases by 0.02 days across agency-connected nodes and 0.03 days across contractor-connected nodes. Thus, participants tend to reallocate resources more when the disrupted projects rely heavily on machinery (instead of relying on labor). This is intuitive because machinery (e.g., excavators) is typically standardized and can be easily transferred to other similar projects, while workers might need to adapt to project-specific conditions. Furthermore, transporting machinery is a straightforward process whereas relocating labor is more complex, requiring accommodations and potential disruptions to the workers' lives. Labor regulations also complicate the reallocation of human labor between projects, particularly across states due to labor laws and collective agreements.

An alternative explanation for the observed results could be that weather-related disruptions require additional machinery capacity rather than labor. However, in Appendix C, where we directly measure disruptions affecting labor, we do not find any evidence of spillover effects on the network, which makes this explanation less plausible.

### 7.2. Government-furnished property

In some projects, the federal government provides equipment and machinery for the execution of tasks—these resources are referred to as government-furnished property and are managed by agency offices. We test whether government-furnished property is a resource that is being reallocated between projects in the event of a network disruption. To that end, we use a binary variable from the USAspending dataset that indicates whether or not a given project relies on government-furnished property. If government-furnished property is indeed one of the resources being reallocated, we should see greater spillover effect on agency nodes when disrupted project utilizes government-furnished equipment. Conversely, the treatment effect should be considerably smaller across contractor-connected nodes when the disrupted project uses government-furnished property (because contractors do not have the authority to (re-)allocate GFP).

Table 7 substantiates our hypothesis—the treatment effect is significantly higher on agency-connected nodes when the project uses government-furnished property, while the effect is statistically and numerically insignificant on the contractor nodes (after controlling for project task).

## 8. Concluding remarks

In this paper, we show that a seemingly localized disruption at a project can spill over and delay other projects that are linked in a "network" through shared contractors and clients. We demonstrate that this spillover effect is caused by participants reallocating resources between projects.

Previous research has recognized and studied spillover effects in supply chains and social networks. For instance, studies have shown how shocks propagate across production networks—including shocks related to innovation and productivity (see Bellamy et al. 2014, Serpa and Krishnan 2018), economic uncertainty (see Osadchiy et al. 2016), disruption events (see Jain et al. 2022, Schmidt and Raman 2022), and financial risk (see Wang et al. 2021).

Supply chain networks are well-defined and observable, but the notion of a project "network" is not well established. There is at best some recognition that contractors own a portfolio of projects and that there is interdependency within this portfolio (see, for example, Girotra et al. 2007's work on pharmaceutical project portfolios). By mapping the network of U.S. federal projects and using the occurrence of severe weather events at a project site as the exogenous source of disruptions, we show how interdependencies don't just exist within project portfolios but also extend to multiple tiers.

Project delays have complex and multi-faceted causes ranging from poor planning by contractors to poor oversight by clients. Our work contributes to the extensive literature on project management by identifying disruptions at connected projects as an under-explored cause of delays. We show that such delay spillovers are statistically significant i.e., project managers face negative "resource reallocation externalities" due to disruptions elsewhere in the network. In Appendix B, we estimate the economic cost of (first-tier) delay spillovers using methods commonly used by project managers and courts to assess the dollar value of a delay. We find that network disruptions and resulting re-allocations impose an additional cost of *at least* \$4,927-\$10,107 per quarter on the connected projects.

Although our analysis relies on data from U.S. public projects, similar resource reallocation dynamics apply to any project network. If, for example, a home renovation project falls behind schedule, its contractor will likely pull resources from other similar projects to bring it up to speed. Doing so will consume slack from other home renovation projects and make them more vulnerable to delays. We argue that this fundamental insight would apply to all project networks – public or private. However, one feature that is idiosyncratic to public projects is that clients (i.e., government agencies) play an active role in managing projects and allocating resources (e.g., through their ability to deploy government-furnished property). Consequently, reallocation spillovers through client nodes may not apply to private projects. We also note that mapping the network of private projects is typically infeasible due to data availability limitations. In contrast, the USAspending dataset allows us to map a large network of projects and quantify the interdependency between them.

Our results also connect to the literature on process flexibility and chaining (Jordan and Graves 1995). While the chaining literature emphasizes the benefits of being able to share resources in a network, we highlight the potential downside of these network connections i.e., our results point to a "dark side" of operational flexibility. In a traditional chaining context—e.g., a chain of production facilities—the decisions are made by a central planner. In a project context, resource reallocation decisions are made by self-interested decentralized agents without considering the negative externalties on other participants in the network.

This prompts the question of whether operational buffers can mitigate reallocation externaltiies. We conjecture that large contractors may be able to invest in safety capacity and pool resources across projects. Therefore, their projects may be able to absorb the effect of a localised disruption without triggering extensive resource reallocations between projects. We provide evidence supporting this claim in Appendix D.

While we identify the problem of resource reallocation externalities in project networks, our analysis also provides insights into tackling this issue. First, we note that contractual incentives that reward or penalize performance on a single project can exacerbate reallocation externalities. We conjecture that “network-based project incentives” that evaluate participants based on their performance across a project network (and not just a single project) could mitigate inefficient resource reallocations.

Furthermore, researchers have recommended “reference class forecasting” as a way of improving project planning and performance (see, for example, Flyvbjerg 2006). Reference class forecasting seeks to predict a project’s budget and schedule based on similar past projects. The results of our study add another dimension to this recommendation: the role of the project network. Specifically, we conjecture that current reference class forecasting methods may be improved if the appropriate reference class included information about concurrent projects in the participants’ networks. Doing so would require participants (such as agencies and contractors) to have greater visibility into their project networks i.e., to which projects they are indirectly connected. This paper, therefore, highlights the value of greater transparency in project networks (similar to the benefits of supply network transparency that has been highlighted in recent literature – see, for example, Kraft et al. 2023).

We acknowledge that network transparency and coordination is more feasible in public project networks, due to government sponsored open data initiatives such as USAspending.gov. We conjecture that our results could enable agency officers to intervene so as to minimize the negative impacts of reallocation externalities. For example, based on the moderating factors we have identified, agency officers may be able to modify the timeline and workload for certain projects if they anticipate a disruption elsewhere in their network.

Projects are routinely delayed and perform poorly. The Project Management Institute estimates that 48% of the projects globally are behind schedule, and $1 million is wasted every 20 seconds collectively due to poor project management.[18] Our research contributes to an improved understanding of what causes project delays, and what could potentially be done to minimize their negative impact. We hope that this work will spur more research on project networks, and lead to more actionable insights on how to allocate resources more efficiently across multiple interconnected projects.

[18] Source: $1 Million Wasted Every 20 Seconds by Organizations around the World

## References


Autor, David H. 2003. Outsourcing at will: The contribution of unjust dismissal doctrine to the growth of employment outsourcing. *Journal of labor economics* **21**(1) 1–42.

Bajari, P., S. Tadelis. 2001. Incentives versus transaction costs: A theory of procurement contracts. *RAND Journal of Economics* **32**(3) 387–407.

Barrot, Jean-Noël, Julien Sauvagnat. 2016. Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *The Quarterly Journal of Economics* **131**(3) 1543–1592.

Baucells, Manel, Yael Grushka-Cockayne, Woonam Hwang. 2024. Managerial mental accounting and downstream project decisions. *Management Science* .

Beaty, Curtis, David Ellis, Brianne Glover, Bill Stockton, et al. 2016. Assessing the costs attributed to project delay during project pre-construction stages. Tech. rep., Texas. Dept. of Transportation.

Bellamy, Marcus A., Soumen Ghosh, Manpreet Hora. 2014. The influence of supply network structure on firm innovation. *Journal of Operations Management* **32**(6) 357–373.

Bertrand, Marianne, Esther Duflo, Sendhil Mullainathan. 2004. How much should we trust differences-in-differences estimates? *The Quarterly journal of economics* **119**(1) 249–275.

Bordat, C., B. G. McCullouch, S. Labi, K. C. Sinha. 2004. An analysis of cost overruns and time delays of indot projects. Tech. rep., Joint Transportation Research Program.

Callaway, Brantly, Pedro H.C. Sant'Anna. 2021. Difference-in-differences with multiple time periods. *Journal of Econometrics* **225**(2) 200–230.

Calvo, Eduard, Ruomeng Cui, Juan Camilo Serpa. 2019. Oversight and efficiency in public projects: A regression discontinuity analysis. *Management Science* **65**(12) 5651–5675.

Cameron, A Colin, Douglas L Miller. 2015. A practitioner's guide to cluster-robust inference. *Journal of Human Resources* **50**(2) 317–372.

Coviello, D., M. Mariniello. 2014. Publicity requirements in public procurement: Evidence from a regression discontinuity design. *Journal of Public Economics* **109** 76–100.

Coviello, Decio, Andrea Guglielmo, Giancarlo Spagnolo. 2018. The effect of discretion on procurement performance. *Management Science* **64**(2) 715–738.

Flyvbjerg, Bent. 2006. From nobel prize to project management: Getting risks right. *Project Management Journal* **37**(3) 5–15.

Flyvbjerg, Bent. 2009. *Optimism and Misrepresentation in Early Project Development*. Palgrave Macmillan UK, London, 147–168.

Funk, Michele Jonsson, Daniel Westreich, Chris Wiesen, Til Stürmer, M Alan Brookhart, Marie Davidian. 2011. Doubly robust estimation of causal effects. *American journal of epidemiology* **173**(7) 761–767.

Girotra, Karan, Christian Terwiesch, Karl T. Ulrich. 2007. Valuing R&D Projects in a Portfolio: Evidence from the Pharmaceutical Industry. *Management Science* **53**(9) 1452–1466.

Goldsmith-Pinkham, Paul, Guido W. Imbens. 2013. Social networks and the identification of peer effects. *Journal of Business & Economic Statistics* **31**(3) 253–264.

Gopal, A., K. Sivaramakrishnan. 2008. On vendor preferences for contract types in offshore software development: The case of fixed price vs. time and materials contracts. *Information Systems Research* **19**(2) 202–220.

Grushka-Cockayne, Yael. 2014. New York City Department of Parks and Recreation. *Darden Business Publishing* Case QA-0815.

Grushka-Cockayne, Yael. 2020. Use Data to Revolutionize Project Planning. *Harvard Business Review* .

Ho, Daniel E., Kosuke Imai, Gary King, Elizabeth A. Stuart. 2007. Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. *Political Analysis* **15**(3) 199–236.

Hull, Theresa S. 2019. Audit of Cost Increases and Schedule Delays for Military Construction Projects at Joint Region Marianas. Tech. rep., Inspector General, US Department of Defense Washington United States.

Jain, Nitish, Karan Girotra, Serguei Netessine. 2022. Recovering global supply chains from sourcing interruptions: The role of sourcing strategy. *Manufacturing & Service Operations Management* **24**(2) 846–863.

