The Propagation of FDA Regulatory Enforcement Through Drug Supply Chains

Karna Malaviya*
University of California, San Diego

Abstract

The Food and Drug Administration (FDA) upholds the safety of US drug supply, but persistent shortages and concentration in manufacturing have raised concerns that the agency's regulatory enforcement actions may disrupt access to essential medicines. I study this trade-off between access and safety by estimating the causal effects of FDA enforcement at upstream drug manufacturing facilities on downstream hospitals and the patients they serve. Linking Medicare claims to newly assembled data on the origin facilities of sterile injectable opioid drugs, I find that enforcement causes substantial production disruptions, reducing volume from the implicated facility by one-third. Yet, responses by other elements of the supply chain provide resilience: hospitals draw on multi-sourced procurement networks to rapidly substitute to alternative suppliers, such that more-exposed hospitals experience no detectable declines in drug supply relative to less-exposed hospitals. Even in high-risk conditions under which supply does decline—among rural hospitals and when enforcement occurs during critical shortages—the disruption does not compromise patient health or quality of care. These findings suggest that, despite concerns about fragility in this setting, FDA enforcement in concentrated supply chains shields patients from drugs produced at non-compliant facilities while maintaining access and quality of care.

*PhD Candidate, Department of Economics. Email: kmalaviya@ucsd.edu. I am grateful to Itzik Fadlon, Jeff Clemens, and Julie Cullen for their guidance and encouragement throughout this project. I also thank Titan Alon, Prashant Bharadwaj, Gaurav Khanna, Fabian Eckert, and participants of the UCSD graduate student lunch workshop for their helpful suggestions. Finally, I thank the Washington Post and Charleston Gazette for their efforts in publicly releasing the DEA ARCOS data and Mohan Ramanujan, Sara Schwarz, Yashna Nandan, and Itzik Fadlon for assistance in securing a data use agreement for Medicare claims data. All errors are my own.

1 Introduction

The concentration and fragility of modern supply chains sharpen the canonical trade-off regulators face between product access and safety (Kwon et al., 2024; Liu et al., 2024; Baldwin et al., 2023). Enforcement actions intended to protect consumers, when directed at high market share firms, may risk production disruptions that compromise access to critical products. Indeed, recent experience illustrates that nationwide shortages can be triggered by an outage at just one single manufacturing facility: infant formula in 2022 (FTC, 2024); hospital IV bags in 2024 and 2017 (Aguero and Allen, 2024); and car brake and fuel line parts in 2012 (Holbrook, 2013), among others. These shocks, which arose from operational failures and natural disasters, underscore a vulnerability of concentrated production: isolated upstream disruptions can have broad downstream consequences.

Understanding how enforcement propagates downstream is therefore essential for evaluating the implications of supply chain concentration for regulators. Unlike natural disasters or operational failures, enforcement actions are a deliberate and routine feature of normally-functioning markets. If enforcement has adverse downstream consequences, regulators may face a dilemma similar to "too-big-to-fail" (Strahan, 2013), whereby they may reduce stringency at important firms to manage shortage risk. For antitrust authorities, this would also highlight tension between efficiency gains from economies of scale and supply chain vulnerability in high fixed cost industries.

This paper examines these implications by providing, to my knowledge, the first causal evidence of the downstream effects of regulatory enforcement by the US Food and Drug Administration (FDA), an influential regulator overseeing products that account for over 20% of US consumer spending (Philipson and Sun, 2008). I focus on the agency's oversight of manufacturing facilities that produce sterile injectable (SI) drugs, which are administered in nearly all inpatient hospitalizations and play a central role in hospital care (Wosińska and Frank, 2023; Lahue et al., 2012). Oversight is high-stakes in this setting because SI drugs are injected directly into the bloodstream, meaning that quality deficiencies are life-threatening. Most importantly, SIs represent a natural setting for studying the access-safety trade-off because FDA officials have repeatedly identified them as a drug class where production disruptions are particularly concerning due to fears of supply chain fragility (Conti and Wosińska, 2025; Serchen et al., 2025; FDA, 2020b; Yurukoglu et al., 2017). Reflecting this concern, the FDA has stated it uses regulatory flexibility—temporarily permitting manufacturers to operate outside regulatory standards—to mitigate SI shortages (Woodcock and Wosinska, 2013).[1]

[1]My measure of FDA enforcement—Official Action Indicated (OAI) inspections representing the agency's initial discovery of non-compliance at a manufacturing facility—is selected so that its extensive margin occurrence is plausibly unaffected by such regulatory flexibility.

Studying how upstream FDA enforcement reverberates through concentrated supply chains requires overcoming two data challenges: the drugs manufactured at specific facilities are guarded as trade secrets (Galdin, 2025; Conti and Wosińska, 2025; Park et al., 2023), and SI use is not reliably recorded in standard healthcare claims datasets because they are typically reimbursed within broader bundled hospital payments (Alpert and Jacobson, 2019). I address these challenges by leveraging data from the Drug Enforcement Administration's (DEA) ARCOS database, which the agency uses to oversee the flow of controlled substances in the US from manufacturers to hospitals, clinics, and pharmacies. Using transaction-level detail in these data, I create a novel link to the origin facilities of SI opioid drugs and calculate high-frequency drug shipment volume at the level of individual hospitals.

The SI opioid drugs captured in the ARCOS data are the second most commonly administered category of inpatient SIs behind anti-infectives (Lahue et al., 2012), and are unrelated to the opioid crisis because their one-off use and administration by trained healthcare providers in a controlled hospital environment limits the scope for patient misuse. In fact, they are the only type of opioid drug considered medically essential by the FDA, owing to their efficacy in managing severe pain for surgeries, cancer treatment, labor, and other acute conditions (FDA, 2020a, 2018a). Like SIs more broadly, their manufacturing is highly concentrated: the top four facilities accounted for over 80% of annual volume over the study period. SI opioids also face persistent shortage risk, with the FDA designating at least one of the three primary active ingredients comprising the market as in shortage for over 60% of the study period.[2]

Combined with records of FDA manufacturing facility inspections and a 20% sample of Medicare claims, I use stacked difference-in-differences empirical strategies to trace the causal effects of FDA regulatory enforcement at upstream manufacturers down the supply chain to hospitals and the patients they serve, documenting three main findings. First, enforcement systematically disrupts production at implicated drug manufacturing facilities. Comparing facilities treated by FDA enforcement with matched controls enforced 1–3 years later and serving non-overlapping opioid drug markets—to avoid spillovers from competitor ramp-up—I find that enforcement sharply reduces drug volume by one-third. Accompanying increases in recalls and drug prices are consistent with the implementation of corrective action to restore compliance.

The extent to which upstream disruptions propagate to downstream hospitals is an empirical question of the resilience provided by other supply chain participants. Competing manufacturing facilities may counterbalance the shock by ramping up to fill the supply void, provided they hold excess capacity and face sufficient market incentives to expand production. Hospitals can also

[2]The FDA designates a drug as in shortage if it deems that demand exceeds commercially available supply. While a shortage signals supply chain strain, it does not necessarily imply measurable utilization reductions: Alpert and Jacobson (2019) found utilization reductions in only 5 out of 19 oncology drug shortages.

mitigate shortages through diversified procurement networks that create natural avenues for substitution to alternative drugs in the event of a disruption. However, each of these responses may be incomplete or slowed by frictions.

My second main finding is that the drug supply of downstream hospitals is remarkably resilient to FDA enforcement. For this analysis, I compare the trajectories of hospitals with different shares of baseline drug supply originating from the enforced facility. Despite facing mechanically larger supply disruptions due to the enforcement action, more-exposed hospitals fully recover by rapidly substituting toward alternative drugs, ultimately experiencing no change in total drug supply relative to less-exposed hospitals even in the short run.

This stability of drug volume is consistent with continuity of supply, not a temporary drawdown of inventories, and hospitals leveraging multi-sourced procurement networks to substitute along less costly margins of adjustment. Specifically, declines at the enforced facility are met with one-for-one increases in drug volume from other existing supply relationships, which provide familiar products for hospitals to integrate into clinical workflows. Together with increases in volume from new supply relationships in the first six months after enforcement, I find that for every 1 percentage point (pp) of baseline drug supply sourced from the enforced facility, hospital shipments from that facility fall by 0.21pp but are offset by a 0.24pp increase among alternative suppliers. However, this resilience is weaker for rural hospitals, which may face greater barriers to substitution due to less flexible procurement and limited bargaining power, and when enforcement action is taken during critical shortages, when the supply chain may lack excess capacity.

Third, despite the operational adjustments implied by rapid substitution, which may lead to rationing, delays in care, and a greater risk of medication errors even when drug access is maintained (Park et al., 2025), I find no evidence of adverse impacts on patients. Comparing index hospital visits before and after FDA enforcement and leveraging cross-sectional variation in their predicted exposure to the shock, I detect no effects on patient health, quality of care, or healthcare costs. These conclusions are robust across a range of specifications, including among rural patients and during critical shortages, when drug supply is most vulnerable. Evidence of patients temporarily shifting toward closer and more commonly used hospitals suggests that reallocations of patients across hospitals may be a margin of adjustment contributing to resilience.

Taken together, the results show that FDA enforcement produces precisely the kinds of large upstream shocks that, in markets with concentrated production and fragile supply chains, have the potential to create difficult trade-offs between decreased drug access in the short-term and long-term safety improvements as the enforced facility takes corrective action. However, my findings imply that such a trade-off does not materialize, on average, thanks to supply-side responses. The rapid reorganization of hospital procurement networks shields patients from drugs produced at non-compliant manufacturing facilities, which may otherwise put them in harm's

way, while maintaining access to essential medicines. Crucially, these adjustments also preserve both patient health and quality of care. Although SI opioids exhibit features of "too-big-to-fail," the evidence therefore suggests that the market adapts quickly to absorb even large shocks.

These findings likely extend beyond opioids to the broader class of SIs, which are produced under similar market conditions, by the same manufacturers, and even in the same facilities. In fact, opioid-specific supply chain rigidities created by DEA regulations (Conti and Wosińska, 2025), such as facility-level annual production limits and heightened security requirements, suggest the resilience I document likely represents a lower bound for SIs.

This paper contributes to three strands of literature. First, it builds on research examining the consequences of disruptions in healthcare, including strikes (Stoye and Warner, 2023; Gruber and Kleiner, 2012), bankruptcies (Antill et al., 2025), provider retirements (Fadlon and Van Parys, 2020), and natural disasters (Ziedan et al., 2022; Deryugina et al., 2020). The findings from this work are mixed: while studies of drug shortages point to resilience (Liebman et al., 2023; Alpert and Jacobson, 2019), Soliman (2025) and Donahoe (2024) document declines in drug access following DEA regulatory enforcement during the opioid crisis. By contrast, the FDA actions I study are part of routine regulatory oversight, outside the context of a public health crisis, allowing me to analyze how supply chains absorb enforcement-induced shocks under normal conditions, and shed light on the access-safety trade-off in concentrated markets.

Second, this paper connects to a growing literature on supply chain propagation, which has primarily used natural disasters to study how shocks travel through input-output networks to affect distant firms (Korovkin et al., 2024; Castro-Vincenzi et al., 2024; Balboni et al., 2023; Carvalho et al., 2021; Boehm et al., 2019; Barrot and Sauvagnat, 2016). This work provides empirical evidence demonstrating that, at least in other settings, shocks can propagate downstream. I extend this literature to the healthcare sector, where disruptions carry especially salient welfare implications because patient health is at stake, and go beyond firm-to-firm propagation by tracing the effects of upstream shocks down to the level of individual consumers. Doing so by mapping the flow of drugs across the supply chain—from manufacturing facilities that produce essential drugs, to hospitals that purchase them, and to patients who are treated by them—is this paper's key empirical contribution.

Third, this paper contributes to research on the FDA's navigation of the core access-safety trade-off. Prior work has primarily focused on the agency's role in new product approvals, where innovation is a central concern, and has generally characterized the FDA as prioritizing safety over access (Grennan and Town, 2020; Rogers, 2023; Budish et al., 2015; Philipson and Sun, 2008; Peltzman, 1973). This paper instead adds to a nascent literature that links drugs to their origin facilities and studies how this trade-off emerges after products have already been approved, especially amid concerns about supply chain fragility, concentration, and drug shortages. Manu-

facturers facing generic drug price competition have sought to reduce production costs through two primary channels. First, they have offshored production to countries with cheaper inputs and weaker regulatory oversight, which increases the risk of shortages and product deficiencies (Galdin, 2025; Noh et al., 2025; Wright et al., 2025). The second, and the focus of this paper, is pursuing economies of scale, which may increase risk as upstream disruptions become mechanically larger.[3] In the spirit of Starc and Wollmann (2025), who show that competition protects against drug price collusion through entry, my findings suggest that responses by other elements in the supply chain provide a source of natural resilience against enforcement-induced disruptions.

The remainder of the paper is organized as follows: Section 2 provides background on FDA regulatory enforcement and SI opioid drugs. Section 3 presents data and descriptive statistics. Section 4 provides first-stage evidence of disruptions caused by FDA enforcement at upstream drug manufacturing facilities. Sections 5 and 6 study the propagation of these disruptions on hospital drug supply and patient outcomes, respectively. Section 7 concludes.

2 Background

2.1 FDA Enforcement and Drug Manufacturing Oversight

The Food and Drug Administration's (FDA) mission is to promote public health by assuring the safety, efficacy, and quality of prescription drugs sold in the United States. Pursuant to this mission, all manufacturing facilities producing for the US market are subject to regular FDA inspections that assess compliance with current Good Manufacturing Practices (cGMP) standards (Conti and Wosińska, 2025; FDA, 2024b, 2023). These inspections comprehensively cover all aspects of the manufacturing process over multiple days—from production and processing to packaging and storage—and, due to limited manufacturer reporting requirements, serve as the agency's primary source of compliance information. Each manufacturing facility is inspected approximately once every two years according to a risk-based site selection rule, and inspections at opioid manufacturing facilities are not pre-announced (FDA, 2024a; Conti and Wosińska, 2025).[4]

FDA regulatory enforcement in the context of this paper refers to inspections resulting in an "Official Action Indicated" (OAI) classification, the worst and rarest outcome, when inspectors determine that a facility is in an unacceptable state of compliance with federal law. Although

[3]The offshoring channel is turned off in my setting as DEA regulations require that finished dosage form manufacturing occur in the US (Conti and Wosińska, 2025).

[4]The FDA's risk-based site selection rule considers compliance history, hazard signals (e.g., adverse patient events, recalls, complaints), product risk, and time since last inspection (FDA, 2024a). FDA inspections of domestic manufacturing facilities are conducted without prior notice, while those of facilities located outside the US are typically pre-announced. DEA regulations require that all finished dosage forms of opioids consumed in the US be manufactured domestically (Conti and Wosińska, 2025), which the FDA inspects without prior notice. Inspections of facilities located outside the United States are pre-announced (Wright et al., 2025).

OAIs do not necessarily imply an imminent risk to patients (FDA, 2024a), they reflect material breakdowns in cGMP standards intended to ensure that drugs are consistently manufactured to specification. Commonly cited deficiencies include contamination risks from leaks or standing water, poorly maintained equipment, and inadequate quality control or operating procedures (FDA, 2018b, 2016, 2010a,b).

Importantly, relative to "No Action Indicated" (NAI) or "Voluntary Action Indicated" (VAI) inspections that signal acceptable compliance, OAIs carry significant potential to disrupt production because they trigger requests for corrective action (Wright et al., 2025; FDA, 2023). For example, enforced facilities may need to upgrade equipment, physically redesign their site, or revise operating procedures. If corrections are not implemented in a timely manner, the FDA may withhold approval of pending drug applications (e.g., NDAs, ANDAs, and BLAs) linked to the enforced facility and suspend its participation in government procurement contracts (FDA, 2024b). Public disclosure of the facility's non-compliant status on the FDA website also creates reputational risk among buyers and investors.