Jordan, William C, Stephen C Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management science* **41**(4) 577–594.

Keogh, Geoffrey, Alan W Evans. 1992. The private and social costs of planning delay. *Urban Studies* **29**(5) 687–699.

Kraft, Tim, Junhao Vincent Yu, Yanchong Zheng, et al. 2023. Supply chain transparency and sustainability. *Foundations and Trends® in Technology, Information and Operations Management* **17**(2) 82–154.

Lewis, Gregory, Patrick Bajari. 2014. Moral Hazard, Incentive Contracts, and Risk: Evidence from Procurement. *The Review of Economic Studies* **81**(3) 1201–1228.

Lovallo, Dan, Daniel Kahneman. 2003. Delusions of Success: How Optimism Undermines Executives' Decisions. *Harvard Business Review* .

McPhie, Neil A.G., Mary M. Rose, Barbara J. Sapin. 2005. Contracting Offiicer Representatives: Managing the Government's Technical Experts to Achieve Positive Contract Outcomes. Tech. rep., U.S. Merit Systems Protection Board.

Monfared, Sam, Andrey Pavlov. 2017. Political risk affects real estate markets. *The Journal of Real Estate Finance and Economics* 1–20.

Office of Inspector General. 2018. The Fort Bliss Hospital Replacement Military Construction Project. Audit Report DODIG-2018-125, Department of Defense.

Osadchiy, Nikolay, Vishal Gaur, Sridhar Seshadri. 2016. Systematic risk in supply chain networks. *Management Science* **62**(6) 1755–1777.

Roy, Dwaipayan, Anant Mishra, Kingshuk K Sinha. 2022. Taxing the taxpayers: An empirical investigation of the drivers of baseline changes in us federal government technology programs. *Manufacturing & Service Operations Management* **24**(1) 370–391.

Schmidt, William, Ananth Raman. 2022. Operational disruptions, firm risk, and control systems. *Manufacturing & Service Operations Management* **24**(1) 411–429.

Serpa, Juan Camilo, Harish Krishnan. 2018. The impact of supply chains on firm-level productivity. *Management Science* **64**(2) 511–532.

Staats, Bradley R., Katherine L. Milkman, Craig R. Fox. 2012. The team scaling fallacy: Underestimating the declining efficiency of larger teams. *Organizational Behavior and Human Decision Processes* **118**(2) 132–142.

Stuart, Elizabeth A, Brian K Lee, Finbarr P Leacy. 2013. Prognostic score–based balance measures can be a useful diagnostic for propensity score methods in comparative effectiveness research. *Journal of clinical epidemiology* **66**(8) S84–S90.

Vartabedian, Ralph. 2021. Years of Delays, Billions in Overruns: The Dismal History of Big Infrastructure. *New York Times*.

Wang, Yixin, Jun Li, Di Wu, Ravi Anupindi. 2021. When ignorance is not bliss: An empirical analysis of subtier supply network structure on firm risk. *Management Science* **67**(4) 2029–2048.

Warren, Patrick L. 2014. Contracting officer workload, incomplete contracting, and contractual terms. *The RAND Journal of Economics* **45**(2) 395–421.

# Appendix A: Robustness analyses

## A.1. Alternative delay metrics

Our models thus far use two delay metrics: $\textit{Delay Days}_{p,t}$ (a continuous variable that measures the length of the delay), and $\textit{Delay Probability}_{p,t}$ (where the dependent variable is an indicator that equals one if there was a reported delay, and is zero otherwise). In this section, we rerun our analysis with two alternative delay measures: $\text{Relative Delay}_{p,t} = 100 \times \frac{\text{Delay Days}_{p,t}}{\text{Initial Duration}_p}$ (following Calvo et al. 2019) and $\text{Percentage Delay}_{p,t} = 100 \times \frac{\text{Delay Days}_{p,t}}{\text{Deadline}_{p,t-1} - \text{Start Date}_p}$. The first measure benchmarks the reported delay days to the baseline project duration, whereas the second measure benchmarks the number of delay days as a function of the current projected execution time (i.e., the expected deadline minus the start date). Table A.1 shows that the results are still significant using these alternative metrics.

**Table A.1 Alternative dependent variables**

| | Agency nodes | | | | | | Contractor nodes | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI | VII | VIII | IX | X | XI | XII |
| *Dep. variable:* | | | | | | | | | | | | |
| Relative delay | 0.78*** | 0.46** | 0.47** | 0.73*** | 0.85*** | 1.41*** | 1.36*** | 1.55*** | 1.27*** | 1.44*** | 1.01*** | 1.85*** |
| | (0.16) | (0.20) | (0.23) | (0.22) | (0.23) | (0.29) | (0.25) | (0.27) | (0.31) | (0.31) | (0.32) | (0.45) |
| Percentage delay | 0.84*** | 0.62*** | 0.59*** | 0.87*** | 0.98*** | 1.46*** | 1.07*** | 1.11*** | 0.88*** | 1.18*** | 0.82*** | 1.58*** |
| | (0.13) | (0.16) | (0.18) | (0.18) | (0.18) | (0.25) | (0.19) | (0.21) | (0.24) | (0.24) | (0.27) | (0.39) |
| Agency FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Contractor FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Task FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| *Matching specification:* | | | | | | | | | | | | |
| Number of bins | 2 | 2 | 3 | 4 | 5 | Sturges | 2 | 2 | 3 | 4 | 5 | Sturges |
| Location | No | State | State | State | State | State | No | State | State | State | State | State |
| Identifiable partitions | 22,219 | 49,299 | 52,508 | 52,461 | 51,412 | 29,884 | 19,158 | 45,183 | 46,756 | 45,944 | 44,674 | 24,006 |
| Observations | 1,118,711 | 706,317 | 583,399 | 491,845 | 432,644 | 139,531 | 1,087,062 | 666,371 | 539,473 | 445,100 | 391,093 | 117,750 |
| $R^2$ | 0.26875 | 0.28503 | 0.29281 | 0.30161 | 0.30566 | 0.35856 | 0.14499 | 0.16632 | 0.17381 | 0.18001 | 0.18277 | 0.20990 |

*Note.* This table presents the estimated coefficients for the spillover effect of a network disruption on delay relative to a project's initial duration, and on delay relative to a project's current duration. Columns I-VI show the effect for projects connected through the agency nodes, and Columns VII-XII show the effect for projects connected through the contractor nodes. Treated and control projects are always matched on two-digit task code, two-digit industry code, and four price categories, and numerical variables (number of bids, initial budget, initial duration, annual revenue, and number of employees) using different levels of coarsening. In Columns I and VII, numerical variables are coarsened into two bins. Columns II-VI and VIII-XII gradually increase the number of bins yielding finer partitions. All specifications include fixed effects for the agency, contractor, and project task. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

## A.2. Alternative matching methodologies

In our main analysis, we employed Coarsened Exact Matching to partition projects. In this section, we supplement our analysis by changing the matching methodology to other popular methodologies. In particular, we re-estimated our main results by using propensity score matching; Lasso-based matching; Tree-based matching; Mahalanobis distance matching; n-caliper matching, with five different calipers ($n = 0.1, 0.2, 0.3, 0.4, 0.5$); and k-nearest-neighbor matching (k-NN), where we let $k = 2, 3, 4, 5$.

Figures A.1 and A.2 show that our results are robust across all methods, both in sign and significance. Note that the figures only report the results for the regression specification that includes all controls and fixed effects. However, the regression results were robust across all other specifications.

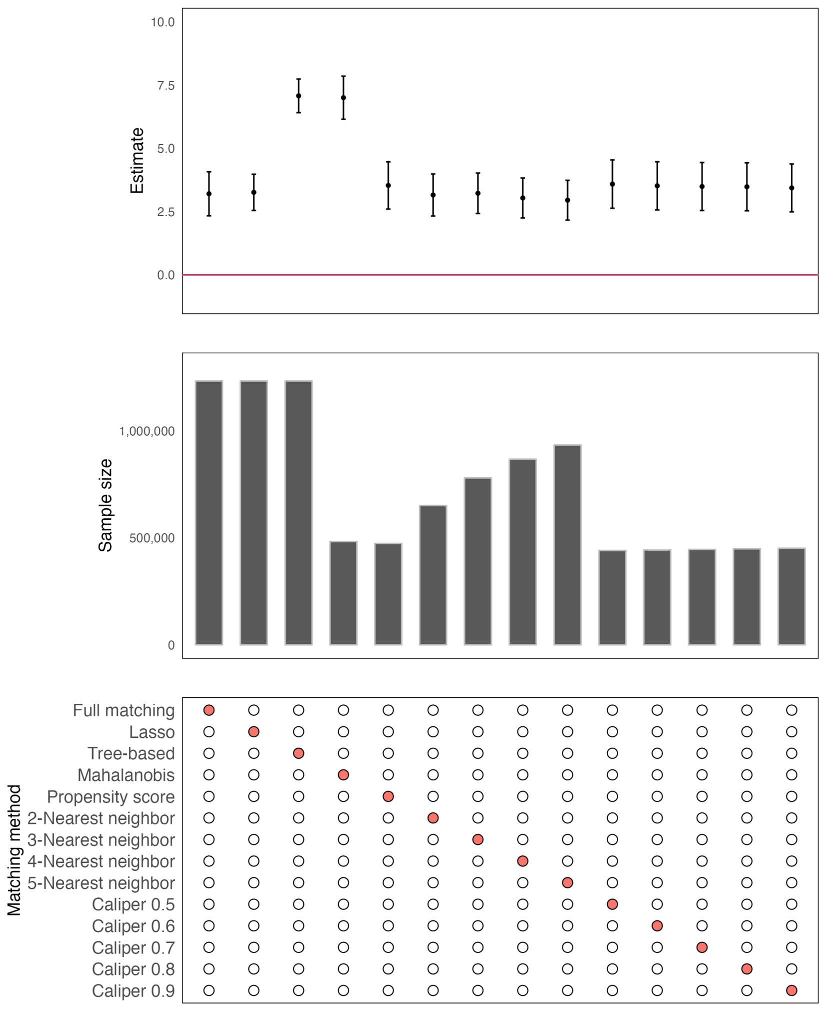


(a) Delay days

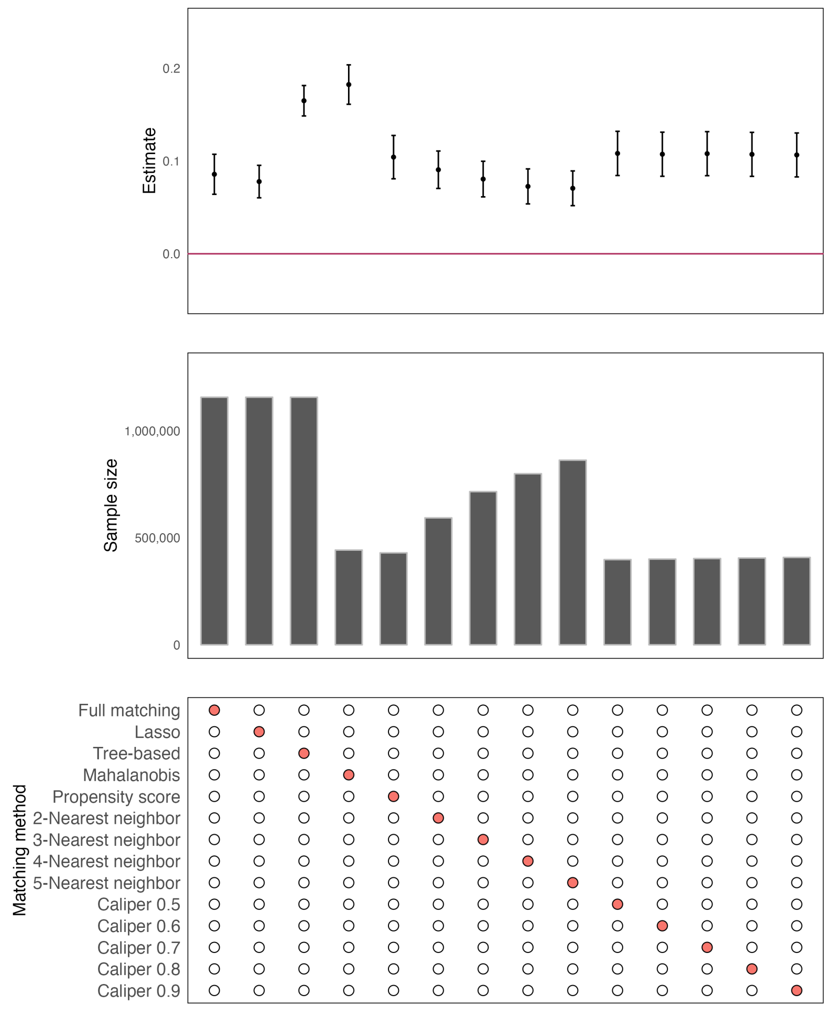


(b) Delay probability

**Figure A.1 Contractor nodes – Spillover effect with alternative matching methodologies.**

These figures show our main results for spillover effect along contractor nodes under different matching techniques. For instance, the first column in each figure shows the results from optimal full matching whereas the last column shows the results from propensity score matching with 0.9 caliper size. The bottom panel illustrates the matching method used (shown in red colored bullets). The middle panel displays the number of observations in each matched sample. The top panel displays the estimated treatment effect and its 95% confidence interval.

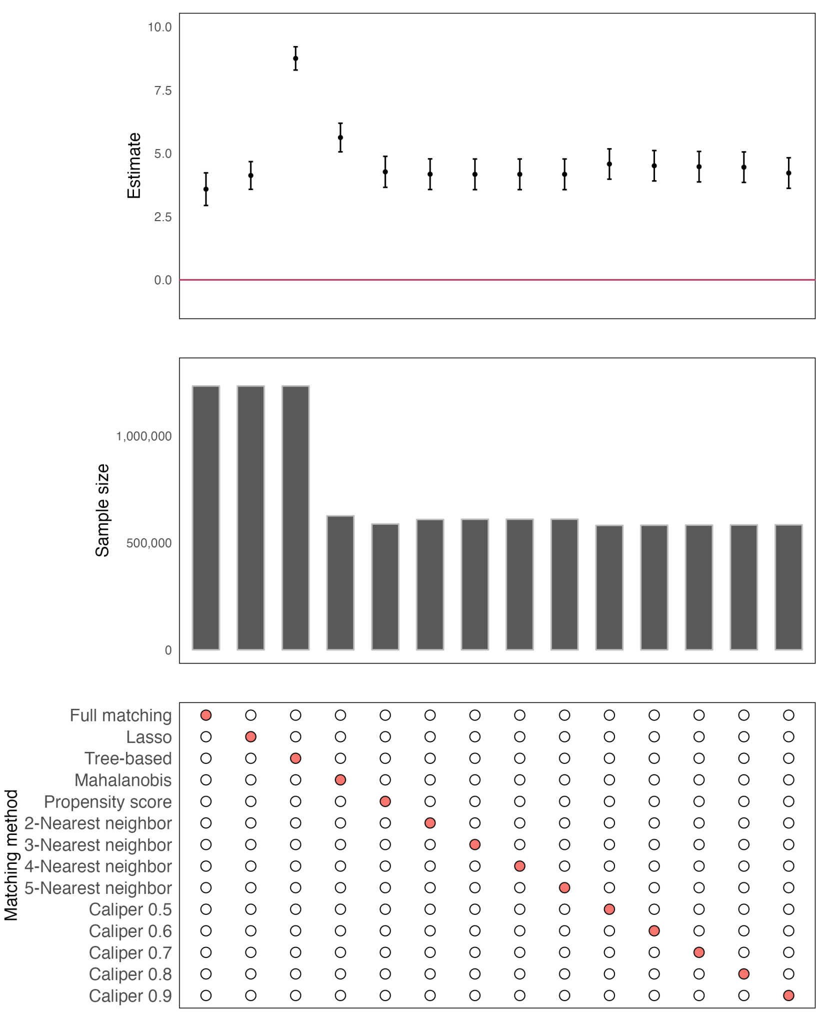


(a) Delay days

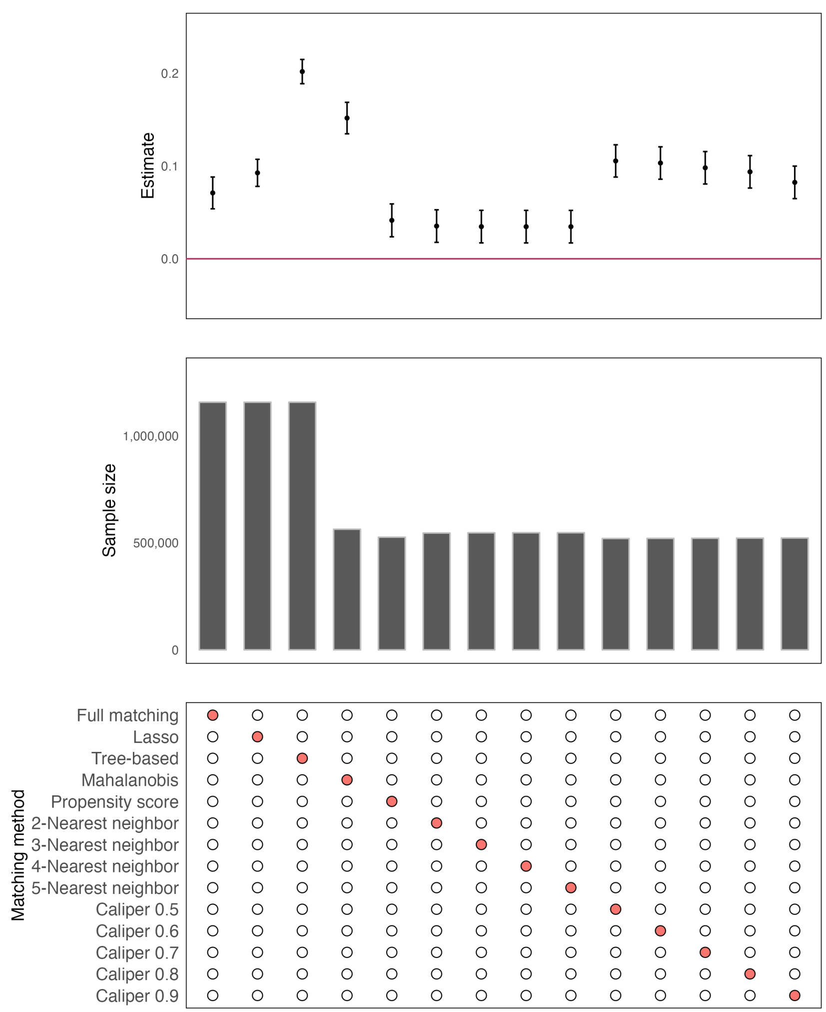


(b) Delay probability

**Figure A.2 Agency nodes – Spillover effect with alternative matching methodologies.**
These figures show our main results for spillover effect along agency nodes under different matching techniques. For instance, the first column in each figure shows the results from optimal full matching whereas the last column shows the results from propensity score matching with 0.9 caliper size. The bottom panel illustrates the matching method used (shown in red colored bullets). The middle panel displays the number of observations in each matched sample. The top panel displays the estimated treatment effect and its 95% confidence interval.

**Figure A.3 Partition granularity**

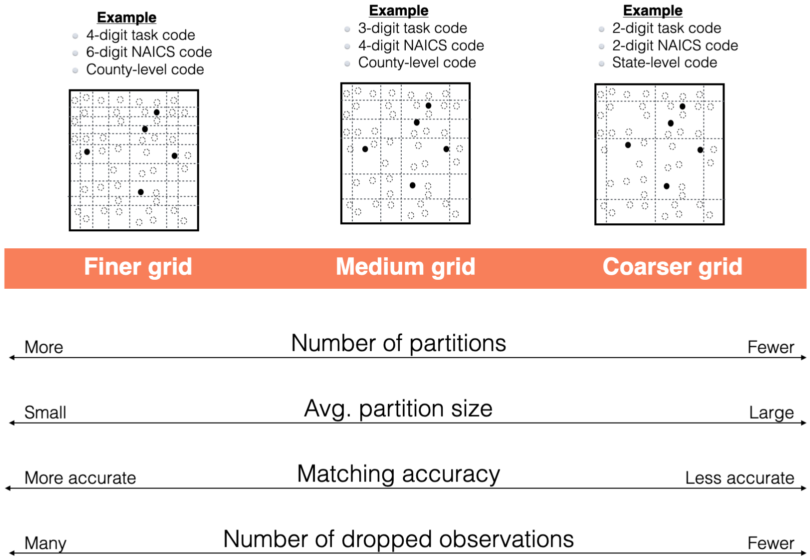


### A.3. Exact matching with coarser partitions

As mentioned in Section 4, our dataset allows us to match projects at different levels of granularity. We could, for example, partition the sample using the 4-digit or 5-digit NAICS code (instead of the 6-digit code), or the 2-digit or 3-digit task-level code (instead of the four-level digit). We could also partition at the county level, instead of requiring it to be at the state level.

Partitioning on higher level variables will result in broader (but fewer) partitions with matches that are slightly less precise. However, it will allow us to utilize most of the observations in our analysis, thereby yielding a more representative estimation (see Figure A.3 for an example). In Table A.2, we present regressions in which we progressively altered the granularity of the partition, transitioning from a very fine to a very coarse partition scheme. As can be seen, the results are largely unaffected by this variation, indicating that our findings are not solely determined by the choice of partition construction.