In the most serious cases when conditions threaten patient safety, the FDA may request recall initiations, issue safety letters to healthcare providers, and make public warnings (FDA, 2024b; American Regent, 2011). Persistent or egregious violations may also be referred to the Department of Justice (DOJ) or state attorneys general for legal action.

2.2 Sterile Injectable Opioid Drugs

Among the prescription opioids covered by the DEA ARCOS data, which enable me to link drugs to their origin manufacturing facility, I focus on sterile injectable (SI) opioids to study how FDA regulatory enforcement propagates through pharmaceutical supply chains. These drugs, and the broader class of SIs produced alongside them, provide an ideal setting for this question because of their medical importance, and the FDA's stated concerns regarding supply chain fragility and shortage risk.

SI drugs are at the core of drug use in hospital settings, where they are administered by healthcare providers intravenously in nearly every hospitalization (Conti and Wosińska, 2025; Wosińska and Frank, 2023). That these drugs are injected directly into the bloodstream, bypassing the body's natural filters like the digestive tract and skin, means that quality deficiencies like contamination can be life-threatening. Opioids are the second most commonly administered SI in inpatient care, and are classified as medically essential by the FDA because they are hospitals' most effective tool for managing severe pain in acute settings, such as surgeries, cancer treatment,

labor, and traumatic injuries (FDA, 2020a, 2018a; Lahue et al., 2012).[5] [6] While they carry risks of adverse effects owing to their opioid active ingredients, SIs have less scope for patient misuse than the outpatient drugs linked to the opioid crisis: they are administered by trained healthcare providers on a one-off basis and in a supervised clinical environment.

Despite their medical importance, SI supply chains are feared to be fragile, creating the potential for FDA enforcement to propagate downstream. In practice, SIs are low-margin products subject to generic price competition, face high fixed and marginal costs due to the need for specialized production lines to ensure sterility, and their manufacturing tends to be highly concentrated. In turn, high storage costs for SI drugs, which are light- and temperature-sensitive, lead hospitals to rely on "just-in-time" inventory practices that provide little buffer against extended disruptions (ASHP, 2024; Wosińska and Frank, 2023; Woodcock and Wosinska, 2013).[7]

For opioid SIs specifically, the top four manufacturing facilities accounted for over 80% of annual volume over the 2007–2019 study period, and at least one of the three primary active ingredients—fentanyl, hydromorphone, and morphine—was designated as in shortage by the FDA for over 60% of that time. Figure 1, summarizing the 16 FDA enforcement events studied in this paper, illustrates that this combination of concentration and persistent shortage risk led to multiple instances of enforcement at high market share facilities while the supply chain was already under strain. Relative to other SIs, the opioid supply chain may be especially vulnerable to such shocks due to an additional layer of DEA regulation (Conti and Wosińska, 2025). For example, manufacturer ramp-up may be limited by DEA quotas, which set facility-specific caps on annual production. Re-purposing production lines after a disruption may also be limited by special rules governing documentation, physical security, and surveillance for opioid drugs.

Finally, the FDA and other stakeholders have repeatedly identified SIs as the drug class raising the greatest concern about the downstream consequences of production disruptions, reflecting an especially acute trade-off between drug access and safety (Conti and Wosińska, 2025; FDA, 2020b; Woodcock and Wosinska, 2013; Kweder and Dill, 2013). As evidence of this concern, the agency reported exercising "regulatory flexibility"—temporarily allowing manufacturers to operate outside regulatory standards to mitigate shortage risk—an average of 49 times per year from 2013 to 2019.[8] For example, the FDA permitted temporary imports of non-FDA-approved hy-

[5] The FDA considers a drug essential if it is used to treat or prevent a serious disease or medical condition for which there is no adequate substitute FDA (2020a). In limited circumstances, SI opioids have non-opioid substitutes that are not observed in the ARCOS data, such as nerve blockers used to treat hip fractures. In such cases, my analysis may understate the true extent of substitution.

[6] Specifically, Lahue et al. (2012) report that the mean number of administrations of injectable medications is 15.39 per inpatient hospitalization. The two most commonly administered classes are anti-infectives and narcotics/analgesics.

[7] An American Society of Health-System Pharmacists (ASHP) survey of hospital professionals responsible for inventory management indicates that roughly 70% of hospitals typically hold less than one month of generic sterile injectable drug inventory on hand, with only 2% holding more than 3 months of supply (ASHP, 2024).

[8] The FDA's annual Drug Shortages reports submitted to Congress beginning in 2013 list the number of instances of

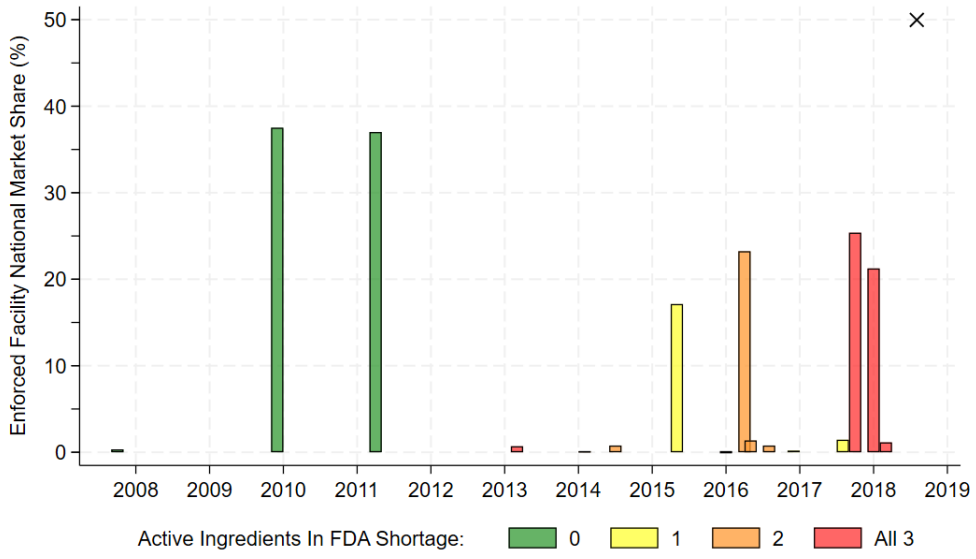


Figure 1. Supply Chain Pressure From FDA Regulatory Enforcement

Notes: Each bar marks one of the 16 FDA enforcement events impacting SI opioid manufacturing facilities in the study sample. The height of each bar represents the national market share held by the enforced facility, and the color of the bar represents how many of the three primary active ingredients comprising SI opioids were in shortage as of the month prior to enforcement. FDA declares a drug to be in shortage when it deems that demand exceeds supply. X corresponds to FDA permitting temporary imports of non-approved injectable hydromorphone from Canada in August 2018, an example of regulatory flexibility.

dromorphone injections from Canada in August 2018 in response to a critical SI opioid shortage caused by persistent drug quality problems (Hospira, Inc., 2018; FDA, 2018a).[9]

This background also motivates measuring FDA enforcement using OAI inspections, which represent the agency's initial discovery of non-compliance and serve as an "intent-to-enforce" starting point for subsequent regulatory action. Since regulatory flexibility as described by the FDA is predicated upon acknowledging non-compliance, it should not affect the issuance of OAIs at the extensive margin. However, it may cause subsequent actions like correction requests, recalls, warning letters, and injunctions to vary endogenously according to supply chain risk.[10] Regulatory flexibility is therefore best understood as part of the downstream supply-side adjustments that shape how markets respond to OAIs.

regulatory flexibility but do not identify affected products, so the extent of their application for SI opioids is not known.

[9] Imports of Canadian hydromorphone injections appear in the DEA ARCOS database from August 2018 through the end of the data in December 2019.

[10] Annual FDA drug shortage reports indicate that regulatory flexibility takes two primary forms: (i) limiting the size of disruptions at facilities with compliance problems, for example by permitting shipments of products with minor defects under special safeguards; and (ii) enhancing substitute availability by expediting review of new or modified products, expanding investigational access, or permitting temporary imports.

3 Data and Descriptive Statistics

Tracing the propagation of FDA regulatory enforcement from upstream drug manufacturing facilities to downstream hospitals and patients requires integrating three primary data sources, illustrated in Figure A.7.

FDA Inspections At the top of the supply chain, I use records of FDA inspections at drug manufacturing facilities spanning January 2000 to June 2024, obtained through a Freedom of Information Act (FOIA) request. The analysis considers inspections between July 2007 and December 2018 to align with coverage from other data sources and allow sufficient pre- and post-periods to estimate dynamic effects. Inspection records report the facility's FDA Establishment Identifier (FEI), name, address, inspection start and end dates, and outcome (i.e., NAI, VAI, or OAI).

Table 1 presents summary statistics for inspections at the 43 manufacturing facilities serving the SI opioid market over the study period.[11] On average, facilities were inspected 7.3 times, roughly once every 18 months, with 13% of inspections resulting in an OAI classification. To study downstream impacts, I focus on OAIs that are a facility's first in a rolling 12-month period (to exclude follow-up OAIs), and combine OAIs at different facilities but in the same month into composite events.[12] Figure 1 depicts the 16 OAI inspections comprising the study sample, where the height of each bar corresponds to the enforced facility's baseline national market share.

DEA ARCOS Database The central link between upstream enforcement and downstream outcomes is provided by the DEA's Automation of Reports and Consolidated Orders System (ARCOS). Ordinarily, the origin facility where prescription drugs are manufactured is protected as a confidential trade secret (Conti and Wosińska, 2025; Galdin, 2025).[13] This research leverages the fact that all drug manufacturers and wholesale distributors are legally required to report transactions involving controlled substances to ARCOS, which the DEA uses to monitor the flow of controlled substances in the US from manufacturing facilities to the pharmacies, hospitals, and clinics where they are dispensed (DEA, 1997). Using an extract covering the universe of prescription opioid shipments from 2006 to 2019, which became public following litigation related to the opioid crisis, I identify a drug's origin manufacturing facility as the one from which wholesale

[11] A manufacturing facility is classified as serving the SI opioid market if it holds at least a 0.1% share of monthly national volume at any point in the study period.

[12] I also restrict to OAIs for which the enforced facility held at least 0.1% SI opioid market share in months 13–18 before enforcement to restrict to facilities reliably serving this market.

[13] For example, even publicly released FDA inspection reports redact information on the drugs produced at the facility and other information considered proprietary or sensitive (see, for example, FDA (2018b)). Although the manufacturer of a particular drug is known, the specific facility where it is produced and whether it is made under contract by another firm is typically not disclosed.

Table 1. FDA Inspections of Manufacturing Facilities in DEA ARCOS Data

	Mean	Median	(SD)
Drug Manufacturing Facilities (N)	43		
Maximum Annual Market Share (%)	5.8	1.5	(10.4)
SI Opioid Recalls	1.7	0.0	(3.4)
Facility-Level FDA Inspection Statistics			
Total Inspections	7.3	5.0	(5.9)
Inspections Per Year	0.8	0.7	(0.7)
No Action Indicated Classifications (NAI)	3.2	3.0	(4.0)
Voluntary Action Indicated Classifications (VAI)	3.2	2.0	(2.8)
Official Action Indicated Classifications (OAI)	0.9	0.0	(1.2)
Study Sample OAIs[†]	0.5	0.0	(0.8)

Notes: Summary statistics consider inspections of drug manufacturing facilities that are under the primary jurisdiction of the FDA, serve the hospital-use sterile injectable opioid market, and have initial opioid drug volume recorded in the DEA ARCOS database at least 18 months prior to inspection. Facilities are classified as serving the hospital-use sterile injectable opioid market if they held at least a 0.1% share of national volume, measured in MMEs, in one or more months. The rate of FDA Inspections Per Year is equal to the number inspections between July 2007 and December 2018 divided by the number of months over this period, restricting to the period after the facility first records positive drug volume in the ARCOS database.

[†] OAIs are included in the study sample if they represent the first at the enforced facility in a rolling 12-month period (to restrict to initial OAIs only, excluding follow-ups) and if the enforced facility held at least 0.01% market share 13–18 months prior to inspection.

distributors are recorded as purchasing the drug. Manufacturing facility FEIs reported in the FDA inspections data are mapped to DEA registrations in the ARCOS data using a fuzzy matching procedure based on addresses and manufacturer names. Over 99.9% of SI opioid volume is matched to an origin manufacturing facility.

From this link to origin facilities, which to the best of my knowledge represents a novel use of the ARCOS data, Figure 2 plots the market shares held by top SI opioid manufacturing facilities over the study period. The top four accounted for over 80% of volume for nearly the full period, with the top facility accounting for roughly one third of market volume. Circular markers in Figure 2 represent instances of FDA enforcement at these top facilities during the study period.

Another key use of the ARCOS data is to provide a reliable measure of SI opioid drug volume at the hospital-month level. Creating such a measure is not feasible with insurance claims datasets commonly used in health economics research because SI drugs may be reimbursed as part of broader bundled inpatient payments. While ARCOS does not capture utilization directly, its transaction-level detail records the arrival of new SI opioid inventory to hospitals, which forms the basis of my analysis of hospital drug supply.

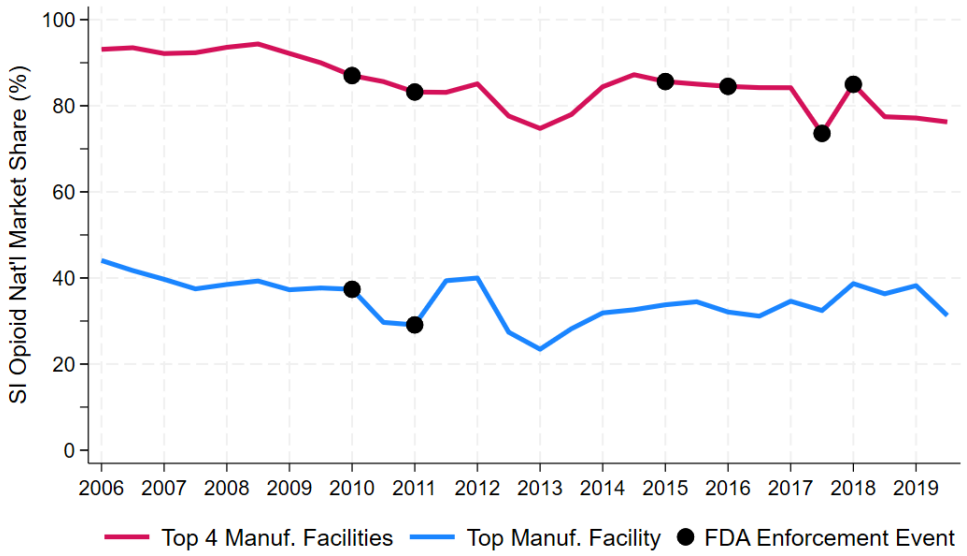


Figure 2. Top 4 Sterile Injectable Opioid Manufacturing Facility Market Share

Notes: Figure 2 is constructed by identifying the top 4 manufacturing facilities by national market share in each half-year period. Circular markers refer to study sample FDA enforcement events impacting top 4 and the top manufacturing facility, respectively.

Medicare Claims To examine impacts on downstream patient health, quality of care, and healthcare costs, I use a 20% sample of administrative Medicare Part A and B billing claims. My preferred outcome is Healthy Days at Home (HDAH), a measure of quality of care and population health developed by Burke et al. (2020) and Lam et al. (2021) in conjunction with the Medicare Payment Advisory Commission (MedPAC). The spirit of HDAH is to classify each patient-day as either "healthy" or "unhealthy," where healthy refers to days in which the beneficiary is alive and not in a facility-based care setting. Though patients may be in poor health even outside of facility-based care, the spirit of HDAH is to focus on the most therapeutically intensive elements of the healthcare system, namely inpatient and post-acute care. My analysis follows the approach taken by (Burke et al., 2023), calculating HDAH indexed to the period following a hospital visit.