**Table A.2 Exact matching with coarser partitions**

| | Agency nodes | | | | | | | Contractor nodes | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI | VII | VIII | IX | X | XI | XII | XIII | XIV |
| *Dep. variable:* | | | | | | | | | | | | | | |
| Delay probability | 0.11*** | 0.11*** | 0.12*** | 0.11*** | 0.11*** | 0.09*** | 0.09*** | 0.11*** | 0.11*** | 0.10*** | 0.09*** | 0.09*** | 0.08*** | 0.05*** |
| | (0.03) | (0.03) | (0.02) | (0.02) | (0.02) | (0.02) | (0.01) | (0.02) | (0.02) | (0.02) | (0.02) | (0.02) | (0.02) | (0.01) |
| Delay days | 4.02*** | 4.05*** | 4.50*** | 4.33*** | 4.20*** | 3.79*** | 3.36*** | 5.72*** | 5.56*** | 5.30*** | 5.09*** | 5.15*** | 4.20*** | 2.06*** |
| | (0.86) | (0.83) | (0.74) | (0.72) | (0.65) | (0.58) | (0.47) | (0.97) | (0.94) | (0.87) | (0.85) | (0.80) | (0.72) | (0.64) |
| *Partition granularity:* | | | | | | | | | | | | | | |
| NAICS code | 5-digit | 4-digit | 3-digit | 2-digit | 2-digit | 2-digit | 2-digit | 5-digit | 4-digit | 3-digit | 2-digit | 2-digit | 2-digit | 2-digit |
| Task code | 4-digit | 4-digit | 4-digit | 4-digit | 3-digit | 2-digit | 2-digit | 4-digit | 4-digit | 4-digit | 4-digit | 3-digit | 2-digit | 2-digit |
| Location | County | County | County | County | County | County | State | County | County | County | County | County | County | State |
| Controls | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Agency FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Contractor FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Identifiable partitions | 29,405 | 30,698 | 33,414 | 35,239 | 37,847 | 41,710 | 39,058 | 32,301 | 33,334 | 36,442 | 38,654 | 41,229 | 44,266 | 37,399 |
| Observations | 225,965 | 242,360 | 294,445 | 313,043 | 378,271 | 512,381 | 844,315 | 255,910 | 270,310 | 321,144 | 341,267 | 403,883 | 529,356 | 834,023 |
| $R^2$/Pseudo $R^2$ | 0.31304 | 0.31281 | 0.30370 | 0.30254 | 0.30653 | 0.29567 | 0.29635 | 0.18021 | 0.17703 | 0.18239 | 0.17927 | 0.17129 | 0.17866 | 0.16563 |

*Note.* This table reports the estimated treatment effect of network disruptions on delay probability and delay days using coarser variables in matching. Columns I-VII report the results for agency nodes and Columns VIII-XIV report the results for contractor nodes. All regression specifications include fixed effects for the contractor, agency, task, and price. Standard errors (in parentheses) are robust and clustered by project. Significance levels: $1\%^{***}, 5\%^{**}, 10\%^{*}$.

Table A.3 Statistics for the unadjusted and adjusted distributions of the data

| Covariate | Type | Standardized bias | | Variance ratio | | Kolmogorov-Smirnov | |
|---|---|---|---|---|---|---|---|
| | | Unadjusted | Adjusted | Unadjusted | Adjusted | Unadjusted | Adjusted |
| Annual Revenue – Contractor | Continuous | 0.04 | 0.00 | 1.28 | 1.00 | 0.03 | 0.03 |
| Number of Employees – Contractor | Continuous | 0.04 | 0.00 | 1.25 | 1.00 | 0.04 | 0.03 |
| Number of bids | Continuous | 0.11 | 0.00 | 1.29 | 1.01 | 0.04 | 0.00 |
| Project budget | Continuous | 0.04 | 0.01 | 1.09 | 1.00 | 0.02 | 0.03 |
| Project duration | Continuous | 0.01 | 0.01 | 1.24 | 1.01 | 0.05 | 0.01 |
| Competitively awarded contract | Categorical | 0.01 | 0.01 | · | · | 0.01 | 0.01 |
| Performance-based Incentives | Categorical | 0.02 | 0.01 | · | · | 0.02 | 0.01 |
| Cost contract | Categorical | 0.01 | 0.00 | · | · | 0.01 | 0.00 |
| Fixed price contract | Categorical | 0.03 | 0.00 | · | · | 0.03 | 0.00 |
| Labor hours contract | Categorical | 0.01 | 0.00 | · | · | 0.01 | 0.00 |
| Time and materials contract | Categorical | 0.01 | 0.00 | · | · | 0.01 | 0.00 |

## A.4. Balancing tests of the matched samples

In our main analysis, we matched treated and control observations using coarsened exact matching. In this section, we verify that, after matching, the treatment and control observations have similar distributional properties across the matching covariates. We verify distributional balance across three statistics of the treatment and control samples: (i) the means; (ii) the variances; and (iii) the cumulative distributions (Stuart et al. 2013):

1. *Balance of means.* To measure the balance of the distributions' first moment, we measure the standardized bias for each variable. We calculate the standardized bias of covariate $x$ by measuring the absolute difference of the means, $\left|\mu_x^{treatment} - \mu_x^{control}\right|$, and dividing this difference by the pooled standard deviation. As a rule of thumb, the adjusted standardized differences should be smaller than 0.1 (Stuart et al. 2013). Figure A.4 and Table A.3 illustrate the standardized bias for each covariate, before and after weighting the distributions. These two exhibits show that the distributional means are balanced.
2. *Balance of variances.* To explore the balance of the distributions' variances, we analyze the ratio of the treatment and control groups' variance. By convention, we place the largest variance in the numerator; a ratio of one means that the variances are perfectly balanced and, as a rule of thumb, a ratio below two is acceptable after adjusting the distributions. Table A.3 shows that the distributions' variances satisfy this requirement.
3. *Balance of cumulative distribution.* We explore the balance of the cumulative distribution functions via the Kolmogorov-Smirnov statistic, which measures the maximum distance between the support of these functions. This statistic ranges from zero (perfect balance) to one (full imbalance). By convention, a value below 0.05 is recommended after adjusting. Table A.3 shows that all our adjusted covariates meet this recommendation.

In summary, the balancing table results show that the treatment and control samples are balanced in their means, variances, and cumulative distributions. Figures A.5 and A.6 further illustrate that the density plots are balanced across several covariates for treated and control observations.

## A.5. Placebo tests

**A.5.1. Inter-temporal placebo: Testing for parallel trends.** Our estimates do not come from a prototypical diff-in-diff model, so testing for the parallel trends assumption is not straightforward. This is for

**Figure A.4 Balancing test results**

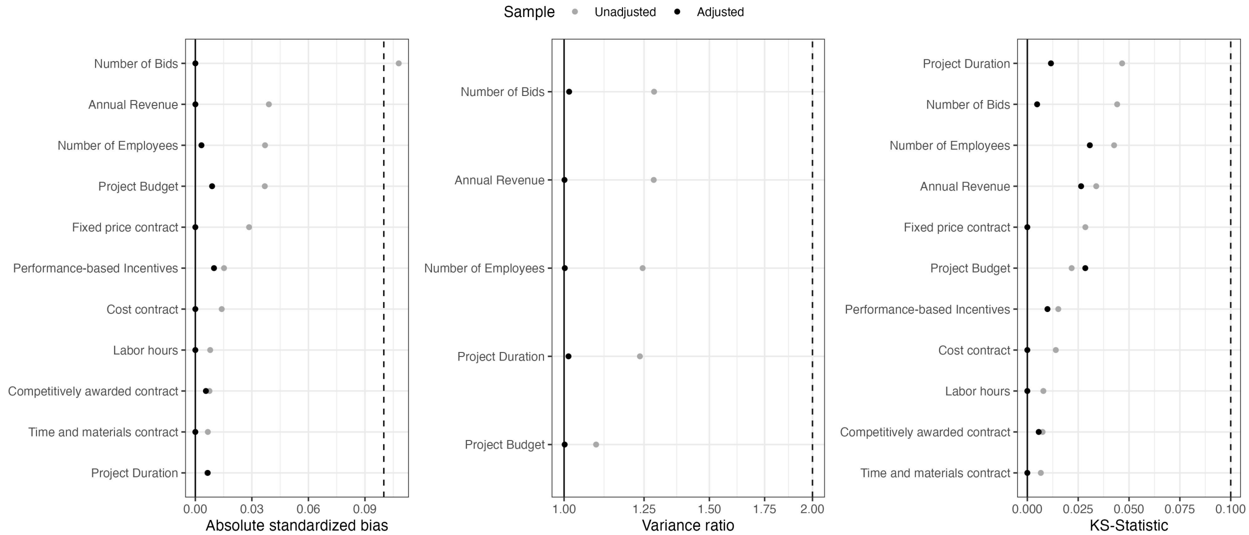


*Note.* This figure illustrates the distributional balance between treated and control observations after matching. The leftmost panel shows the balance of means, the middle panel shows the balance of variances, and the rightmost panel shows the balance of cumulative distribution functions.

**Figure A.5 Density plots – Contractor**

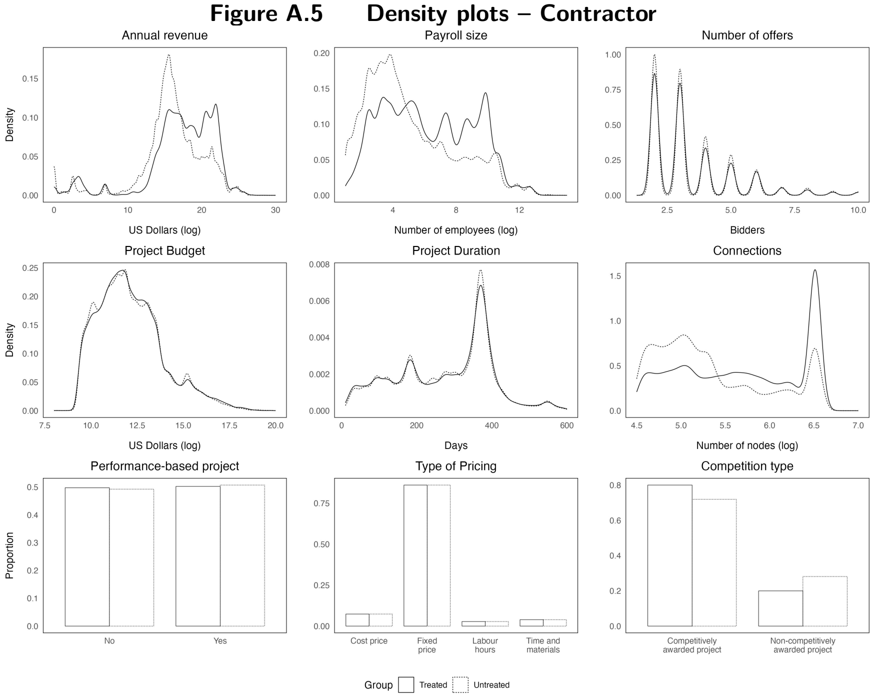


*Note.* This figure illustrates the distribution plots of the treated and control observations for contractor nodes after matching.

three reasons: (i) observations in our sample can be 'treated' repeatedly (a given project can be disrupted at any point in the time series); (ii) the treatment is not coordinated along the cross-section (i.e., disruptions hit different projects in different periods); and (iii) the treatment effect dissipates across time (i.e., the impact of a disruption is not permanent).