Hospital-level SI opioid drug volumes from ARCOS are matched to Centers for Medicare and Medicaid Services (CMS) provider numbers using a fuzzy matching procedure based on addresses and hospital names, enabling me to link upstream shocks to downstream patient outcomes. Over 95% of hospital admissions are matched to a hospital DEA registration in the ARCOS data.

4 Upstream Manufacturing Facility Disruption

I begin by estimating the causal effects of FDA regulatory enforcement on drug volume originating from the implicated manufacturing facility. This can be understood as a first stage: en-

forcement can only propagate to downstream hospitals and patients if it first disrupts production at the upstream facility. Specifically, I focus on FDA inspections resulting in an OAI classification, which formally declare unacceptable regulatory compliance and require the enforced facility to implement potentially costly and time-consuming corrective action.

4.1 Empirical Strategy

The central challenge to identifying causal effects is that FDA regulatory enforcement is not randomly assigned to manufacturing facilities. Two sources of endogeneity are particularly salient. First, *whether* a facility receives an OAI—rather than a VAI or NAI—depends on the presence of serious regulatory violations, which are likely correlated with latent manufacturing problems that may make enforced facilities inherently more disruption-prone. Second, *when* a facility receives an OAI is shaped by the FDA's risk-based inspection rule, which prioritizes sites that emit hazard signals like recalls, complaints, and adverse patient events. This raises the concern that FDA enforcement may coincide with disruptions that are already underway as facilities proactively respond to production problems associated with hazard signals.

To address these concerns, I implement a difference-in-differences design that compares facilities treated by an OAI with matched controls that go on to receive an OAI one to three years later. Restricting to treatment and control facilities on a shared enforcement trajectory supports the plausibility of parallel trends, which is empirically assessed with the pre-trends falsification test.[14] To isolate the effects of FDA enforcement from confounding events, I implement my analysis in monthly event time. Assuming that the precise month of enforcement is plausibly exogenous within a narrow window, such as due to delays in assessing hazard signals or variation in inspector availability, anticipatory declines prior to the OAI would suggest the presence of confounding events. On the other hand, effects emerging sharply after the inspection can be attributed to enforcement. As a placebo test, I replicate this design for VAI inspections, in which the FDA uncovers some violations but deems acceptable compliance.

My baseline specification draws on OAI inspection events across all opioid manufacturing facilities represented in the DEA ARCOS data, not only those producing SI opioids.[15] This broad

[14] See Coyne et al. (2024) for further discussion of the stacked difference-in-differences research design.

[15] Drawing on World Health Organization Anatomical Therapeutic Chemical (ATC) codes and clinical uses on the drug's FDA-approved drug label, I classify non-SI opioids in the DEA ARCOS data into three markets capturing distinct therapeutic uses. The first consists of combination opioid pills that pair an opioid ingredient with a non-opioid analgesic, such as hydrocodone combined with acetaminophen or ibuprofen, to treat moderate pain. The second market is pure opioid pills, like oxycodone and morphine tablets, which are used to treat more severe pain and carry higher misuse potential. Note that combination opioids cannot be substituted for more potent pure opioids because their non-analgesic components become toxic at higher dosages. The third market includes specialized dosage forms like transdermal patches, nasal sprays, lozenges, and liquid solutions, which are intended for patients with advanced care needs. For example, patients who are highly opioid-tolerant, have difficulty swallowing, or require continuous pain management over multiple days.

design is important for several reasons. To begin with, credible identification requires matched controls that serve strictly different opioid markets and therefore do not directly compete with the enforced facility. Since competing facilities may ramp up when rivals are disrupted, including them risks biasing estimates. Drawing on the full set of enforcement events also enables the analysis to provide, to my knowledge, the first systematic evidence of the causal effects of FDA enforcement, a contribution with intrinsic policy relevance given interest in understanding the drivers of supply chain risk (CRS, 2023; FDA, 2020b; CRS, 2018).[16] [17] Finally, the broader sample of opioid markets allows for the analysis of impacts on drug prices using Medicare Part D prescription fills data, which do not reliably capture SI opioids because they are typically reimbursed through bundled inpatient payments. To complement my analysis of downstream propagation, which focuses on enforcement events at SI opioid manufacturing facilities where the access-safety trade-off is thought to be most acute, I also estimate a first-stage specification restricted to these events. The results are robust and validate the existence of a first-stage upstream disruption in the downstream analysis.

The causal effects of FDA OAI inspections on the enforced facility are estimated using the following stacked event study specification,

$$Y_{f,m,t,\tau} = \sum_{r=-6, r\neq -1}^{12} \beta_r \, \text{Enforced}_{f,m,\tau} \times \mathbf{1}\left\{r \text{ Months Since Enforcement}\right\}_{f,t,\tau} + \delta_{f,m,\tau} + \eta_{r,\tau} + X_{f,m,t,-\tau} + \varepsilon_{f,m,t,\tau} \tag{1}$$

where $Y_{f,m,t,\tau}$ measures the outcomes (e.g., drug volume, recalls, prices) of manufacturing facility f in opioid market m in month t around FDA OAI inspection τ, with event time indexed by r.[18] Aggregating outcomes to the facility-by-market level, rather than to the facility level, preserves cross-market differences in drug volume and prices, preventing results from being systematically driven by high-level markets.[19] Accordingly, the regression includes two-way facility-by-market-by-event fixed effects $\delta_{f,m,\tau}$ and event-by-event time fixed effects $\eta_{r,\tau}$. The $\text{Enforced}_{f,m,t}$ treatment indicator equals 1 for facilities enforced in event τ and 0 for future-treated matched controls.[20] The vector of controls $X_{f,t,-\tau}$ captures the impacts of OAI inspections with event

[16] An FDA report on the root causes of drug shortages lists quality issues, increases in demand, natural disasters, and product discontinuations as the most common sources of drug shortages from 2013–2017. This analysis sheds light on the role of regulatory enforcement within quality issues, the most common cause of shortages, which may also refer to recalls and other product defects (FDA, 2020b).

[17] Wu and Wang (2023) also study the impacts of FDA inspections on manufacturing facility outcomes, but focus on the relationship between inspection frequency and drug recalls using a distance-to-facility as an instrumental variable, rather than the impact of inspections resulting in enforcement.

[18] The OAI inspection index τ refers to each opioid market-by-month (m, t) with an FDA enforcement event.

[19] About 30% of facilities in my sample serve more than one market.

[20] For a given treatment group OAI inspection, matched controls must (1) have no OAI inspection in event months

horizons overlapping with τ, as drug volume from both treatment and control facilities may respond to disruptions associated with these events.[21] Following guidance in Wing et al. (2024), standard errors are clustered at the manufacturing facility level f. Inference is conducted using the cluster wild bootstrap because there are 45 distinct manufacturing facility clusters (Roth et al., 2023; Cameron et al., 2008).

4.2 Results

Figure 3 plots the differential evolution of drug volume between enforced facilities and future-treated matched controls, where drug volume is measured in milligrams of morphine equivalents (MMEs, a standard potency-adjusted measure of opioid quantity) and divided by baseline total market volume to yield a market share interpretation. Treatment and control facilities evolve along parallel trends prior to enforcement, followed by a persistent decline in drug volume from the enforced facility beginning sharply in the month after the OAI. Difference-in-difference estimates reported in Table 2 show that drug volume from enforced facilities decreases by the equivalent of 2pp of market share in the year after enforcement (33% of the baseline mean).

To explore mechanisms, panel (a) of Figure A.1 shows that FDA enforcement is followed by an increase in active Class I and II drug recalls originating from the enforced facility.[22] Overall, recalls rise by 0.2 units in the year after enforcement (equivalent to one new recall for every five enforced facilities, or 51% of the baseline mean), with timing closely tracking the onset of drug volume disruptions. This pattern is consistent with FDA statements indicating that the agency requests voluntary recalls after inspections that uncover product issues posing an immediate risk of consumer harm. Together, the timing of the decline in drug volume and rise in recalls in the post-enforcement period supports the interpretation that the disruption is caused by FDA enforcement rather than confounding events.

Panel (b) examines impacts on drug prices, measured as the logarithm of total Medicare Part D prescription spending per MME dispensed.[23] The results suggest an approximately 25% in-

$r \in [-12, 12]$ to prevent contamination from past enforcement; (2) go on to receive an OAI inspection in event months $r \in [13, 36]$; and (3) must produce for opioid markets that are not served by the implicated facility. To exclude volatility from manufacturing facilities with low production volumes, treatment and control facilities are required to record positive drug volume in at least two months the six months prior to the FDA treatment group enforcement event and hold at least 0.01% national market share in the period 13–18 months prior to it. See Appendix B.3 for a detailed discussion of the matching procedure.

[21] The controls consists of national market shares held by the manufacturing facility implicated by overlapping events, capturing the size of the disruptions, and interacts them with event time indicator variables reflecting the timing of those overlapping events.

[22] The FDA classifies recalls into classes I–III, where class I and II represent situations where use of the recalled product may cause adverse health consequences. Class III recalls are defined as situations in which use of the product is not likely to cause adverse consequences.

[23] For this calculation, drug volume in MMEs is measured with Medicare Part D data rather than the DEA ARCOS database, and total prescription costs include both payments by the insurer and the beneficiary. Because of missing

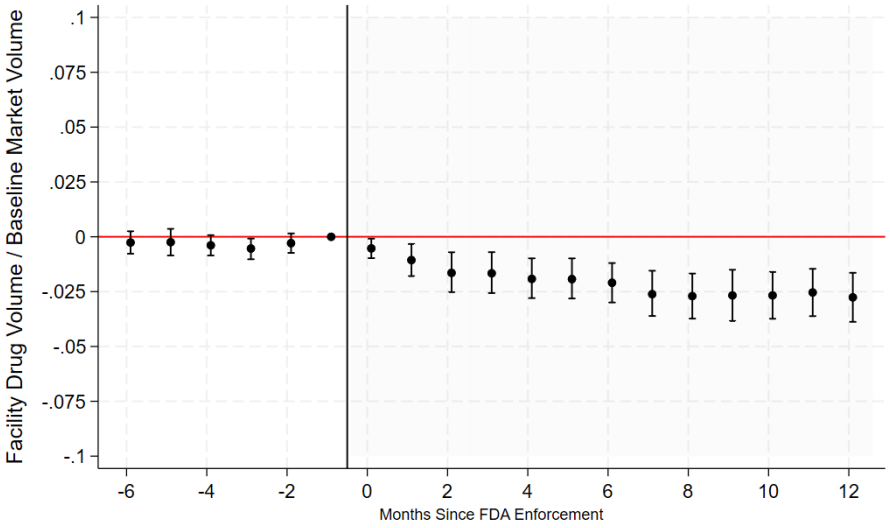


Figure 3. Impact of FDA OAI Inspections on Enforced Facility Drug Volume

Notes: Figure 3 plots the differential evolution of manufacturing facility-by-opioid market outcomes between enforced manufacturing facilities that matched controls that were enforced 1–3 years later and serve opioid drug markets that do not overlap with those of the enforced facility. Error bands represent 90% confidence intervals using standard errors calculated with the wild cluster bootstrap and manufacturing facility-level clusters. The outcome variable is national drug volume originating from the implicated facility measured in milligrams of morphine equivalents divided by mean total market volume in the baseline period 13–18 months prior to enforcement. The results can therefore be interpreted as percentage point changes in baseline market share.

Table 2. Impact of FDA OAI Inspections on Enforced Facility Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)
	All Opioid Manufacturing Facilities				**SI Facilities Only**	
	Market Share	Active Recalls	Log(Enforced $/MME)	Log(Rest of Market $/MME)	Market Share	Active Recalls
Enforced× 0–12 Months Post	−0.02***	0.20*	0.26*	−0.00	−0.02**	0.46**
N	4,465	4,465	760	760	1,330	1,330
R^2	0.98	0.90	0.97	0.99	0.98	0.92
Dep. Var. Pre-Mean	0.06	0.39	−3.15	−3.03	0.11	0.80
Facility-by-Event FEs	✓	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓	✓
Overlapping Event Controls	✓	✓	✓	✓	✓	✓
Wild Cluster Bootstrap 95% CI	[-.025, -.011]	[.012, .415]	[-.001, .587]	[-.057, .048]	[-.0338, -.004]	[.054, .895]
Wild Cluster Bootstrap p-value	0	0.08	0.10	0.99	0.03	0.03
# of Events	55	55	15	15	15	15
# of Facilities	45	45	19	19	31	31

Notes: * 10%, ** 5%, *** 1%. Table 2 reports the results of difference-in-differences estimates comparing enforced facilities with matched controls enforced 1–3 years later and serving non-overlapping opioid drug markets (see equation 1). Outcomes include drug volume (divided by baseline total market volume), active class I and II drug recalls, and the logarithms of facility drug prices and rest of market drug prices. Columns 5 and 6 report the results when restricting to FDA enforcement events impacting SI opioid manufacturing facilities, which are further studied in the analysis of downstream propagation. All inference is conducted using the wild cluster bootstrap with manufacturing facility-level clusters.

crease in enforced facility drug prices in the year after enforcement, statistically significant at the 10 percent level, while prices of drugs in the same market but produced by other manufacturing facilities remain unchanged. This asymmetry suggests that enforcement raises production costs at the implicated facility, such as due to implementation of corrective action or the higher costs of operations that fully comply with FDA standards, while the remaining market remains competitive. Indeed, opioid markets are composed of many therapeutically equivalent active ingredients and dosage forms, creating scope for substitution. However, these effects on prices may not generalize to SIs, which are manufactured in a comparatively more concentrated market structure and reimbursed through bundled payments rather than per unit dispensed.

Placebo estimates replicating this empirical strategy for VAI inspections yield statistically insignificant and near-zero effects on drug volume, recalls, and prices (see Appendix Section B.1). The contrast between the effects of OAI and VAI inspections suggest that the regulatory actions associated with OAIs—corrective remediation, recall requests, hindrance to application approvals, exclusion from government procurement, etc.—are likely responsible for the disruptions.

Finally, because my analysis of the downstream impacts of FDA enforcement focuses on SI opioids, Table 2 also reports the results of specifications restricting to this subset of treatment events.[24] The results validate that the first-stage disruption is robust in this sample: drug volume from enforced facilities decreases by 2pp of market share (18% of the baseline mean) and active Class I and II recalls increase by 0.46 (58% of the baseline mean) in the year after enforcement. Dynamic difference-in-differences estimates in Appendix Figure B.1, albeit less precise given the smaller number of enforcement events (15 compared to 55, using non-SI opioid manufacturing facilities purely as controls), also show the sharp onset of treatment effects. These disruptions are further reinforced in the analysis of hospital drug supply in Section 5, which documents a persistent decline in enforced facility drug volume among more-exposed hospitals.

Taken together, the first-stage results demonstrate that FDA regulatory enforcement from OAI inspections systematically disrupts production at implicated manufacturing facilities, reducing drug volume by one-third relative to baseline and prompting increases in recalls and facility-specific drug prices. These findings support the hypothesis that concentration may complicate the FDA's core trade-off between drug access and safety, as enforcement against dominant facilities can cause large upstream supply shocks. Non-compliance therefore poses a dual challenge for the FDA: consumers may be harmed not only by exposure to sub-standard drugs, but also by

manufacturing facility-by-month drug price data, including for the entire class of sterile injectable opioid drugs, this analysis uses a smaller number of FDA enforcement events than for drug volume and recalls. Specifically, to prevent from results being driven by compositional differences, I restrict the regression sample to drug price trajectories that are non-missing throughout the event horizon.