**Figure A.6 Density plots — Agency**

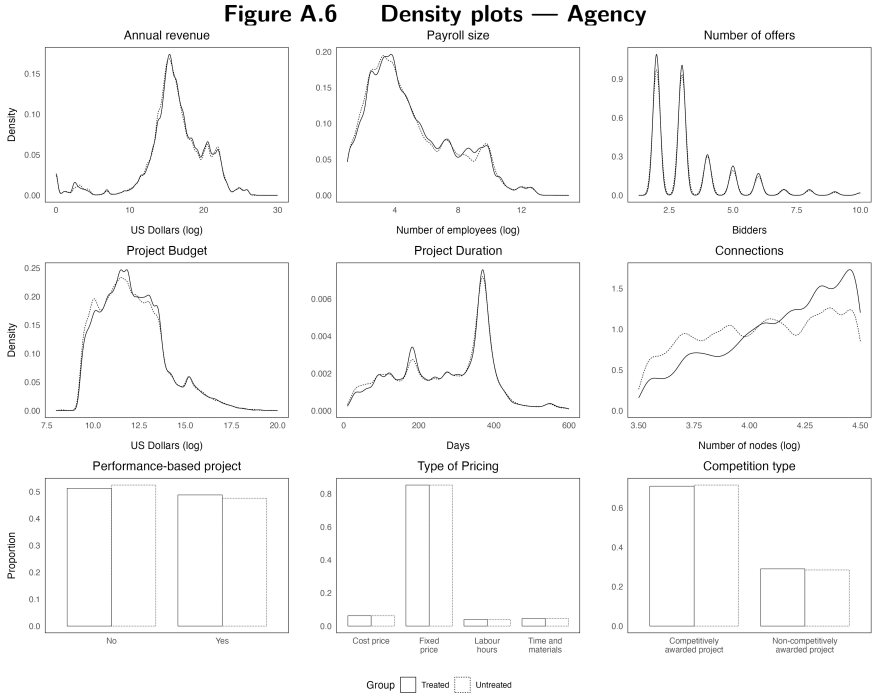


*Note.* This figure illustrates the distribution plots of the treated and control observations for agency nodes after matching.

These three peculiarities mean that any "off-the-shelf" parallel-trends test—like Autor (2003)'s placebo test or a lead-lag test—will not work in our setting. We can, however, still test the fundamental assumption of a diff-in-diff estimator—that absent a treatment, the outcome of the two groups would have been identical over time—by performing an adaptation of Monfared and Pavlov (2017)'s inter-temporal placebo test. In this test, we create 40 synthetic samples, each being identical to the true sample except that we pretend that the disruption occurred during a different period. Twenty of these datasets are called *lagged* synthetic samples, while the remaining twenty are called *lead* synthetic samples, where (i) the $n^{th}$ lagged sample pretends that the treatment date occurred $n$ periods *before* the true event date and (ii) the $n^{th}$ lead sample pretends that the treatment date occurred $n$ periods *after* the true event date. The idea is to shift the treatment date back and forth, one quarter-year period at a time (for a maximum of twenty periods each way), and then determine how this shift would affect the diff-in-diff coefficients. Figure A.7 plots the value of all forty placebo coefficients (i.e., Periods, -20,...,-2,-1,1, 2, ,..., 20) and their corresponding 95% confidence interval. For ease of interpretation, we normalize the value of the true coefficients to 1. According to Monfavred and Pavlov's test, if the coefficients are significant and the parallel-trends assumption is valid, then all lagged placebo regressions would be insignificant (i.e., there are no anticipatory trends). In contrast, the lead placebo regressions would be significant but the effect should dwindle as time passes. Figure A.7 confirms this pattern, by showing a lack of anticipatory effects and a lead treatment effect that lingers for two periods (or six months).

**A.5.2. Cross-sectional placebo test.** We also conduct a cross-sectional placebo test to determine if our results are artifacts of spurious correlations in the data. We create 10,000 samples by randomly assigning disruption occurrences across observations via Bernoulli trials. These 10,000 synthetic samples are

**Figure A.7 Inter-temporal placebo test.**

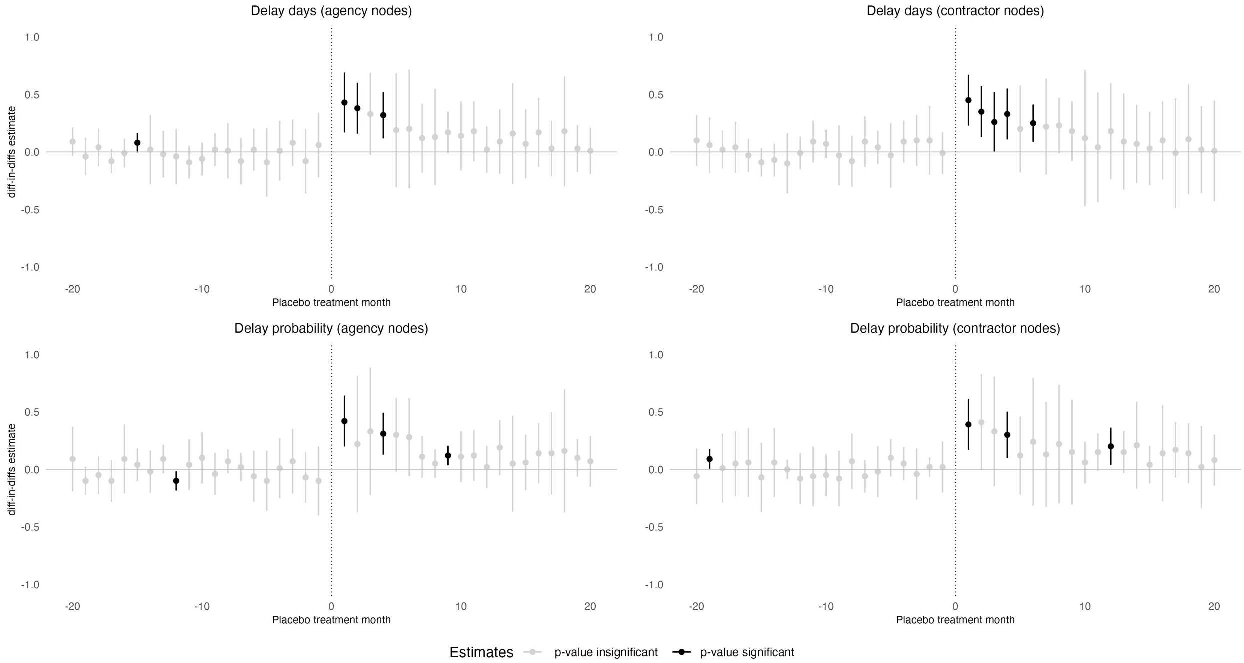


*Note.* For forty synthetic datasets, we re-estimate our models by artificially setting the disruption time period to be different from the "true" event date. Twenty of these samples set the "placebo" treatment date before the true event, and twenty set it to be after the true event. We plot the distribution of p-values for the diff-in-diff coefficient of the forty placebo estimates. When the p-value exceeds 0.1, it is shaded gray; if the draw of a placebo regression yields a p-value below 0.1, it is shaded black. All estimates are drawn from regression models including all fixed effects and control variables.

**Figure A.8 Cross-sectional placebo test.**

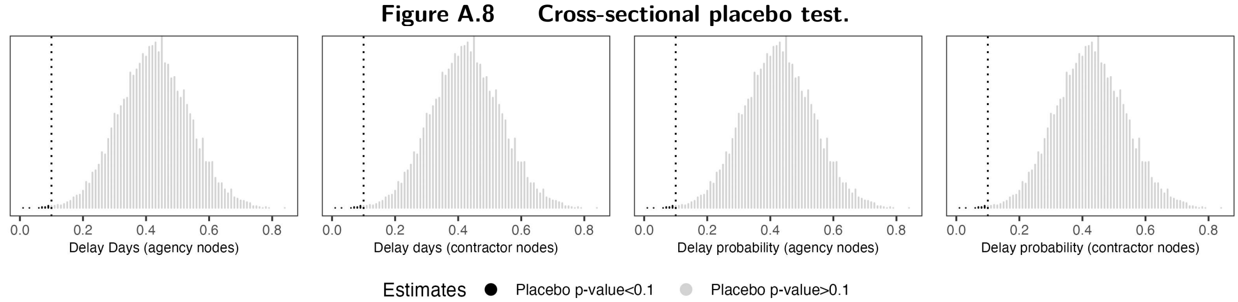


*Note.* For 10,000 synthetic datasets, we re-estimate our model specifications. We plot the distribution of p-values for the diff-in-diff coefficient of the 10,000 placebo estimates. When the p-value exceeds 0.1, it is shaded gray; if the draw of a placebo regression yields a p-value below 0.1, it is shaded black. Each of the four plots contains the distribution using a different dependent variable. All estimates are draw from regressions models including all fixed effects and control variables.

all identical to our true sample except that observations are randomly sorted into the "treated" and "control" groups. Figure A.8 presents the p-values of the 10,000 placebo estimates. If our true estimates were artifacts of spurious correlation, then a substantial number of the synthetic samples would have low p-values. But Figure A.8 shows that a vast majority of the placebo estimates are not even significant at the 10% level.

## A.6. Alternative standard errors

In our main analysis, we accounted for serial correlation by clustering the standard errors at the project level consistent with the recommendation in Bertrand et al. (2004) and Cameron and Miller (2015). In this section, we assess the robustness of our findings to alternate levels of clustering. Table A.4 shows that our estimates are still positive and statistically significant when we cluster the standard errors at the contractor, agency, county, and task level, or use a combination of contractor and agency level clusters.

We also account for potential temporal correlation in the data, and consider alternative measures of standard errors. Table A.5 shows that our results are robust to these measures, including heteroscedasticity and autocorrelation consistent standard errors and heteroscedasticity-consistent standard errors—see Cameron and Miller (2015) for a detailed discussion of these issues.