[24] Due to the lack of a matched control for one event, which targeted a manufacturing facility holding 1.3% market share at baseline, the first-stage analysis of SI opioid manufacturing uses 15 events rather than the 16 studied in Sections 5 and 6.

short-term reductions in access as enforcement restores compliance. Sections 5 and 6 quantify these downstream effects in the context of SI opioids, which belong to the broader class of SIs that the FDA has characterized as having particularly vulnerable supply chains.

5 Downstream Hospital Drug Supply

Next, I turn to estimating how disruptions to upstream manufacturing facilities propagate downstream, beginning with impacts on hospital drug supply. Propagation depends on the relative sizes of two opposing forces: the FDA enforcement-induced disruption at the implicated facility and countervailing supply-side responses by other supply chain participants that may lend resilience (Acemoglu and Tahbaz-Salehi, 2025; Carvalho et al., 2021; Barrot and Sauvagnat, 2016). These responses operate at two levels of the supply chain—upstream manufacturers and downstream hospitals—and their speed and extent may be limited by frictions. This discussion provides a framework for interpreting and generalizing the empirical estimates that follow, which capture the *net* reduced-form effects of these opposing forces on hospital drug supply.

At the manufacturing level, the key question is whether competing facilities expand production to fill the supply void left by the enforced facility. This response depends on the availability of excess production capacity and the strength of market incentives to ramp up. Since SI drugs are typically mature generics without patent protection, excess capacity arises from the presence of multiple manufacturers approved to produce them.[25] Generic drug price competition means that price increases following disruptions are plausible, though they may not materialize in the short-run because hospital procurement contracts often set maximum prices.[26] Nevertheless, fixed costs in SI production create economies of scale that may independently motivate production increases.

Next, hospitals may reorganize their procurement operations to substitute to alternative drugs. This response depends on the availability and closeness of substitutes (i.e., the input-specificity of disrupted drugs), adjustment costs to incorporate them into clinical workflows, and the price elasticity of hospital demand. SI drugs have both perfect substitutes—bioequivalent drugs from other manufacturers—and imperfect substitutes that are therapeutically equivalent but differ in potency, dosage form (e.g., vial, syringe, ampule), or active ingredients. While the former can be used interchangeably, the latter may compromise efficacy or increase the risk of medication errors due to reduced product familiarity.[27] Additionally, substitution toward suppliers already

[25] By contrast, branded drugs are produced by a single monopolistic manufacturer, so excess capacity depends on the firm's own redundancies, which may be incentivized by high markups.

[26] FDA found sustained price increases for 18% of drugs that went into shortage between 2013 and 2017 (FDA, 2020b).

[27] Not all SI opioids are interchangeable for all clinical indications (e.g., morphine may be unsuitable for older patients with renal impairment) and converting across active ingredients requires careful conversion of concentrations as overdoses can be deadly. The ASHP warns that providers should "exercise extreme caution when purchasing products that are not regularly stocked" (ASHP, 2018).

in hospitals' procurement networks is likely less costly, avoiding the search, staff re-training, and workflow adjustments associated with integrating new suppliers. Although the medical importance of SI drugs contributes to inelastic demand, hospitals may reduce utilization through rationing, delaying elective procedures, reducing waste, and approving new substitutions (Vo et al., 2021). The use of "just-in-time" inventories adds to the urgency of substitution.

Given this framework, SI opioids provide a plausible stress-test within the broader class of SIs where the propagation of FDA regulatory enforcement is especially likely. Supply chain rigidities imposed by the additional layer of DEA regulations may limit excess capacity at non-enforced facilities.[28] Although medically essential, the demand of SI opioids may be more elastic than other SIs, like anti-infectives, because opioids provide palliative rather than curative care. Even short-lived reductions in drug supply caused by enforcement may carry meaningful welfare implications when considering the collective weight of dozens of FDA OAIs issued every year.

5.1 Empirical Design

Since hospitals' treatment status switches "on-and-off" repeatedly, up to once for each of the 16 FDA regulatory enforcement events during the study period, I estimate the causal effects of enforcement on the drug supply of downstream hospitals using a stacked difference-in-differences design. This approach leverages three sources of variation: (i) temporal variation before and after enforcement, (ii) cross-sectional variation in hospitals' reliance on the enforced facility for drug supply at baseline, and (iii) collective evidence from 16 distinct enforcement events.[29] Intuitively, the design compares drug supply trajectories across hospitals with differing levels of exposure to a given FDA enforcement action.

For each hospital h and upstream FDA enforcement event τ, exposure is measured by the share of a hospital's baseline drug supply that originated from the enforced facility.[30] [31] This quantifies the hospital-specific disruption caused by upstream enforcement, capturing the mechanical loss in volume if the enforced facility shut down and hospitals made no substitutions. This construction parallels the "shift-share" approach widely used to study the impacts of distant

[28] If manufacturers view quota increases as permanent, the quota system may strengthen incentives to ramp up over longer time horizons, as disruptions would offer an opportunity to permanently improve economies of scale.

[29] While the effects of a single FDA enforcement action could be analyzed as a standalone case study using a comparable difference-in-differences approach, the objective of this paper is to characterize the *systematic* effects of enforcement across the full set of events affecting SI opioids during the study period. Pooling events together in a stacked design mitigates concerns that results reflect idiosyncratic conditions surrounding any single enforcement event. However, this approach introduces complexity in the form of overlapping event horizons across FDA enforcement actions that create scope for contamination, which I address explicitly in the empirical specification.

[30] The FDA enforcement event index τ consolidates all enforcement actions that take place in the same month across different upstream manufacturing facilities into a single event.

[31] The baseline period spans the period 13–18 months prior to enforcement. This creates a six-month buffer between the construction of the treatment variable and the start of the event horizon in event month $r = -6$, thereby allowing for the resolution of potential mean-reversion.

shocks in the trade literature (Borusyak et al., 2025, 2022; Autor et al., 2016), where manufacturing facility-level FDA enforcement shocks are interacted with hospital-level supplier shares.

Although hospital supply networks may be formed endogenously—for example, hospitals that are more reliant on low-compliance suppliers may face resource constraints or differ in quality of care—causal identification relies on parallel trends across hospitals with different supplier networks in the counterfactual world without the FDA enforcement action. This assumption is plausible because of the separation between downstream hospitals and the enforcement actions: not only are hospitals not the targets of FDA enforcement, those actions take place multiple levels upstream in the supply chain. It is therefore reasonable to conceive of enforcement timing as quasi-random from the perspective of downstream hospitals.

The principal threat to identification arises from the FDA's risk-based targeting of inspections, which may result in enforcement systematically following a period of latent deterioration in production quality at the implicated manufacturing facility. Consequently, more- and less-exposed hospitals may experience differential deterioration in the quality of their drug supply over the event horizon independently of FDA action. Hazard signals like adverse drug events and healthcare provider complaints potentially associated with quality changes may also create a link between downstream outcomes and upstream enforcement. Although statements by FDA officials indicate that OAIs typically preemptively address potential problems before they manifest into realized ones (FDA, 2023), my empirical strategy quantifies this concern with the parallel pre-trends falsification test. To the extent that there is an inherent lag between the emergence of quality issues and upstream FDA regulatory enforcement, the systematic worsening of drug supply may appear as "anticipatory" responses. For example, hospitals may recognize deviations in drug quality through firsthand experience and begin substituting away from the enforced manufacturer prior to the OAI.[32]

The stacked data structure creates an event-specific panel dataset containing the monthly trajectories of hospital drug volume around each FDA enforcement action.[33] Since there is no variation in treatment timing across hospitals—enforcement always occurs in event month $r = 0$—concerns about differential treatment timing from the modern difference-in-differences literature do not apply. I therefore estimate the following stacked two-way fixed effects difference-in-

[32]The FDA's consideration of hazard signals at downstream hospitals within their risk-based inspection targeting rule could also lead to reverse causality, as adverse events at high-exposure hospitals in the pre-period may cause upstream FDA enforcement. The parallel pre-trends test also offers empirical evidence regarding this channel.

[33]The stacked data structure allows for the straightforward construction of hospital variables specific to each enforcement event, including treatment variables, outcome variables, hospital characteristics such as the manufacturing facilities with which it had an existing supply relationship at baseline.

differences model (Wing et al., 2024),[34]

$$Y_{h,t,\tau} = \sum_{r=-6, r\neq -1}^{12} \beta_r \, \text{HospitalExposure}_{h,\tau} \times \mathbf{1}\left\{r \text{ Months Since Enforcement}\right\}_{t,\tau} + \delta_{h,\tau} + \eta_{r,\tau} + X_{h,t,-\tau} + \varepsilon_{h,t,\tau} \tag{2}$$

where $Y_{h,t,\tau}$ measures drug volume shipped to hospital h in month t around FDA OAI event τ, with monthly event time indexed by r. The treatment variable $\text{HospitalExposure}_{h,\tau}$ measures the share of hospital h's baseline drug supply accounted for by the facility implicated by event τ. Two-way fixed effects include hospital-by-event fixed effects $\delta_{h,\tau}$, which absorb time-invariant level differences across hospitals within each event, and event-by-event time fixed effects $\eta_{r,\tau}$, which capture shocks common to all hospitals at each event time r of event τ. The vector $X_{h,t,-\tau}$ controls for potential contamination from enforcement events $-\tau$ with overlapping time horizons.[35] To provide a percentage point (pp) change in drug volume interpretation, I divide outcomes by each hospital's mean monthly drug volume at baseline, and weight hospitals by this mean to reduce volatility from low-volume hospitals.[36] Standard errors are clustered at the hospital level.

Table 3 summarizes the quasi-experimental variation leveraged in this analysis. Panel B characterizes the exposure of hospitals' baseline drug supply to upstream FDA enforcement. On average, hospitals relied on the enforced facility for more than 0.01% of drug supply in six of 16 FDA enforcement events. Among these exposed hospitals, the mean share of drug supply from the enforced facility is 26%, and the mean share in hospitals' maximum exposure event is 50%.

Panel C summarizes hospitals' procurement patterns in the six months preceding enforcement. Hospitals receive shipments of SI opioids in nearly every month, with average drug volumes remaining stable relative to the baseline period spanning 13–18 months before enforcement. Procurement exhibits multi-sourcing: at baseline, hospitals source from an average of 3.9 distinct manufacturing facilities (only 3% of hospitals single-source) and 1.6 distributors per month. The mean monthly Herfindahl-Hirschman Index (HHI) of 4,467 indicates a level of diversification marginally greater than dividing purchases evenly across two manufacturing facilities.

[34] The (unstacked) two-way fixed effects analog, which removes event indices τ from the fixed effects and treatment variable, is $Y_{h,t} = \sum_{r=-6,r\neq-1}^{12} \beta_r \, \text{HospitalExposure}_h \times \mathbf{1}\left\{r \text{ Months Since Enforcement}\right\}_t + \delta_h + \eta_r + \varepsilon_{h,t}$. This specification could hypothetically be used to study each of the 16 FDA enforcement events in isolation, setting aside concerns about overlapping treatment.

[35] This vector models the treatment effects of overlapping events with the same structure as the focal event τ: with event time-by-HospitalExposure interaction terms reflecting the timing and hospital-specific exposure to $-\tau$. The construction of hospital exposure using the lagged baseline period $r \in [-18, -13]$ means that $X_{h,t,-\tau}$ is constructed exclusively using information about hospitals' supply networks from the pre-enforcement period, even if the overlapping event occurs late in the focal event horizon (e.g., $r = 12$). This avoids the problem of "forbidden controls," whereby hospital exposure to a future event may be shaped by substitution caused by the focal event.

[36] Results are robust to weighting hospitals equally.

Table 3. Hospital-Level Regression Sample Summary Statistics

	Mean	(SD)
Panel A: Hospital Characteristics		
Hospitals (N)	3,729	
% Teaching Center	30	
% Rural Hospital	33	
Panel B: Exposure to FDA Enforcement Events		
Events with HospitalExposure $\geq 0.1\%$	6	(2)
% HospitalExposure Among HospitalExposure $\geq 0.1\%$	26	(20)
Maximum % HospitalExposure Across All Events	50	(16)
Panel C: Pre-Enforcement Monthly Hospital Procurement		
Drug Volume / Hospital's Baseline Mean	1.0	(1.2)
% of Months with SI Opioid Volume > 0	99	(10)
# of Manuf. Facility Suppliers	3.9	(1.5)
% Single-Sourcing	3.3	
HHI across Suppliers (0–10,000)	4,455	(1,734)
# of Distributors	1.6	(0.6)
% Fentanyl	49	(24)
% Hydromorphone	27	(19)
% Morphine	19	(16)

Notes: Panel A of Table 3 reports the distinct hospitals in the sample, and the share classified as rural (in a non-metro Rural-Urban Continuum ZIP code) and teaching centers. Panel B summarizes the exposure of hospitals' drug supply to FDA enforcement events using months 13–18 before enforcement. Panel C summarizes hospitals' procurement practices in months 1–6 before enforcement.

Hospitals also order meaningful quantities of all three primary active ingredients comprising SI opioids. Taken together, these patterns suggest that despite concentration in manufacturing, hospital procurement networks maintain natural avenues of substitution, across suppliers and active ingredients, providing a foundation to help mitigate upstream disruptions.

5.2 Results

Figure 4 plots the evolution of drug supply across more- and less-exposed hospitals around FDA regulatory enforcement events, revealing substantial resilience in the pharmaceutical supply chain. Consistent with the first stage disruption documented in Section 4, more-exposed hospitals experience a discontinuous decline in drug volume originating from the enforced facility starting sharply in the month of enforcement. Regression estimates imply that for every 1pp of baseline drug supply sourced from the enforced facility, shipments from that facility fall by 0.21pp of total

baseline volume in the first six months after enforcement, growing to 0.36pp in months 7–12.[37] [38]

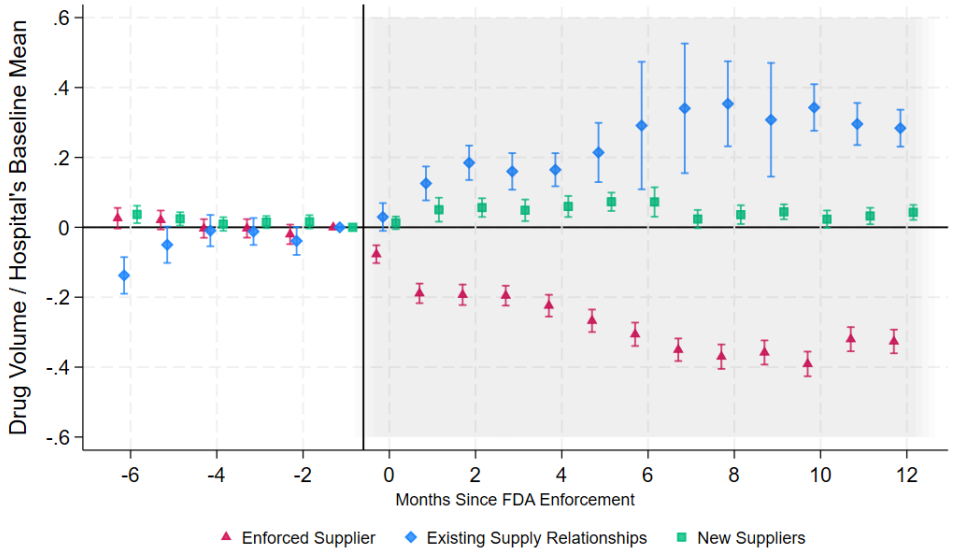


Figure 4. Impact of FDA Enforcement on Decomposed Hospital Drug Supply

Notes: Figure 4 plots the evolution of hospital sterile injectable opioid drug volume around FDA OAI inspections. The dynamic difference-in-differences estimates corresponding to $\hat{\beta}_r$ from equation 2 compares the trajectories of hospitals differentially reliant on the enforced upstream manufacturing facility at baseline for drug supply. The figure plots the results of three separate regressions, one using each element of the following total drug volume decomposition as the outcome variable: $q_{Total} = q_{OAI} + q_{ExistingOther} + q_{NewOther}$. Error bands represent 95% confidence intervals with standard errors clustered at the hospital level.