**Table A.4 Alternative cluster levels**

| | Agency nodes | | | | | Contractor nodes | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI | VII | VIII | IX | X |
| Delay days | $3.74^{***}$ | $3.74^{***}$ | $3.74^{***}$ | $3.74^{***}$ | $3.74^{***}$ | $4.32^{***}$ | $4.32^{***}$ | $4.32^{***}$ | $4.32^{***}$ | $4.32^{***}$ |
| | (0.61) | (0.87) | (0.73) | (0.63) | (0.85) | (1.11) | (1.53) | (1.06) | (1.01) | (1.53) |
| Cluster level | Contractor | Agency | County | Task | Agency + Contractor | Contractor | Agency | County | Task | Agency + Contractor |
| Agency FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Contractor FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Task FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 491,846 | 491,846 | 491,846 | 491,846 | 491,846 | 445,101 | 445,101 | 445,101 | 445,101 | 445,101 |
| $R^2$ | 0.26 | 0.26 | 0.26 | 0.26 | 0.26 | 0.11 | 0.11 | 0.11 | 0.11 | 0.11 |

*Note.* This table reports the estimated treatment effect of network disruptions on delay days by clustering the standard errors (in parentheses) at different levels. Significance levels: $1\%^{***}, 5\%^{**}, 10\%^{*}$.

**Table A.5 Alternative standard errors**

| | Agency nodes | | | | Contractor nodes | | | |
|---|---|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI | VII | VIII |
| Delay days | $3.74^{***}$ | $3.74^{***}$ | $3.74^{***}$ | $3.74^{***}$ | $4.32^{***}$ | $4.32^{***}$ | $4.32^{***}$ | $4.32^{***}$ |
| | (0.57) | (0.86) | (0.53) | (0.51) | (0.81) | (1.30) | (0.78) | (0.75) |
| SE Type | Newey-West | Driscoll-Kraay | Degree of Freedom | Bell-Mcaffrey | Newey-West | Driscoll-Kraay | Degree of Freedom | Bell-Mcaffrey |
| Agency FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Contractor FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Task FE | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 491,846 | 491,846 | 491,846 | 491,846 | 445,101 | 445,101 | 445,101 | 445,101 |
| $R^2$ | 0.26 | 0.26 | 0.26 | 0.26 | 0.11 | 0.11 | 0.11 | 0.11 |

*Note.* This table reports the estimated treatment effect of network disruptions on delay days by considering alternative measures of the standard errors (reported in parentheses). Significance levels: $1\%^{***}, 5\%^{**}, 10\%^{*}$.

## Appendix B: Measuring the economic costs of delay spillovers

The costs of a delay to a project can manifest in many ways: there can be direct, indirect, and opportunity costs. For example, if a traffic jam delays our evening commute home, there are direct costs (cost of gas), indirect costs (additional wear and tear on car) and opportunity costs (less time at home to spend with family or to relax). In other words, delay costs ripple outward – starting with costs that are easy to attribute to the delay and to estimate, and moving to indirect and opportunity costs that are harder, if not impossible, to estimate or even to identify or attribute to the original delay.

This idea, that delay costs ripple outward in the manner described above, is recognized in studies that attempt to estimate the economic significance of delays. For example, the Texas Dept of Transportation (Beaty et al. 2016) developed a model to estimate the costs associated with delay of highway projects.

In our setting, in order to estimate the costs of delay spillovers, we first start by conceptualizing the different categories of delay costs. Keogh and Evans (1992) provide a simple but useful categorization. Project delay costs can be classified as leading to "private costs" and "social costs." Private costs refer to the costs of delays that are incurred by the organizations that are directly involved (e.g., the contractors, agencies, etc.); these costs can be direct (e.g., additional labor costs) or indirect (e.g., overheads incurred while the project is delayed). Social costs refer to the costs of delay that are externalized, i.e., borne by society at large. These costs can include the impact on citizens (e.g., commuters when a road construction project is delayed), etc. Beaty et al. (2016) point out, however, that in public projects, all costs are ultimately borne by the general public.

There are some examples of studies that attempt to estimate social costs of a delay. Using a case study of highway repair projects, Lewis and Bajari (2014) argue that "the daily social cost imposed by the construction would be 175,000 hours. Valuing time at $10 an hour, this implies a social cost of $1.75 million per day."

In our context, it might be impractical, if not impossible, to attempt to quantify the social costs of delays. This is because of the wide variety of projects and users that are part of our data set. Nonetheless, we acknowledge that these costs can be significant. Keogh and Evans (1992) also note that social costs are extremely hard to estimate and that estimates of delay costs are often confined to the estimation of private costs. We follow a similar approach here and restrict our estimation of costs to the direct and indirect cost borne by the contractors. This estimate, of course, would be a lower bound on the true delay costs.

### Measuring direct and indirect costs of project delays

A straightforward approach to estimating the direct costs of a delay is to simply look at the cost overrun (i.e., the actual project cost – the initial estimated project cost). The benefit of this approach is its simplicity. However, cost overruns may incorrectly estimate the direct costs, e.g., if some direct costs cannot be billed to the project.

For projects where labor costs are the major component of direct costs of delay, estimating the additional labor cost due to project delays provides a conservative estimate of delay costs.

Indirect costs are even more challenging to account for and estimate. This is because indirect costs are often overheads that a contractor allocates to its projects. How would a delay affect the allocation and recovery of the additional overhead costs that a contractor may incur?

U.S. Federal Courts have accepted a formula called the Eichleay formula as a valid method to estimate overhead costs that cannot be directly attributed to a particular project but are incurred due to the extended duration of the project caused by the delay. Overhead costs typically include administration costs, utilities, depreciation, taxes, insurance, etc., which are not directly billable to a specific project but are necessary business costs. The Eichleay formula is calculated via a three-step procedure:

- Step 1: Total Project Billings ÷ Total Company Billings × Total Home-Office Overhead During Actual Contract Period = Overhead Allocated to Project
- Step 2: Overhead Allocated to Project ÷ Actual Days of Project Performance (including delay) = Rate of Overhead Allocated to Project Per Day
- Step 3: Rate of Overhead Allocated to Project Per Day × Number of Days Delayed = Amount of Overhead Allocated to Project due to delay.

**Using data to estimate the economic costs of delay spillover**

We estimate the direct and indirect costs of a delay using three methods commonly used by the project managers and the courts to assess the damages caused by a delay.

**1. Eichleay formula:** For the first estimate, we use the labor cost of delay to estimate the direct cost and we estimate the indirect cost using the Eichleay formula.

The labor cost of delay (i.e., direct costs of delay) was estimated by retrieving the average hourly wage in every industry from the Bureau of Labor Statistics, and the average overhead cost and company billings from filing records. We, then, matched this with project records to estimate the cost of labor in the project's industry. We computed the labor cost for a project delay as follows:

$$\text{Direct cost of labor} = 8 \times \text{Hourly wage per worker} \times \text{Days of delay} \times \text{Employees per project}$$

We obtain the overhead allocated to a project by re-writing the first two steps of Eichleay formula:

$$\text{Daily overhead allocated to project} = \frac{\text{Total project cost}}{\text{Actual project duration}} \times \frac{\text{Total overhead}}{\text{Total company billings}}$$

Once the overhead allocated to a project is defined, the indirect costs are calculated as follows.

$$\text{Indirect costs} = \text{Daily overhead allocated to project} \times \text{Delay days}$$

**2. Modified Eichleay formula:** For the second estimate, we again use the labor cost of delay to estimate the direct cost and we estimate the indirect cost using a Modified Eichleay formula.

The modified Eichleay formula uses a slightly different formula for the estimating the overhead allocated to a project. The formula is as follows:

$$\text{Daily overhead allocated to project} = \frac{\text{Total project cost}}{\text{Initial project duration}} \times \frac{\text{Total overhead}}{\text{Total company billings}}$$

Note that the denominator in the modified Eichleay formula is the initial project duration and not the actual project duration. Again, the indirect costs are calculated by multiplying the daily overhead allocated to a project with delay days.

**3. CD3 formula:** A third approach we use to estimate delay costs is called the "cost of delay divided by duration" or "CD3". This is a commonly used approach by contractors, and it does not require much data. The CD3 approaches estimates the sum of direct and indirect costs as follows.

$$\text{Direct Cost} + \text{Indirect cost} = \frac{\text{Total project cost}}{\text{Initial project duration}} \times \text{Delay days}$$

Once we have been able to estimate the economic costs of a project delay, we can quantify the economic costs of a delay spillover. In order to do this, we first measure the size of spillovers per delay day. We then multiply the per-day total cost by this magnitude. Put differently:

$$\text{Economic costs of delay spillover} = \text{Cost of delay per day} \times \text{Spillover magnitude}$$

We adapt the above methods using our data. In particular, to estimate the direct cost of a delay, we retrieved the average hourly wage in every industry from the Bureau of Labor Statistics, and proxy for the average overhead cost and company billings using filing records from Compustat. We, then, matched this with project records to estimate the cost of labor in the project's industry. Note that this approach can only provide a lower bound on the direct cost of delays as it does not account for some inputs. However, our intent here is to illustrate a methodology which can be used to estimate delay costs, and a lower bound will at least provide some idea about the economic impact of a delay.

Our estimates of spillover costs using the above three methods are shown in Table B.1. Further, Figure B.1 illustrates the spillover cost for select task types in our data.

**Table B.1 Expected cost of delay spillover per period**

| | Eichleay Formula | | | Modified Eichleay Formula | | | $CD^3$ |
|---|---|---|---|---|---|---|---|
| | Direct costs | Indirect costs | Total Cost | Direct costs | Indirect costs | Total Cost | Total Cost |
| Treated projects | \$33,753.11 | \$9,461.42 | \$43,214.53 | \$33,753.11 | \$24,065.65 | \$57,818.76 | \$70,778.47 |
| Control projects | \$30,327.72 | \$7,959.81 | \$38,287.53 | \$30,327.72 | \$20,344.41 | \$50,672.13 | \$60,671.65 |
| **Spillover cost** | | | **\$4,927** | | | **\$7,146.63** | **\$10,106.82** |

*Note.* This table shows the estimated cost of delay spillovers using three different methods: Eichleay formula, Modified Eichleay formula, and $CD^3$ approach.

**Figure B.1 Illustration of spillover costs by task type**

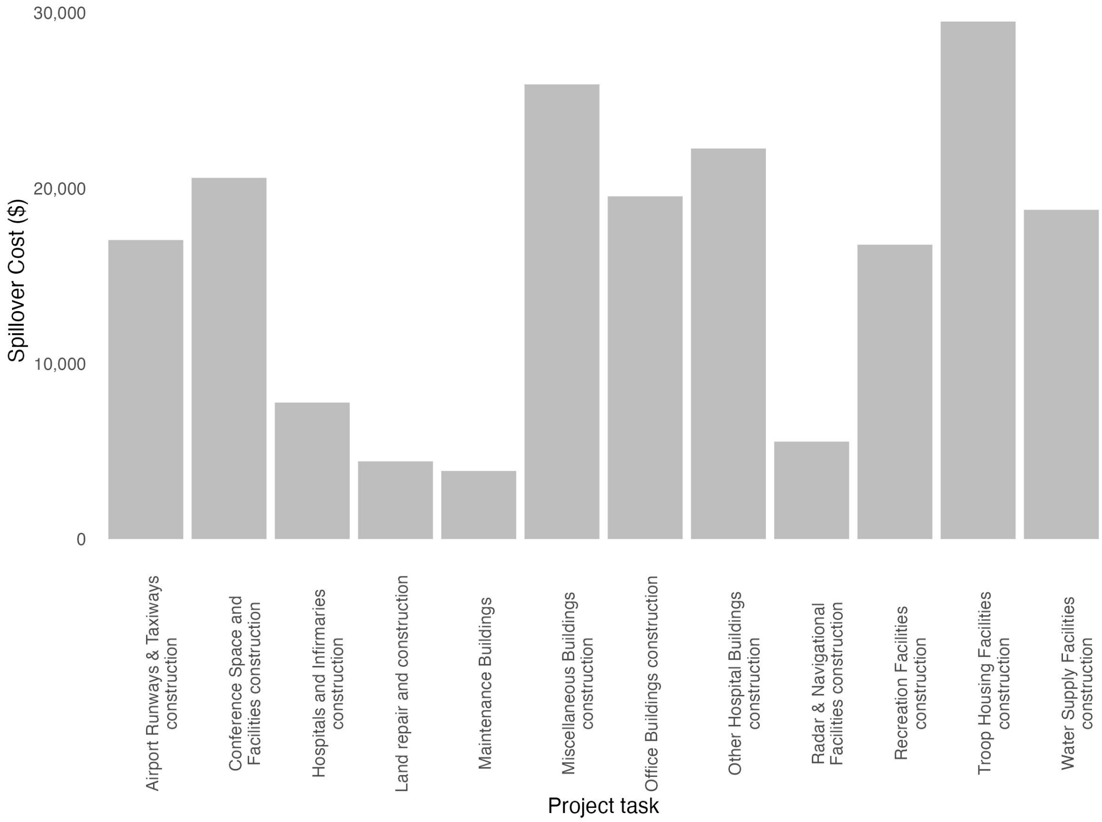


*Note.* This figure illustrates the estimated cost of delay spillovers for select project tasks.

## Appendix C: Non-weather disruptions

The main analysis focused on weather disruptions, given that doing so allows us to obtain cleaner results. In this section, we examine the impact of other types of disruptions in this analysis. Unlike weather-related disruptions, which can be readily identified, categorizing other disruptions is not as straightforward. Therefore, we had to establish a taxonomy. Recall that we observe a text description wherein the agency officer describes the event corresponding to a disruption. We examined these text descriptions, both manually and through word-frequency searches, to discern patterns in the disruption events. From this exercise, we identified the following frequently-occurring root causes.

1. **Bureaucratic issues:** Issues related to bureaucratic hurdles, paperwork and invoice issues, legal requirements, or policy changes that affect the project. This includes changes in building codes, permits, environmental regulations, or zoning restrictions that necessitate modifications to the project plan.

   *Example: Modification to extend completion date to 30 June 2015 due to delays resolving regulatory comments on Site 7 sampling plan.*
2. **Labor/personnel issues:** Issues related to labor strikes, shortage of skilled workers, or a sudden departure of key personnel that impacts the project's progress.

   *Example: Add additional test fits, redesign, documentation, and construction management required to execute the revised space plan design and additional services related to the original scope of work due to labor dispute.*
3. **Worksite issues:** Issues related to a malfunction or failure of machinery or equipment used in the project. This includes, for example, crane malfunction, discovery of asbestos on construction site, a generator failure, or a software glitch that affects project operations.

   *Example: Date change due to asbestos- and arsenic- related issues at the building.*
4. **Supply chain issues:** Issues related to shipping problems, material stockouts, or quality issues with the supplied materials.

   *Example: The purpose of this modification is to a grant time extension due to late delivery and unexpected modifications the crane's runway due to the location of an electrical panel. A time extension has been granted.*

We next categorize the text descriptions from the modification records into the four types of disruptions identified above. Given the size of our dataset, it is infeasible to tag all disruptions manually. Therefore, we used string matching to do a raw categorization of the data using commonly appearing strings. For instance, whenever we identified "delay due to labor," or "staff" issues, we proceeded to categorize the disruption as a labor disruption. Similarly, when we saw terms like "asbestos," "electrical", "code violation," "machine," we proceeded to tag the disruption as a worksite problem. When we saw terms like "paperwork", "permit", "administrative" issues, we classified it as a bureaucratic disruption.[19]

[19] To further validate our findings, we also conducted a secondary test via an Amazon Mechanical Turk (MTurk) task. Specifically, we asked Mturk participants to read a subsample of modification records and classify the reported disruptions in them. Although the results from this analysis were qualitatively similar, we do not report them due to the small sample size and potential for labeling inaccuracies.

**Table C.1 Diff-in-diff effect by disruption type**

| | Bureaucratic disruptions | | Worksite disruptions | | Labor disruptions | | Supply chain disruptions | |
|---|---|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI | VII | VIII |
| Delay days | 3.03*** | 3.74*** | 6.14*** | 6.51*** | -1.38 | -1.26 | 6.77*** | 6.93*** |
| | (1.15) | (1.14) | (1.29) | (1.28) | (6.89) | (6.77) | (1.45) | (1.40) |
| Controls | No | Yes | No | Yes | No | Yes | No | Yes |
| Fixed effects | Yes | Yes | Yes | Yes | Yes | Yes | Yes | Yes |
| Observations | 451,890 | 451,890 | 339,605 | 339,605 | 97,363 | 97,363 | 400,995 | 400,995 |
| $R^2$ | 0.26 | 0.28 | 0.37 | 0.38 | 0.55 | 0.56 | 0.31 | 0.33 |

*Note.* This table presents the estimated coefficients for diff-diff effect on delay days by the type of disruption. Each specification includes agency, contractor, task, and county fixed effects, and includes or excludes project level controls. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

We re-run our analysis by matching treatment and control projects using these new definitions of disruptions, and performing a diff-in-diff regression that retrieves the treatment effect by root cause—Table C.1 shows the results.

We find evidence that bureaucratic, worksite, and supply chain disruptions spill over to delay other projects. Labor disruptions, on the other hand, do not cause significant spillover effect. This finding is also consistent with our analysis that shows contractors reallocate resources from machine-intensive tasks, as opposed to labor intensive ones. While these results do provide some evidence that disruptions caused by non-weather related issues also propagate in the project network, we caution against causal interpretation of these findings due to potential for endogeneity, reporting biases, and omissions in the data (as previously discussed in Section 3).

## Appendix D: Additional results and summary statistics

### D.1. Which tasks are more prone to reallocation externalities?

The propensity to experience a weather shock and the ease of resource reallocation may depend on the task being performed. This means that the effects studied in this paper could be highly contingent on the nature of the work. But which types of tasks are more prone to experiencing disruption externalities? To investigate this issue, our initial approach was to disaggregate the analysis and determine how the treatment effect varied as a function of the task category. That is, we obtained a treatment effect for every task category, using the 3-digit task code.

Figure D.1 displays the treatment effect by task type via a tree map, while Table D. 1 presents the five types of tasks with the largest treatment effect and the five types of tasks with the most negative treatment effect. This analysis shows allows us to see that the treatment effect is positive across most task categories, meaning that the effect is fairly generalized across the network instead of being secluded to a specific subset of tasks. From this disaggregated analysis, however, we cannot identify a clear pattern regarding which types of tasks are more prone to delays.

**Figure D.1 Treatment effect by task type of the disrupted project**

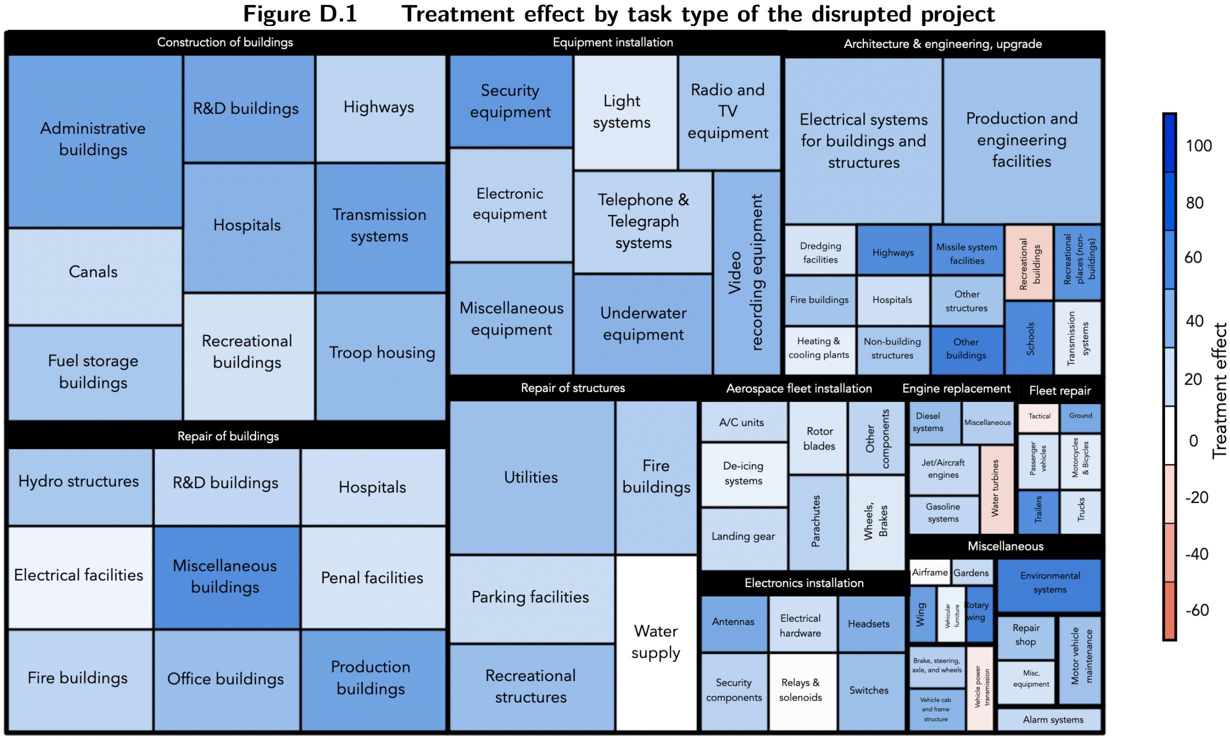


*Note.* This figure displays how the average treatment effect varies by task category of the disrupted project.