However, these losses are offset by an immediate rise in drug volume from alternative suppliers. For each 1pp of baseline hospital drug supply sourced from the enforced facility, shipments from non-enforced facilities increase by 0.24pp of total baseline drug volume in months 0–6 after enforcement and 0.38pp in months 7–12. This pattern of rapid substitution is driven primarily by increased volume from manufacturing facilities in hospitals' existing supply networks, which increases approximately one-for-one with declines from the enforced facility (see Table 4).[39] New supply relationships are also important, accounting for 15% of the supply increase in the first six months after enforcement, though their role diminishes thereafter.

On net, these rapid procurement adjustments allow more-exposed hospitals to fully recover

[37] Equivalently, hospitals that were fully reliant on the enforced facility experience declines of 21% and 36% of baseline volume over the first and second six months, respectively.

[38] Drug volume is measured in milligrams of morphine equivalents (MMEs), a standard potency-adjusted measure of opioid shipments, divided by the hospital's total drug volume in the baseline $r \in [-18, -13]$ period, giving the results a percentage point change relative to baseline interpretation.

[39] A manufacturing facility is classified as being in a hospital's existing supply network if it accounts for any positive drug volume in the baseline $r \in [-18, -13]$ period. New supply relationships are therefore free to form beginning in event month $r = -12$, meaning that the parallel pre-trends test for this outcome provides a meaningful test for the differential development of new supply relationships across hospitals.

Table 4. Effect of FDA Enforcement on Hospital Drug Supply

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Volume Type / Hospital Baseline Total Volume						
	Total Drug Supply	Enforced	Existing	New	Supplier HHI	Ingred. Sub.	# Distributors
HospitalExposure× 0–6 Mo. Post	0.034	−0.211**	0.208**	0.037**	−80.401	0.275	0.005
	(0.020)	(0.010)	(0.018)	(0.010)	(57.265)	(0.706)	(0.029)
HospitalExposure× 7–12 Mo. Post	0.023	−0.356**	0.362**	0.017	−315.642**	0.683	0.053
	(0.031)	(0.013)	(0.030)	(0.009)	(84.947)	(0.845)	(0.042)
N	1,064,893	1,064,893	1,064,893	1,064,893	1,064,893	1,064,893	1,064,893
R^2	0.281	0.621	0.300	0.426	0.453	0.370	0.658
Hospital-by-Event FEs	✓	✓	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓	✓	✓
Dep. Var. Pre-Mean	0.970	0.100	0.820	0.050	3,927.280	15.430	1.930
Ind. Var. Mean	0.230	0.230	0.230	0.230	0.230	0.230	0.230
Ind. Var. SD	0.200	0.200	0.200	0.200	0.200	0.200	0.200
Ind. Var. Range	[0, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]
# of Enforcement Events	16	16	16	16	16	16	16
# of Providers	3,720	3,720	3,720	3,720	3,720	3,720	3,720

Notes: * 5%, ** 1%. Table 4 reports the results of stacked difference-in-differences estimates comparing hospitals that were differentially reliant on the enforced facility for drug supply at baseline (13–18 months before enforcement; see equation 2). Outcomes include total drug volume (divided by baseline mean hospital volume), drug volume by type of origin facility, HHI of drug supply across manufacturing facilities, active ingredient substitution index, and number of supplying distributors. Non-enforced suppliers are categorized as existing (rather than new) supply relationships based on the presence of positive volume at baseline. Parentheses denote standard errors clustered at the hospital level.

from mechanically larger volume losses caused by FDA enforcement at the upstream manufacturing facility. Indeed, I find no statistically significant differences in total drug supply between more-exposed hospitals and their less-exposed counterparts (see Figure A.2), with 95% confidence intervals ruling out differential declines greater than 0.5% in 0–6 months and 4% in months 7–12 months after FDA enforcement.[40] Notably, since the DEA ARCOS data measures drug shipments delivered to hospitals rather than utilization, this stability of hospital drug volume is consistent with continuity of supply, not hospitals temporarily drawing down on inventories.

To examine mechanisms, columns 5–7 of Table 4 show how FDA regulatory enforcement impacts the structure of hospital procurement networks. The supplier Herfindahl-Hirschman Index (HHI), which measures concentration in hospital drug supply across manufacturing facilities, remains stable in the first six months after FDA enforcement before declining by over 300 HHI units (7.1% of the baseline mean) in the following six months. This reduction is consistent with more balanced hospital purchases across manufacturing facilities, which may help guard against future disruptions. Other margins, including substitution across active ingredients and changes in the number of supplying distributors, are suggestive of hospital responses but are not statistically significant, on average.[41]

[40] Interpreting point estimates as differences between more- and less-exposed hospitals is consistent with continuous difference-in-differences effectively comparing above- and below-average exposure units (Callaway et al., 2024).

[41] Substitution across active ingredients is measured using an index that tracks changes in the active ingredient

In isolation, a limitation of these results is that cross-sectional comparisons across differentially exposed hospitals may mask a general equilibrium "missing intercept," whereby reductions to drug supply may be evenly absorbed through redistribution between more- and less-exposed hospitals. Concerns about such general equilibrium effects are somewhat tempered by the results of an extension of the first stage analysis studying the dose-response relationship between the national market share of the enforced facility and aggregate nationwide drug volume from competing facilities (see Appendix Section B.2). The findings of this analysis, pictured in Appendix Figure B.3, suggest that enforcement at higher market share facilities naturally results in larger declines in implicated facility volume, but crucially also triggers greater increases in competing facility volume. That competitor ramp-up appears at the aggregate level suggests that the substitution to alternative manufacturing facilities found in this section is fueled by true volume increases rather than pure redistribution across hospitals. However, wide confidence intervals from this national-level analysis mean that it cannot gauge whether competitor ramp-up fully offsets disruptions at implicated facilities.

In summary, the absence of measurable differences in total drug supply between more- and less-exposed hospitals belies a rapid reorganization of procurement networks that shields hospitals from upstream manufacturing disruptions. Substitution occurs primarily along low-cost margins of adjustment—toward manufacturers already in hospitals' baseline supply networks and without differential changes in active ingredient composition—with familiar products likely decreasing adjustment costs to integrate alternative drugs. These patterns of substitution are enabled by hospitals' multi-sourced procurement operations, which provide natural avenues for substitution when a supplier is disrupted. Since single-sourcing is likely the most cost-effective but least resilient form of procurement, one interpretation of these findings is that hospitals may internalize supply chain risk and design procurement networks capable of agile substitution.

5.3 Heterogeneity in Drug Supply Resilience

Having shown that upstream FDA enforcement does not disrupt hospital drug supply, on average, I next examine heterogeneity along dimensions that shape the supply-side responses underlying this resilience. First, ramp-up by alternative manufacturers depends on the availability of excess production capacity, which may be limited when enforcement actions place greater stress on the supply chain. I therefore examine heterogeneity by the national market share of the enforced facility, which captures the amount of production capacity directly implicated, and

composition of a hospital's drug supply. For each SI opioid active ingredient, the change in share of hospital drug supply it accounts for is calculated between the current month and the same month in the previous year. The index is equal to the sum of the absolute values of these changes across active ingredients. Stability in ingredient composition leads an index value close to zero, with increases corresponding to changing composition.

by the degree of ongoing drug shortages at the time of enforcement, when manufacturers may already be operating near capacity to meet demand.

Second, I analyze factors influencing hospitals' capacity to substitute to alternative drugs, including the baseline diversification of their supply networks and institutional characteristics related to procurement sophistication, such as 340B Drug Pricing Program participation, teaching status, and hospital size. I also consider distributional factors that may lead to disadvantages in hospital resources and bargaining power, focusing on rural versus urban hospitals and on the demographic composition of their patient populations, including the shares of dual-eligible and Black, Hispanic, and Native American beneficiaries.

Figure A.3 summarizes these results, reporting both the baseline post-period treatment effect on hospital drug supply (i.e., overall treatment effects on drug supply net of the heterogeneity characteristic) and the heterogeneity interaction (i.e., deviation from baseline treatment effects based on the heterogeneity characteristic).[42]

With two important exceptions, discussed below, I find that hospitals are broadly shielded from declines in drug supply following FDA enforcement. For example, hospitals maintain drug supply regardless of the enforced facility's national market share, baseline procurement operations, and socioeconomic compositions of the patients they serve. Nevertheless, two contexts—rural hospitals and when enforcement occurs during an ongoing critical shortage—point to vulnerabilities in competitor ramp-up and hospital substitution when statistically significant declines in drug supply can arise. Importantly, declines in both of these settings emerge only after a delay rather than immediately after FDA enforcement, which may provide valuable lead time for hospitals to prepare for periods of diminished drug access.

Critical Shortages Hospital drug supply is more vulnerable when FDA enforcement occurs during a critical shortage, defined as periods when all three primary active ingredients of SI opioids are simultaneously in shortage in the month prior to the OAI inspection.[43] Declines in shipments to more-exposed hospitals, which emerge beginning two quarters after enforcement (see panel (a) of Figure A.4), are consistent with capacity constraints that limit upstream alternative manufacturers from sustaining elevated production for extended periods.

Indeed, the results show supply chains working hard during critical shortages to maintain

[42] Markers positioned next to estimated coefficients denote the statistical significance of the sum of the baseline and heterogeneity interaction estimates, indicating whether results are consistent with non-zero effects on hospital drug supply based on the heterogeneity characteristic.

[43] This definition of critical shortages is intended to capture times when the supply chain is already under significant strain at the time of FDA enforcement, as the supply chain is already grappling with shortages and there is no clear readily available active ingredient that hospitals can substitute to following a disruption. The results are qualitatively and quantitatively similar when defining critical shortages as two or more of the three active ingredients being in shortage as of the month prior to enforcement.

drug volume after FDA enforcement. Although such enforcement actions are associated with larger declines in implicated facility volume, on the order of 30 percentage points more than in other enforcement events, the decline is initially offset by increased supply from alternative manufacturers (panel (b) of Figure A.4). Two quarters after enforcement, however, supply from these alternative facilities falls, consistent with capacity constraints, while volume from the enforced facility remains depressed. Dynamic patterns also suggest that hospitals take more aggressive action to search for alternative supply during critical shortages, underscoring heightened supply chain strain (panels (c) and (d) of Figure A.4). At baseline, more-exposed hospitals do not adjust along the margins of connections to distributors or active ingredient substitution, possibly reflecting supply shortfalls being resolved by upstream ramp-up (see Table 4). During critical shortages, by contrast, more-exposed hospitals increase the number of distributors they source from (by 15% of the baseline in quarters 2 and 3 after enforcement) and also engage in greater substitution across active ingredients (by 30% in quarters 0 and 1 after enforcement).

Rural Hospitals More-exposed rural hospitals experience significant drug supply reductions in the second and third quarters after FDA enforcement, driven by losses in volume from manufacturing facilities in their baseline supply networks (see Panels (a) and (b) of Figure A.5).[44] A plausible explanation for these patterns lies in contracting dynamics and limited substitution capacity. Manufacturers and wholesalers typically allocate supply according to historical purchasing patterns, explaining the short-term resilience of rural hospitals (Conti and Wosińska, 2025; Wosińska and Frank, 2023).[45] Over time, however, allocations can be revised as hospitals vie for shipments against the backdrop of supply uncertainty and shortage risk (Park et al., 2024; Hantel et al., 2019), when rural hospitals may face disadvantages in sustaining substitution due to weaker bargaining power, fewer resources, and limited procurement flexibility.

Consistent with this interpretation, rural hospitals exhibit weaker substitution responses following FDA enforcement relative to their other more-exposed counterparts. While more-exposed hospitals generally respond by expanding the number of distributors they source from, rural hospitals tend to narrow their supply networks (see panel (c) of Figure A.5). Furthermore, rural hospitals display less substitution across the active ingredients comprising SI opioid drugs (see panel (d) of Figure A.5). Together, these patterns suggest limitations in procurement flexibility that leave rural hospitals more vulnerable to disruptions caused by FDA enforcement.

Overall, heterogeneity analysis reveals that while hospital drug supply is generally resilient to FDA enforcement, vulnerabilities to competitor ramp-up and hospital substitution emerge in rural settings and during critical shortages. Though the delayed emergence of these declines until

[44] These quarters are indexed to capture the periods spanning six to twelve months after FDA enforcement.

[45] In fact, the results suggest that rural hospitals actually see a short-term increase in drug supply immediately after enforcement, which is driven by maintaining greater drug volume from the enforced facility.

several months after enforcement may provide lead time to protect patients against adverse consequences, losses to drug supply carry clear scope to impact patient care. I therefore pay special attention to heterogeneity along these same dimensions in my analysis of patient outcomes.

6 Downstream Patient Outcomes

Finally, I estimate the causal effects of FDA regulatory enforcement on downstream patient outcomes, including health, quality of care, and healthcare costs. Whether upstream production disruptions caused by enforcement decrease, increase, or have no effect on hospital drug supply, their welfare implications depend on how they ultimately impact patient care.

Having shown that hospital drug supply is resilient, on average, there is limited scope for patient impacts through fluctuations in drug access outside of rural hospitals and during critical shortages. Instead, the most plausible channel is through the hospital responses generating resilience. Substitution may entail operational adjustments that increase the risk of errors, delay care, or lead to rationing as healthcare providers adapt to unfamiliar products (Park et al., 2025; Hantel et al., 2019).[46] On the other hand, recalling that FDA enforcement seeks to protect consumers from non-compliant products, patient outcomes may improve if substitution away from enforced facilities enhances the safety and quality of hospital drug supply.

6.1 Empirical Strategy

My empirical strategy estimates the impacts of FDA enforcement on the outcomes of patient hospital visits, which are the natural unit of observation for studying disruptions to SI opioids because they are administered on a one-off basis in acute hospital care. To ensure full visibility into patient outcomes and avoid contamination from previous hospitalizations, I focus on "index" visits: patients' first emergency department or inpatient claim within a 12-month period of continuous enrollment in Medicare fee-for-service Parts A and B.[47]

Using this unit of analysis, I implement a stacked difference-in-differences design that compares hospital visits before and after FDA enforcement, leveraging cross-sectional variation in the exposure of drug supply in the patient's local hospital environment. One of this design's key considerations is that enforcement may not only influence care within hospitals, but also which hospital patients visit. For example, hospitals navigating supply chain strain may adjust

[46] Medication errors are a particularly salient issue in this setting because Weiss and Elixhauser (2013) found that, in 2011, opiates and narcotics were the third leading cause of hospital adverse drug events (injuries due to drugs being prescribed or administered improperly).

[47] This construction leads to the exclusion of repeat hospitalizations from the set of index visits, though 30-day readmissions are studied as an outcome variable. Index visits also exclude patients with frequent hospitalizations ("frequent flyers"), who may not have a 12-month gap in their hospitalizations.

admissions or discharge rules according to their capacity to provide care, potentially generating variation in travel costs, quality of care, and patient-provider familiarity.

Accordingly, index visit exposure is constructed by using the empirical distribution of hospital visits in the pre-period (January 2006–June 2007) as a data-driven predictor of the hospitals that patients would have visited in the absence of FDA enforcement.[48] For index hospital visit i by a patient residing in ZIP code z, predicted exposure to enforcement event τ is

$$\mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau} = \sum_{H} s^{pre}_{z(i),h} \times \text{HospitalExposure}_{h,\tau}$$

where $s^{pre}_{z(i),h}$ is the fixed share of visits by patients residing in ZIP code $z(i)$ to hospital h in the pre-period, and $\text{HospitalExposure}_{h,\tau}$ is hospital h's share of baseline drug supply from the facility enforced in event τ. Intuitively, this construction reflects the degree to which hospitals in a patient's local environment happened to rely on the enforced facility for drug supply at baseline, assuming that patients continue visiting hospitals following pre-period patterns.[49]

The advantage of $s^{pre}_{z(i),h}$, compared to approaches like geographic proximity instruments (e.g., Antill et al. (2025) and Einav et al. (2025)), is that realized hospital visits embody the full set of factors—quality, capacity, specialization, referral networks, distance—that endogenously shape hospital market shares. This endogeneity does not threaten identification because visit probabilities are pre-determined with respect to FDA enforcement, and Appendix D.2 demonstrates that the pre-period hospital visit distribution and the predicted hospital exposure it generates are highly predictive of realized hospital visits and hospital exposure, respectively.[50]

I estimate the following specification under the identifying assumption that the outcomes of index visits with different predicted hospital exposure would have evolved along parallel trends in the absence of enforcement

$$\begin{aligned} Y_{i,t,\tau} = \sum_{r=-6, r\neq -1}^{12} \beta_r \, \mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau} \times \mathbf{1}\left\{r \text{ Months Since Enforcement}\right\}_{t,\tau} \\ + \delta_{z(i),\tau} + \eta_{r,\tau} + X_{z(i),t,-\tau} + W_i + \varepsilon_{i,t,\tau} \end{aligned} \tag{3}$$

[48] Conditioning on the actual hospital visited would risk understating negative effects if, for example, patients who would have normally visited a high-exposure hospital instead shifted to a low-exposure hospital. While FDA enforcement may seem unlikely to affect hospital choice—since sterile injectable drugs represent only one component of hospital services and their availability may not be sufficient to impact hospitals' admissions decisions alone—I allow this to be part of the causal mechanism.

[49] The measure also resembles the spirit of a shift-share design: the "shifts" are hospital-specific exposure to upstream FDA enforcement and the "shares" are the fixed pre-period hospital visit probabilities.

[50] A first stage regression of exposure at the hospital actually visited on (i) predicted exposure and (ii) the national market share of the upstream implicated facility indicates that every percentage point of predicted exposure is associated with a highly statistically significant 0.87pp increase in actual exposure. By comparison, a 1pp increase in the national market share of the implicated facility, capturing the overall exposure of hospital drug volume to enforcement, is associated with only a 0.13pp increase in actual exposure.

where $Y_{i,t,\tau}$ is the outcome of index hospital visit i in month t around FDA OAI inspection event τ.[51] Two-way fixed effects include patient ZIP code-by-event fixed effects $\delta_{z(i),\tau}$ and event-by-event time fixed effects $\eta_{r,\tau}$. The vector $X_{z(i),t,-\tau}$ controls for event time-by-predicted exposure interactions from overlapping FDA enforcement events $-\tau$, and W_i are primary diagnosis fixed effects. Standard errors are clustered at the ZIP code level, where treatment is assigned.

To focus on index visits most clinically relevant to SI opioid disruptions, I use pre-period data to train a cross-validated LASSO machine learning model on patient demographics (age, age squared, sex, and race) and primary diagnosis to predict whether a hospital visit involved anesthesia or operating room use, a likely proxy of SI opioid utilization informed by input from an emergency physician. Index visits in the top quartile of predicted probability serve as the regression sample, capturing the highest ex-ante likelihood of utilization, while those in the bottom quartile provide a placebo group unlikely to be directly impacted by the disruptions. Appendix D.1 details the prediction exercise and documents substantial differences in the severity and hospital characteristics of top- and bottom-quartile visits.

As robustness checks, I replicate the analysis using the full sample of index visits and by constructing predicted hospital exposure using pre-period hospital visit distributions specific to the top predicted SI opioid utilization quartile. To address small cell sizes, I implement an Empirical Bayes procedure that shrinks quartile-specific ZIP code hospital visit probabilities toward the distribution of top-quartile patients in the same Hospital Service Area.

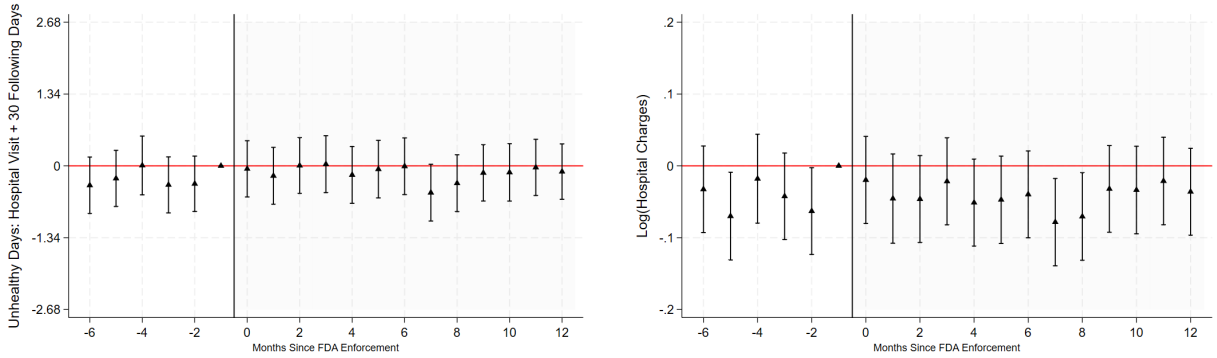
6.2 Results

Like the analysis of drug supply, the patient-level results highlight hospitals' resilience to upstream FDA enforcement. Across a range of outcomes, I find no detectable harms to patient health, quality of care, or healthcare costs following enforcement. Suggestive evidence of temporary shifts toward closer and more commonly used hospitals points to changes in the types of hospitals visited as a margin of adjustment that may contribute to this resilience.

I begin by examining how enforcement impacts patient health and quality of care (see Table A.1). My preferred outcome is unhealthy days away from home (hereafter, unhealthy days), a quality-of-care measure developed in conjunction with CMS and MedPAC that counts the number of days in the 30 after hospital arrival that the patient spends in a healthcare facility (e.g., hospital or skilled nursing facility) or is deceased (Burke et al., 2020, 2023). The advantage of this measure is that it aggregates outcomes across a broad range of healthcare settings, capturing the net effects of initial length of stay, readmissions, post-discharge care, and mortality.

[51] The (unstacked) two-way fixed effects analog, which removes event indices τ from the fixed effects and treatment variable, is $Y_{i,t} = \sum_{r=-6, r\neq -1}^{12} \beta_r \, \mathbb{E}\left[\text{HospitalExposure}\right]_{z(i)} \times \mathbf{1}\left\{r \text{ Months Since Enforcement}\right\}_t + \delta_{z(i)} + \eta_r + W_i + \varepsilon_{i,t}$. This specification could be used to study each of the 16 FDA enforcement events in isolation, setting aside concerns about overlapping treatment.

Event study estimates plotted in Figure 5 show that differences in unhealthy days between more- and less-exposed index hospital visits remain stable throughout the event horizon: more-exposed visits experience a 0.15-day rise in unhealthy days (1.4% of the 10.7-day baseline mean) in the first six months after FDA enforcement and a 0.01-day rise in the following six months (0.1% of the baseline). Corresponding 95% confidence intervals rule out increases of more than 3.4% and 2.2% over these periods, respectively.



(a) Unhealthy Days **(b)** Log(Hospital Charges)

Figure 5. Impact of FDA OAI Inspections on Patient Health

Notes: Figure 5 plots the evolution of patient health outcomes around FDA OAI inspections. The dynamic difference-in-differences estimates corresponding to $\hat{\beta}_r$ from equation 3 compares the trajectories of hospital visits by patients residing in ZIP codes with local hospital environments that were differentially reliant on the implicated upstream manufacturing facility at baseline for their drug supply. The treatment variable used to make this comparison, $\mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau}$, is a convex combination of hospital-level exposure to FDA enforcement event τ and patient ZIP code-level hospital visit probabilities, which are based on the fixed empirical distribution of hospital visits in the pre-study period. Error bands represent 95% confidence intervals with standard errors clustered at the patient ZIP code level.

Though less precisely estimated owing to their low frequency, small and statistically insignificant point estimates for 90-day all-cause mortality and 30-day readmissions are also consistent with no systematic deterioration in patient outcomes. Using the full sample of index visits provides additional precision. The 95% confidence intervals rule out increases in 90-day mortality of more than 4.9% and 6.2% in months 0–6 and 7–12 after enforcement, respectively, and increases of more than 1.8% and 1.5% in 30-day readmissions over the same periods.

These patterns are robust across a range of alternative specifications. Specifically, there is no detectable evidence of declines in patient health or quality of care when (i) constructing predicted exposure with quartile-specific hospital visit distributions (see Appendix Table D.1), (ii) using the exposure of the hospital actually visited (see Appendix Table D.2), and (iii) when analyzing the full set of index hospital visits rather than restricting to the top quartile most likely to be treated with SI opioids within standard care (see Appendix Tables D.3 and D.4).

Next, I find that FDA enforcement has no detectable impacts on the logarithm of hospital

charges or Medicare reimbursements associated with index visits. These null findings may appear surprising given the operational adjustments undertaken by more-exposed hospitals in response to the upstream supply disruption, most notably reorganizing their procurement networks. However, a limitation of this analysis is that claims-based cost measures may not be well-suited to fully capturing these adjustments. Hospital charges reflect list prices and so embody changes in the intensity or complexity of billed services. As such, they may not reflect staffing, administration, or drug acquisition expenses that are not specific to a single hospital visit. Nevertheless, to the extent that hospitals do in fact bear such adjustment costs, null effects on reimbursement indicate that they are not directly passed on to the Medicare program or its beneficiaries.

Although patient outcomes and costs are unaffected, there is suggestive evidence of temporary shifts to the characteristics of hospitals visited after FDA enforcement. At baseline, patients in the top quartile of predicted SI opioid utilization disproportionately receive care from more distant and higher-resource hospitals: they are less likely to visit their geographically closest or a rural hospital and are more likely to be admitted to teaching centers (Appendix Table D.7).[52] These patterns are modestly attenuated following FDA enforcement among patients in more-exposed ZIP codes. In the first six months after enforcement, point estimates suggest that such hospital visits are 1.5pp more likely to occur at the patient's nearest hospital (a 4% increase relative to baseline), 1.2pp less likely to occur at a teaching center (a 2.4% decline), 0.3pp more likely to occur at a rural hospital (a 2.4% increase), and involve 2% shorter travel distances (see Table A.2). Event study estimates indicate that these effects are concentrated between two and six months after enforcement and dissipate afterward (see Appendix Figure D.1).

Interestingly, these shifts do not reflect avoidance of high-exposure hospitals: point estimates imply that index visits in ZIP codes with hospitals that were fully reliant on the implicated facility for baseline drug supply shift to hospitals with only 0.3–0.7pp lower exposure. Instead, patients appear to be directed toward hospitals with higher baseline visit probabilities. One possible explanation for these patterns is that enforcement temporarily strains hospital environments, prompting adjustments to capacity, scheduling, or referral patterns that may limit patients' willingness or ability to travel to more distant hospitals. Placebo estimates for bottom-quartile visits show smaller and statistically insignificant effects (see Appendix Table D.6), reinforcing the interpretation that the hospital visited is a margin of adjustment limited to patients most affected by SI opioid disruptions.

In summary, FDA enforcement does not cause detectable declines in patient health, quality of care, or healthcare costs. These conclusions are robust across index visit samples and exposure definitions, with confidence intervals ruling out large adverse impacts. Since these findings

[52]Distance is measured as the geodesic distance between the centroid of the patient's ZIP code and hospital latitude and longitude.

reflect the net effects of potential harms from supply disruptions on one hand, and benefits from potential improvements in the safety and quality of drug supply following substitution on the other, it is not clear whether both forces are muted or influence patient outcomes in an offsetting manner. Nevertheless, the results underscore hospitals' resilience to upstream FDA regulatory enforcement: they not only rapidly reshuffle their procurement networks to maintain supply, but also manage to do so while maintaining the quality and cost of care they provide.

6.3 Heterogeneity Analysis

The hospital-level analysis in Section 5 identifies rural hospitals and FDA enforcement taken during critical shortages as high-risk settings where drug supply is vulnerable to declines. To test whether resilience in patient outcomes also holds in these settings, which most threaten patient care due to potential losses in drug access, I examine heterogeneity across these dimensions on health, quality of care, and cost outcomes. For completeness, I also assess heterogeneity by the national market share held by the implicated manufacturing facility, capturing the magnitude of supply chain strain created by enforcement, as well as along patient characteristics linked to distributional concerns, including dual-eligibility for Medicare and Medicaid, sex, and among Black, Hispanic, and Native American beneficiaries.

Figure A.6 presents the results, reporting the baseline post-enforcement treatment effects on patient outcomes (i.e., overall treatment effects on patient outcomes net of the heterogeneity characteristic) and the heterogeneity interaction (i.e., deviation from baseline treatment effects based on the heterogeneity characteristic).[53] [54]

Across nearly all dimensions, the findings reinforce the conclusion that hospitals effectively shield patients from potential adverse consequences of upstream FDA regulatory enforcement. Even under high-risk conditions when hospital drug supply is most vulnerable to decline—among rural patients and during critical shortages—I find small and statistically insignificant impacts on all patient outcomes studied, including unhealthy days, 90-day mortality, 30-day readmissions, and healthcare costs. Similar patterns of robustness hold across the other dimensions of heterogeneity studied (see Figure A.6).[55]

[53] Markers positioned next to estimated coefficients denote the statistical significance of the sum of the baseline and heterogeneity interaction estimates, indicating whether results are consistent with non-zero effects on patient outcomes based on the heterogeneity characteristic.

[54] Like the baseline regression specification in the main analysis of patient outcomes, heterogeneity specifications use the sample of index visits in the top quartile of predicted sterile injectable opioid utilization.

[55] Two isolated findings suggesting the existence of heterogeneous effects merit mention. First, I find a reduction in healthcare costs specific to dual-eligible beneficiaries, appearing as reductions to both the logarithm of hospital charges and reimbursements. The implications of this finding are ambiguous, potentially reflecting modest improvements to care or, given that dual-eligible beneficiaries may bring lower reimbursement than other Medicare beneficiaries (Keohane and Hwang, 2022), reductions in services provided. However, because the cost reductions are not accompanied by changes to unhealthy days, mortality, or readmissions, the evidence suggests no impacts

7 Conclusion

This paper links detailed supply chain data from the DEA ARCOS database with Medicare claims to causally trace how FDA regulatory enforcement against drug manufacturing facilities propagates to downstream hospitals and patients. Despite large disruptions at the enforced facility, the feared fragility of SI supply chains, and evidence of supply chain propagation in other settings, I show that the supply of SI opioid drugs is resilient to FDA enforcement. My findings therefore suggest that, in this context, supply chain concentration has not meaningfully complicated the FDA's trade-off between drug access and safety. Owing to supply-side responses by other elements of the supply chain, including manufacturer ramp-up and hospital substitution, the short-term reductions to drug access that might accompany long-term safety improvements as enforced facilities implement corrective actions do not appear to have materialized.