### D.2. Contractor size and reallocation externalities

Our results show that localized disruptions at one project propagate to delay other projects that are connected in the project network due to self-interested resource reallocation by project participants. This prompts the question of how operational concepts such as slack time or safety capacity affect these reallocation

**Table D. 1** **Diff-in-Diff estimates by task type: top-5 and bottom-5 task codes**

**Top-5 (Three-digit task codes)**

| Code | Task category | Treatment effect |
|---|---|---|
| C12 | Architecture & Engineering – Non-building structures | 88.44 |
| 254 | Vehicular Equipment Components – Furniture and accessories | 84.73 |
| C1C | Architecture & Engineering – Schools | 81.34 |
| C1L | Architecture & Engineering – Highways, roads, streets, bridges, and railways | 78.65 |
| Z19 | Maintenance and Repair of Buildings – Miscellaneous buildings | 77.56 |

**Bottom-5 (Three-digit task codes)**

| Code | Task category | Treatment effect |
|---|---|---|
| C1F | Architecture & Engineering – Recreational buildings | -26.05 |
| 283 | Engines & Turbines – Water turbines and water wheels | -22.64 |
| Z21 | Maintenance and Repair of Non-buildings – Dams | -16.70 |
| 155 | Aerospace Craft And Structural Components – Space vehicles | -11.59 |
| 230 | Motor Vehicles, Cycles, Trailers – Ground effect vehicles | -11.31 |

**Table D. 2** **Diff-in-diff estimates as a function of contractor size**

| | Treatment effect (in days) | | | | | |
|---|---|---|---|---|---|---|
| | I | II | III | IV | V | VI |
| Ln(1+Annual Revenue) | -1.32*** | -1.43*** | -1.30*** | | | |
| | (0.16) | (0.18) | (0.18) | | | |
| Ln(1+Employees) | | | | -2.29*** | -2.58*** | -2.65*** |
| | | | | (0.37) | (0.38) | (0.40) |
| Controls | Yes | Yes | Yes | Yes | Yes | Yes |
| Task FE | Yes | Yes | Yes | Yes | Yes | Yes |
| Industry FE | No | Yes | Yes | No | Yes | Yes |
| County FE | No | No | Yes | No | No | Yes |
| Num treated | 25,875 | 25,870 | 25,870 | 25,875 | 25,870 | 25,870 |
| $R^2$ | 0.12 | 0.15 | 0.18 | 0.12 | 0.15 | 0.18 |

*Note.* This table presents the estimated coefficients for the treatment effect as a function of the contractor size. Each specification controls for the project's initial budget and duration, and number of offers received; and includes or excludes fixed effects. The standard errors (reported in parentheses) are robust and clustered at the project level. Significance levels: 10% (*), 5% (**), and 1% (***)

externalities. One could conjecture, for example, that larger contractors have a bigger resource pool and can invest in operational buffers. Therefore, projects operated by large contractors may be able to absorb the effect of a localised disruption without negative reallocation externalities. Since we do not observe the resources that contractors allocate to a project, we cannot directly observe the slack time or the level of safety capacity available. However, we can proxy for operational buffers by using two measures of contractor size in our data—the annual revenue and the number of employees. We examine how the treatment effect varies with respect to these two variables. Table D. 2 shows that, indeed, the reallocation externalities reduce as the size of the contractor increases.

### D.3. Summary statistics

In this section, we present the summary statistics for the projects and weather data used in our analyses (see Figure D.2 and Tables D. 3 and D. 4), as well as on the partitioning process (in Table D. 5).

**Figure D.2 Map of weather events by type**

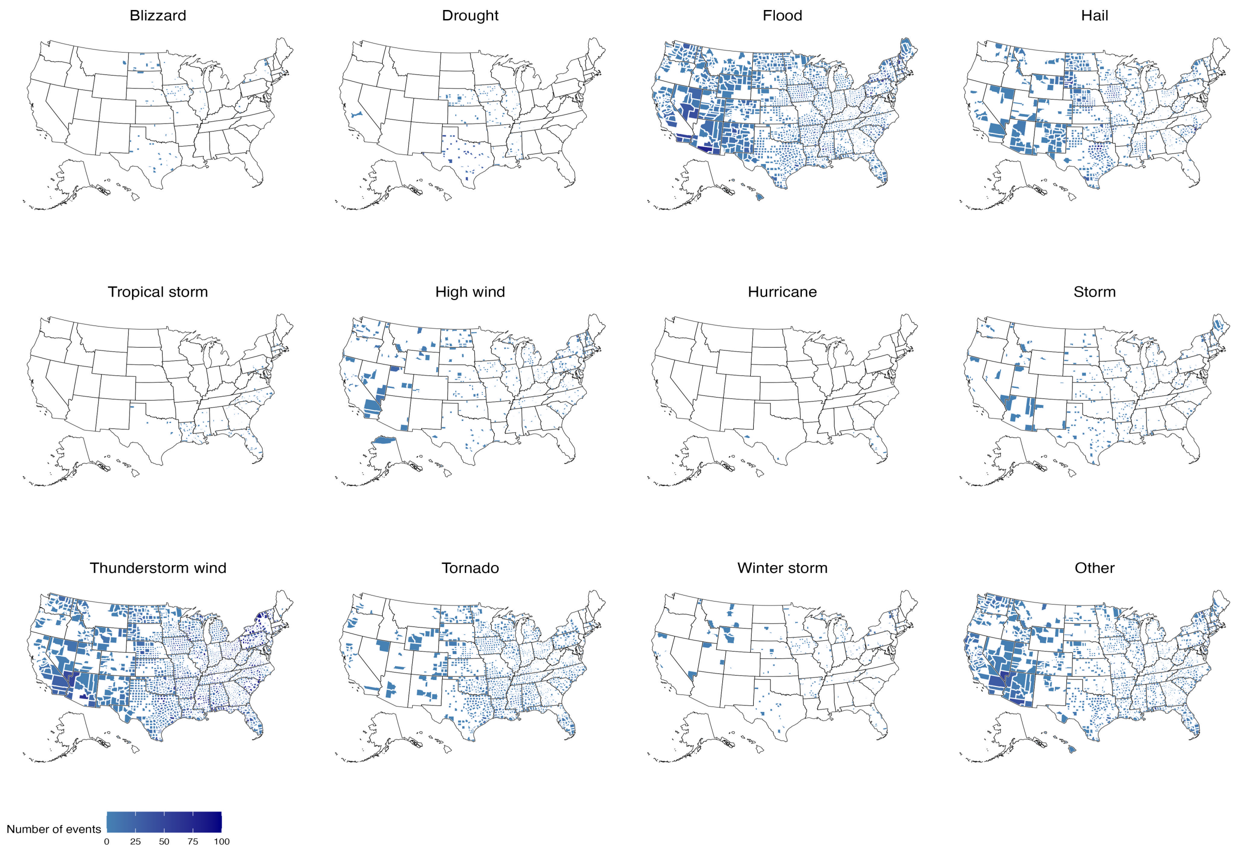


*Note.* Number of reported weather-related events, by county and type. Darker shades represent more events.

**Table D. 3 Summary statistics: Project characteristics (drawn from the project records dataset) before matching**

| Variable | Type | Unit | Mean | St Dev |
|---|---|---|---|---|
| Project budget | Continuous | Dollars (100,000s) | 5.27 | 11.18 |
| Project duration | Continuous | Days | 286.54 | 272.68 |
| Number of bids | Count | Bids | 3.45 | 3.71 |
| Number of employees | Count | Employee (100s) | 147.77 | 355.48 |
| Annual revenue | Continuous | Dollars (in millions) | 5822.71 | 15881.17 |
| Competitively awarded project | Categorical | {0,1} | 0.76 | 0.43 |
| Cost-plus contract | Categorical | {0,1} | 0.08 | 0.26 |
| Fixed price contract | Categorical | {0,1} | 0.88 | 0.33 |
| Labor hours contract | Categorical | {0,1} | 0.02 | 0.15 |
| Time and Materials contract | Categorical | {0,1} | 0.02 | 0.15 |
| Number of projects | 2,484,188 | | | |
| Number of contractors | 124,026 | | | |
| Number of agency offices | 3,559 | | | |
| Number of tasks | 1,188 | | | |
| Number of counties | 2,970 | | | |
| Number of agencies | 65 | | | |
| Number of sub-agencies | 173 | | | |
| Avg. delay caused by a weather disruption | 48.71 days | | | |
| Sample Timespan | Jan 01, 2011 to Sep 30, 2015 | | | |

**Table D. 4 Summary statistics: Weather records**

| Variable | Type | Unit | Mean | St Dev |
|---|---|---|---|---|
| Number of deaths | Count | Person | 0.010 | 0.330 |
| Number of injuries | Count | Person | 0.050 | 2.930 |
| Property damage | Continuous | Dollars (in millions) | 0.499 | 44.792 |
| Event duration | Continuous | Days | 1.840 | 6.640 |
| Wildfire | Categorical | $\{0,1\}$ | 0.006 | 0.080 |
| Flash flood | Categorical | $\{0,1\}$ | 0.060 | 0.240 |
| Tornado | Categorical | $\{0,1\}$ | 0.020 | 0.140 |
| Ice storm | Categorical | $\{0,1\}$ | 0.004 | 0.070 |
| Drought | Categorical | $\{0,1\}$ | 0.060 | 0.230 |
| Number of severe weather events (damage $\geq$\$500,000) | | 5,808 | | |
| Time-span | | Jan 1, 2011 to Dec 31, 2018 | | |

**Table D. 5 Summary statistics: Storm-hit project's network connections**

| Number of Connections | Percentage of projects |
|---|---|
| 0 | 0.57 |
| 1 | 0.12 |
| 2 | 0.06 |
| 3 | 0.04 |
| 4 | 0.03 |
| 5 | 0.03 |
| 6 | 0.02 |
| 7 | 0.02 |
| 8 | 0.02 |
| 9 | 0.01 |
| $\geq$10 | 0.14 |

*Note.* This table shows the number of network connections that a given storm-hit project concurrently had at a given time (not the sum of connections over the entire project duration).

**Table D. 6 Partitions in Step 1**

| Partition size | Number of partitions |
|---|---|
| 1 | 272,403 |
| 2 | 58,876 |
| 3 | 26,327 |
| 4 | 14,873 |
| 5 | 9,879 |
| 6 | 6,875 |
| 7 | 5,026 |
| 8 | 3,924 |
| 9 | 3,116 |
| $\geq$10 | 27,788 |

**Table D. 7 Period specific partitions in Step 3**

| Partition size | Number of partitions | % Treated |
|---|---|---|
| 1 | 986,666 | 19.85 |
| 2 | 175,822 | 21.17 |
| 3 | 70,199 | 21.90 |
| 4 | 37,204 | 22.40 |
| 5 | 22,758 | 22.86 |
| 6 | 15,527 | 22.77 |
| 7 | 10,848 | 22.85 |
| 8 | 8,197 | 23.07 |
| 9 | 6,194 | 24.47 |
| $\geq$ 10 | 41,489 | 29.48 |

**Table D. 8 Identifiable partitions in Step 4**

| Partition size | Number of partitions |
|---|---|
| 1 | 0 |
| 2 | 17,644 |
| 3 | 10,581 |
| 4 | 7,010 |
| 5 | 4,899 |
| 6 | 3,703 |
| 7 | 2,862 |
| 8 | 2,250 |
| 9 | 1,910 |
| 10 | 16,424 |