These results plausibly generalize to the broader category of SI drugs, which share the conditions of concentrated production, recurrent quality problems, and frequent shortages. Given that resilience in this setting is driven by substitution, a natural direction for future research is to examine drug classes with limited substitutability, such as breakthrough branded therapies or specialty drugs. A key empirical question is whether the absence of close substitutes heightens the risk of downstream propagation for these drugs, or whether the potential loss of high-margin sales incentivizes manufacturers to invest in redundancies and more robust production.

A second avenue for future work is to comprehensively assess the impacts of FDA enforcement on hospital costs. While null effects on reimbursement suggest little direct impact on spending borne by the Medicare program or its beneficiaries, claims data do not capture the full set of adjustment costs hospitals may incur when incorporating substitutes. Whether arising from increased drug acquisition expenses, labor needs, or inventory management, these costs represent a policy-relevant parameter for fully understanding the implications of FDA regulatory enforcement. By clarifying the conditions under which enforcement risks compromising short-run outcomes, this paper and the proposed directions for future work can help contribute to a more complete understanding of the FDA's access-safety trade-off in concentrated supply chains.

on care. Second, estimates suggest a possible rise in unhealthy days among female patients. However, the effect is not statistically significant when combined with the baseline treatment effect, and event studies reveal no clear dynamics consistent with a deterioration in outcomes.

References

Acemoglu, Daron and Alireza Tahbaz-Salehi (2025) "The Macroeconomics of Supply Chain Disruptions," *Review of Economic Studies*, 92 (2), 656–695.

Aguero, David and Delia Allen (2024) "Weathering the storm: commentary on the Hurricane Helene IV fluid shortage," *The Journal of Pediatric Pharmacology and Therapeutics*, 29 (6), 667–669.

Alpert, Abby and Mireille Jacobson (2019) "Impact of oncology drug shortages on chemotherapy treatment," *Clinical Pharmacology & Therapeutics*, 106 (2), 415–421.

American Regent, Inc. (2011) "Important Drug Information Notice Regarding Potassium Phosphates Injection, USP, 3 mM/mL Phosphorous; 4. 4 mEq/mL Potassium: 5 mL, 15 mL and 50 mL Single Dose Vials," https://www.prnewswire.com/news-releases/important-drug-information-notice-regarding-potassium-phosphates-injection-usp-3-html.

Antill, Samuel, Jessica Bai, Ashvin Gandhi, and Adrienne Sabety (2025) "Healthcare Provider Bankruptcies," technical report, National Bureau of Economic Research.

ASHP (2018) "Injectable Opioid Survey Report," technical report, American Society of Health-System Pharmacists, https://www.ashp.org/-/media/assets/drug-shortages/docs/drug-shortages-svp-opioid-survey-report.pdf, Accessed September 18, 2025.

——— (2024) "Severity and Impact of Sterile Fluid Shortages," technical report, American Society of Health-System Pharmacists, https://news.ashp.org/-/media/assets/drug-shortages/docs/2024/ASHP-Sterile-Fluid-Shortage.pdf, Accessed September 18, 2025.

Autor, David H, David Dorn, and Gordon H Hanson (2016) "The China shock: Learning from labor-market adjustment to large changes in trade," *Annual review of economics*, 8 (1), 205–240.

Balboni, Clare, Johannes Boehm, and Mazhar Waseem (2023) "Firm adaptation and production networks: Structural evidence from extreme weather events in Pakistan."

Baldwin, Richard, Rebecca Freeman, and Angelos Theodorakopoulos (2023) "Hidden exposure: Measuring US supply chain reliance," *Brookings Papers on Economic Activity*, 2023 (2), 79–167.

Barrot, Jean-Noël and Julien Sauvagnat (2016) "Input specificity and the propagation of idiosyncratic shocks in production networks," *The Quarterly Journal of Economics*, 131 (3), 1543–1592.

Boehm, Christoph E, Aaron Flaaen, and Nitya Pandalai-Nayar (2019) "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake," *Review of Economics and Statistics*, 101 (1), 60–75.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2022) "Quasi-experimental shift-share research designs," *The Review of economic studies*, 89 (1), 181–213.

——— (2025) "A practical guide to shift-share instruments," *Journal of Economic Perspectives*, 39 (1), 181–204.

Budish, Eric, Benjamin N Roin, and Heidi Williams (2015) "Do firms underinvest in long-term research? Evidence from cancer clinical trials," *American Economic Review*, 105 (7), 2044–2085.

Burke, Laura G, Ryan C Burke, Ciara E Duggan, Jose F Figueroa, E John Orav, and Edward R Marcantonio (2023) "Trends in healthy days at home for Medicare beneficiaries using the emergency department," *Journal of the American Geriatrics Society*, 71 (10), 3122–3133.

Burke, Laura G, E John Orav, Jie Zheng, and Ashish K Jha (2020) "Healthy days at home: a novel population-based outcome measure," in *Healthcare*, Vol. 8 (1), 100378, Elsevier.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna (2024) "Difference-in-differences with a continuous treatment," technical report, National Bureau of Economic Research.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller (2008) "Bootstrap-based improvements for inference with clustered errors," *The review of economics and statistics*, 90 (3), 414–427.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi (2021) "Supply chain disruptions: Evidence from the great east japan earthquake," *The Quarterly Journal of Economics*, 136 (2), 1255–1321.

Castro-Vincenzi, Juanma, Gaurav Khanna, Nicolas Morales, and Nitya Pandalai-Nayar (2024) "Weathering the storm: Supply chains and climate risk," technical report, National Bureau of Economic Research.

Conti, Rena M and Marta E Wosińska (2025) "The Economics of Generic Drug Shortages: The Limits of Competition," *Journal of Economic Perspectives*, 39 (2), 79–102.

Coyne, David, Itzik Fadlon, Shanthi P Ramnath, and Patricia K Tong (2024) "Household Labor Supply and the Value of Social Security Survivors Benefits," *American Economic Review*, 114 (5), 1248–1280.

CRS (2018) "Drug Shortages: Causes, FDA Authority, and Policy Options," https://crsreports.congress.gov/product/pdf/IF/IF11058.

——— (2023) "The Federal Role in Addressing the Adderall Drug Shortage," https://crsreports.congress.gov/product/pdf/IN/IN12156.

DEA (1997) "ARCOS Registrant Handbook."

Deryugina, Tatyana, Jonathan Gruber, and Adrienne Sabety (2020) "Natural disasters and elective medical services: How big is the bounce-back?", technical report, National Bureau of Economic Research.

Donahoe, J Travis (2024) "Supplier Enforcement and the Opioid Crisis," *Available at SSRN 5260245*.

Einav, Liran, Amy Finkelstein, and Neale Mahoney (2025) "Producing health: measuring value added of nursing homes," *Econometrica*, 93 (4), 1225–1264.

Fadlon, Itzik and Jessica Van Parys (2020) "Primary care physician practice styles and patient care: Evidence from physician exits in Medicare," *Journal of health economics*, 71, 102304.

FDA (2010a) "Establishment Inspection Report: Hospira, Inc.," Establishment Inspection Report, January, Summary of FDA inspection conducted in January 2010 at Hospira, Inc. Facility in Rocky Mount, NC.

——— (2010b) "FDA Form 483: Hospira, Inc.," FDA Form 483, January, FDA Form 483 Issued after FDA inspection conducted in January 2010 at Hospira, Inc. Facility in Rocky Mount, NC.

——— (2016) "FDA Form 483: Hospira, Inc.," FDA Form 483, May, FDA Form 483 Issued after FDA inspections conducted in May and June 2016 at Hospira, Inc. Facility in McPherson, KS.

FDA (2018a) "Statement from Douglas Throckmorton, M.D., deputy center director for regulatory programs in FDA's Center for Drug Evaluation and Research, on the agency's response to ongoing drug shortages for critical products," https://www.biospace.com/statement-from-douglas-throckmorton-m-d-deputy-center-director-for-regulatory-pro

FDA (2018b) "FDA Form 483: Hospira, Inc.," FDA Form 483, July, FDA Form 483 Issued after FDA inspections conducted in July and August 2018 at Hospira, Inc. Facility in McPherson, KS.

FDA (2020a) "Drug and Biologic Essential Medicines, Medical Countermeasures, and Critical Inputs for the List Described in Section 3(c) of the Executive Order 13944."

——— (2020b) "Drug Shortages: Root Causes and Potential Solutions," https://www.fda.gov/media/131130/download?attachment.

——— (2023) "Inspection Basics," https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/inspection-basics.

——— (2024a) "FDA's Risk-Based Approach to Inspections," https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/inspection-basics/fdas-risk-based-approach-inspections.

——— (2024b) "Understanding FDA Inspections and Data," https://www.fda.gov/drugs/news-events-human-drugs/understanding-fda-inspections-and-data-09062023.

FTC (2024) "Market Factors Relevant to Infant Formula Supply Disruptions," https://www.ftc.gov/system/files/ftc_gov/pdf/infant_formula_report_final.pdf, Accessed: 2024-10-14.

Galdin, Anais (2025) "Resilience of Global Supply Chains and Generic Drug Shortages."

Grennan, Matthew and Robert J Town (2020) "Regulating innovation with uncertain quality: information, risk, and access in medical devices," *American Economic Review*, 110 (1), 120–161.

Gruber, Jonathan and Samuel A Kleiner (2012) "Do strikes kill? Evidence from New York state," *American Economic Journal: Economic Policy*, 4 (1), 127–157.

Hantel, Andrew, Mark Siegler, Fay Hlubocky, Kevin Colgan, and Christopher K Daugherty (2019) "Prevalence and severity of rationing during drug shortages: a national survey of health system pharmacists," *JAMA internal medicine*, 179 (5), 710–711.

Holbrook, Emily (2013) "A world of supply chain risks," *Risk Management*, 60 (3), 6–8.

Hospira, Inc. (2018) "Temporary Importation of Hydromorphone Hydrochloride Injection, USP 2 mg/mL Ampules from Canada to Address U.S. Drug Shortage," Hospira, Inc. Communication to Healthcare Providers, August, Dear Healthcare Provider Letter.

Keohane, LM and Ann Hwang (2022) "Payment policy and the challenges of Medicare and Medicaid integration for dual-eligible beneficiaries," *Health Affairs Health Policy Brief*.

Korovkin, Vasily, Alexey Makarin, and Yuhei Miyauchi (2024) "Supply Chain Disruption and Reorganization: Theory and Evidence from Ukraine's War," *Available at SSRN 4825542*.

Kweder, Sandra L and Susie Dill (2013) "Drug shortages: the cycle of quantity and quality," *Clinical Pharmacology & Therapeutics*, 93 (3), 245–251.

Kwon, Spencer Y, Yueran Ma, and Kaspar Zimmermann (2024) "100 years of rising corporate concentration," *American Economic Review*, 114 (7), 2111–2140.

Lahue, Betsy J, Bruce Pyenson, Kosuke Iwasaki, Helen E Blumen, Susan Forray, and Jeffrey M Rothschild (2012) "National burden of preventable adverse drug events associated with inpatient injectable medications: healthcare and medical professional liability costs," *American health & drug benefits*, 5 (7), 1.

Lam, Miranda B, Kristen E Riley, Jie Zheng, E John Orav, Ashish K Jha, and Laura G Burke (2021) "Healthy days at home: a population-based quality measure for cancer patients at the end of life," *Cancer*, 127 (22), 4249–4257.

Liebman, Eli, Emily C Lawler, Abe Dunn, and David B Ridley (2023) "Consequences of a shortage and rationing: evidence from a pediatric vaccine," *Journal of Health Economics*, 92, 102819.

Liu, Ernest, Vladimir Smirnyagin, and Aleh Tsyvinski (2024) "Supply Chain Disruptions and Supplier Capital in US Firms."

Noh, In Joon, John Gray, George Ball, Zachary Wright, and Hyunwoo Park (2025) "Are all generic drugs created equal? An empirical analysis of generic drug manufacturing location and serious drug adverse events," *Production and Operations Management*, 10591478251319691.

Park, Minje, Anita L Carson, and Rena M Conti (2025) "Linking Medication Errors to Drug Shortages: Evidence from Heparin Supply Chain Disruptions Caused by Hurricane Maria," *Manufacturing & Service Operations Management*.

Park, Minje, Anita L Carson, Erin R Fox, and Rena M Conti (2024) "Stockpiling at the onset of the covid-19 pandemic: An empirical analysis of national prescription drug sales and prices," *Management Science*, 70 (10), 6483–6501.

Park, Minje, Rena M Conti, Marta E Wosińska, Ergun Ozlem, Wallace J Hopp, and Erin R Fox (2023) "Building resilience into US prescription drug supply chains," *Health Affairs Forefront*.

Peltzman, Sam (1973) "An evaluation of consumer protection legislation: the 1962 drug amendments," *Journal of Political Economy*, 81 (5), 1049–1091.

Philipson, Tomas J and Eric Sun (2008) "Is the Food and Drug Administration safe and effective?" *Journal of Economic Perspectives*, 22 (1), 85–102.

Rogers, Parker (2023) "Regulating the Innovators: Approval Costs and Innovation in Medical Technologies."

Roth, Jonathan, Pedro HC Sant'Anna, Alyssa Bilinski, and John Poe (2023) "What's trending in difference-in-differences? A synthesis of the recent econometrics literature," *Journal of Econometrics*, 235 (2), 2218–2244.

Serchen, Josh, David Hilden, Jenny R Silberger, Health, and Public Policy Committee of the American College of Physicians (2025) "Bolstering the Medication Supply Chain and Ameliorating Medication Shortages: A Position Paper From the American College of Physicians," *Annals of Internal Medicine*.

Soliman, Adam (2025) "Disrupting Drug Markets: The Effects of Crackdowns on Rogue Opioid Suppliers," *American Economic Journal: Economic Policy*, 17 (4), 165–91, 10.1257/pol.20230640.

Starc, Amanda and Thomas G Wollmann (2025) "Does entry remedy collusion? Evidence from the generic prescription drug cartel," *American Economic Review*, 115 (5), 1400–1438.

Stoye, George and Max Warner (2023) "The effects of doctor strikes on patient outcomes: evidence from the English NHS," *Journal of Economic Behavior & Organization*, 212, 689–707.

Strahan, Philip E (2013) "Too big to fail: Causes, consequences, and policy responses," *Annu. Rev. Financ. Econ.* 5 (1), 43–61.

Vo, Peter, Daniel A Sylvia, Loay Milibari et al. (2021) "Management of a parenteral opioid shortage using ASHP guidelines," *American Journal of Health-System Pharmacy*, 78 (5), 426–435.

Weiss, Audrey J. and Anne Elixhauser (2013) "Characteristics of Adverse Drug Events Originating During the Hospital Stay, 2011," hcup statistical brief #164, Agency for Healthcare Research and Quality, https://hcup-us.ahrq.gov/reports/statbriefs/sb164.jsp, Healthcare Cost and Utilization Project (HCUP).

Wing, Coady, Seth M Freedman, and Alex Hollingsworth (2024) "Stacked Difference-in-Differences," Working Paper32054, National Bureau of Economic Research, 10.3386/w32054.

Woodcock, Janet and Marta Wosinska (2013) "Economic and technological drivers of generic sterile injectable drug shortages," *Clinical Pharmacology & Therapeutics*, 93 (2), 170–176.

Wosińska, Marta E and Richard G Frank (2023) "Federal Policies to Address Persistent Generic Drug Shortages."

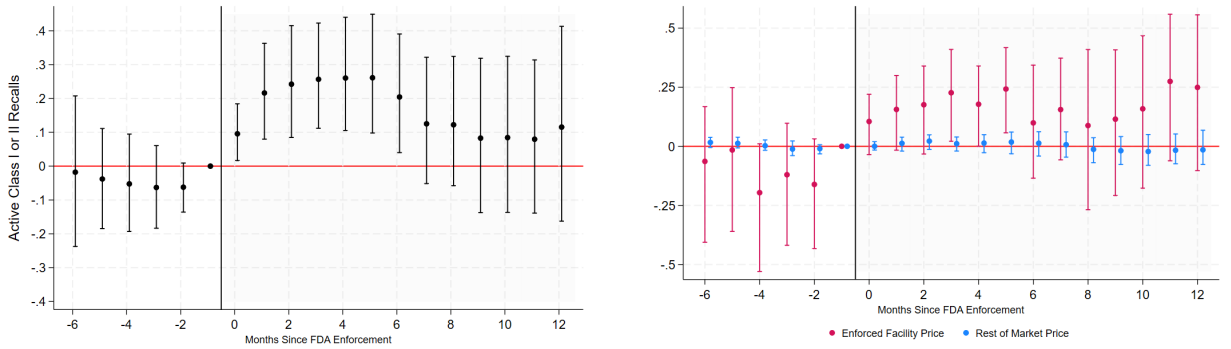
Wright, Zachary, John Gray, George Ball, and In Joon Noh (2025) "Preannounced Regulatory Inspections: FDA Oversight and Drug Quality Risk," *Available at SSRN 5252874*.

Wu, Anqi Angie and Yixin Iris Wang (2023) "The more monitoring, the better quality? Empirical evidence from the generic drug industry," technical report.

Yurukoglu, Ali, Eli Liebman, and David B Ridley (2017) "The role of government reimbursement in drug shortages," *American Economic Journal: Economic Policy*, 9 (2), 348–382.

Ziedan, Engy, Kosali I Simon, and Coady Wing (2022) "Mortality effects of healthcare supply shocks: evidence using linked deaths and electronic health records," technical report, National Bureau of Economic Research.

A Supplementary Tables and Figures



(a) Active Class 1 and 2 Recalls **(b)** Log Enforced Facility and Rest of Market Drug Prices

Figure A.1. Impact of FDA OAI Inspections on Secondary Outcomes

Notes: Figure A.1 plots the differential evolution of manufacturing facility-by-opioid market outcomes between enforced manufacturing facilities that matched controls that were enforced 1–3 years later and serve opioid drug markets that do not overlap with those of the enforced facility. Error bands represent 90% confidence intervals based on standard errors calculated using the wild cluster bootstrap with manufacturing facility-level clusters. Panel (a) plots the evolution of active class I and II recalls of drugs originating from each manufacturing facility. Class I and II recalls refer to recalls where use of the drug has the potential to cause adverse health consequences, and recalls are counted as active between their initiation and termination months. Panel (b) plots the evolution of log drug prices, including those of drugs originating from the implicated facility and of drugs produced by other manufacturing facilities the same opioid market. Drug prices are calculated by the sum of all Medicare Part D prescription spending divided by milligrams of morphine dispensed from those prescriptions.

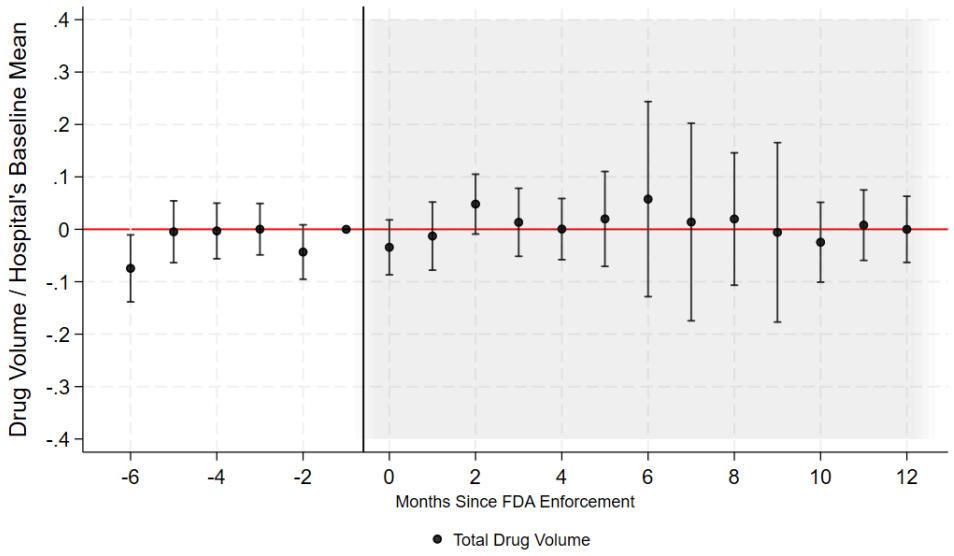


Figure A.2. Impact of FDA Enforcement on Total Hospital Drug Supply

Notes: Figure A.2 plots the evolution of hospital sterile injectable opioid drug volume around FDA OAI inspections. The dynamic difference-in-differences estimates corresponding to $\hat{\beta}_r$ from equation 2 compares the trajectories of hospitals differentially reliant on the enforced upstream manufacturing facility at baseline for drug supply. The figure plots the results when using total drug volume as the outcome variable. Error bands represent 95% confidence intervals with standard errors clustered at the hospital level.

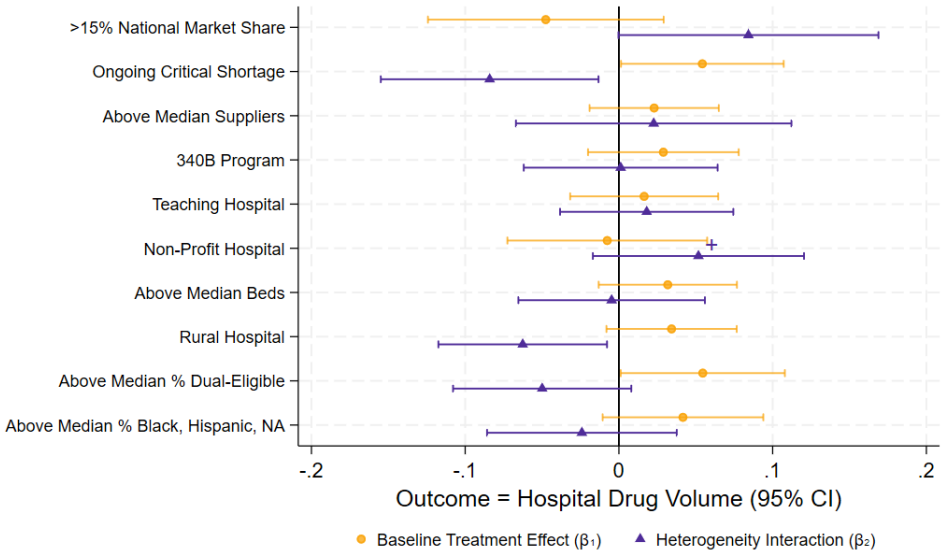
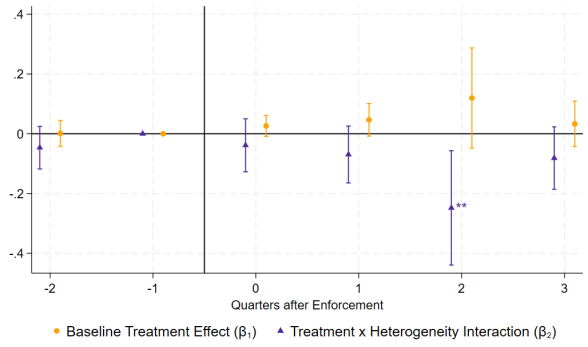


Figure A.3. Heterogeneity in Resilience of Hospital Drug Supply

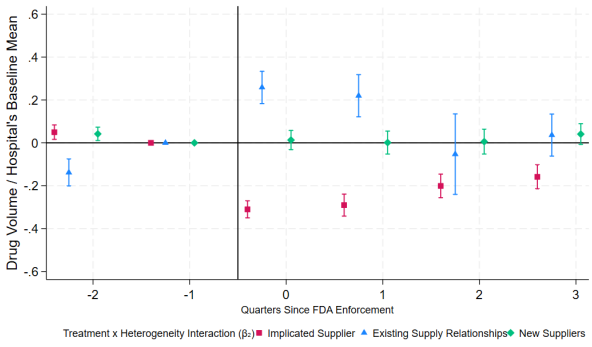
Notes: Figure A.3 examines heterogeneity in the impacts of FDA regulatory enforcement on hospital drug supply based on the degree of stress imposed by the enforcement action on the supply chain, characteristics of hospital procurement operations, and by distributional concerns. All estimates are from difference-in-difference regression equations with the form:

$$Y_{h,t,\tau} = \beta_1 \text{Post}_{t,\tau} \times \text{HospitalExposure}_{h,\tau} + \beta_2 \text{Post}_{t,\tau} \times \text{HospitalExposure}_{h,\tau} \times \text{Heterogeneity}_{h,\tau} + \delta_{h,\tau} + \eta_{r,\tau} + X_{h,t,-\tau} + \varepsilon_{h,t,\tau}$$

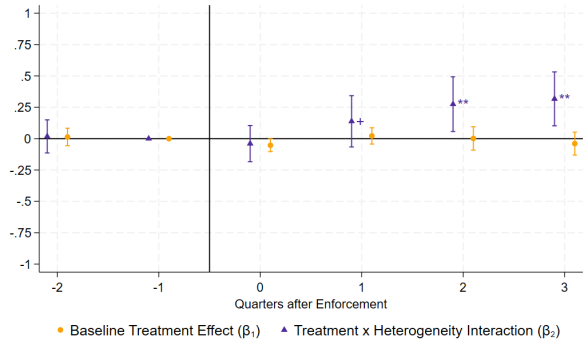
where β_1 represents the baseline post-period treatment effect of FDA enforcement (plotted in orange with circular markers) and β_2 represents the heterogeneity interaction capturing the deviation in treatment effects by each heterogeneity characteristic (plotted in purple with triangular markers). Symbols plotted in the figure denote the statistical significance of the test $\beta_1 + \beta_2 = 0$, indicating whether the baseline treatment effect and heterogeneity effect are collectively different from zero: + 10% * 5% ** 1% *** 0.1%.



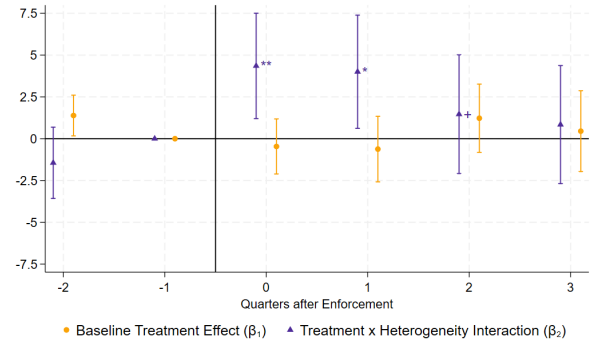
(a) Total Hospital Drug Supply



(b) Drug Supply Decomposed by Supplier Type



(c) Supplying Distributors



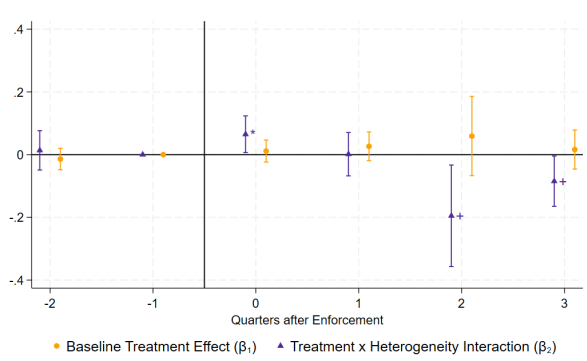
(d) Active Ingredient Substitution Index

Figure A.4. Heterogeneity in Hospital Drug Supply During Critical Shortages

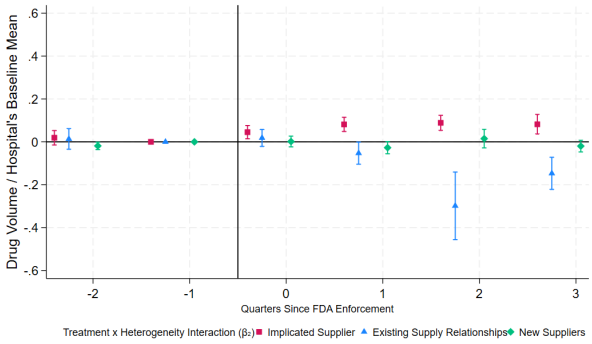
Notes: Figure A.4 studies the dynamic heterogeneous effects of FDA regulatory enforcement on hospital drug supply when enforcement occurs during an ongoing critical shortage, which is defined as all three primary active ingredients of sterile injectable opioids being in shortage as of the month prior to enforcement. All estimates are from difference-in-difference regression equations with the form:

$$
\begin{aligned}
Y_{h,t,\tau} = \sum_{q=-2}^{3} & \beta_q^1 \text{HospitalExposure}_{h,\tau} times \mathbf{1}\left\{q \text{ Quarters Since Enforcement}\right\}_{t,\tau} \\
& + \beta_q^2 \text{HospitalExposure}_{h,\tau} \times \mathbf{1}\left\{q \text{ Quarters Since Enforcement}\right\}_{t,\tau} \times \text{CriticalShortage}_{\tau} \\
& + \delta_{h,\tau} + \eta_{r,\tau} + X_{h,t,-\tau} + \varepsilon_{h,t,\tau}
\end{aligned}
$$

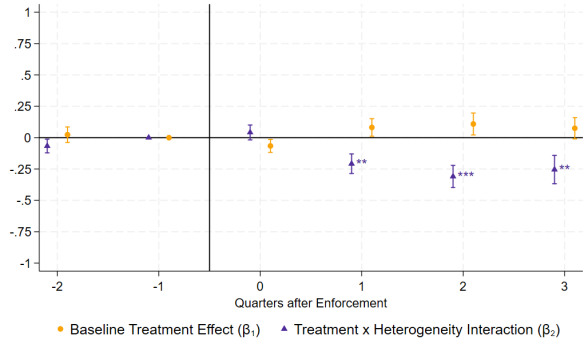
where β_q^1 represents the baseline treatment effect q quarters after FDA enforcement (plotted in orange with circular markers) and β_q^2 represents the heterogeneity interaction capturing the deviation in treatment effects when FDA enforcement occurs during an ongoing critical shortage (plotted in purple with triangular markers). Quarters are indexed such that $q = 0$ includes the month of FDA enforcement and the two following months. Panels (a), (c), and (d) plots β_q^1 and β_q^2 when the outcome variable is total hospital drug supply, number of wholesale distributor suppliers, and active ingredient drug substitution index, respectively. Symbols plotted in the figure denote the statistical significance of the test $\beta_q^1 + \beta_q^2 = 0$, indicating whether the baseline treatment effect and heterogeneity effect are collectively different from zero: + 10% * 5% ** 1% ***. Panel (b) plots β_q^2 for decomposed measures of hospital drug volume, including from the enforced manufacturing facility and other suppliers, capturing deviations from baseline treatment effects during critical shortage enforcement events.



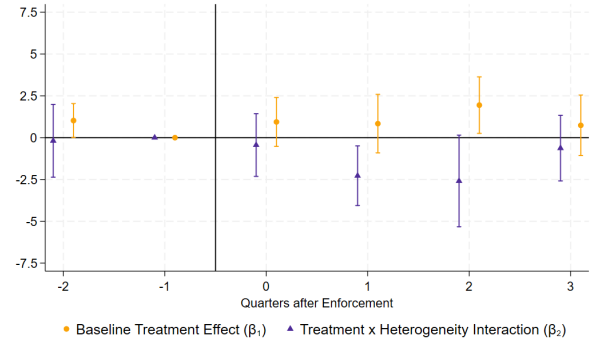
(a) Total Hospital Drug Supply



(b) Decomposed Hospital Drug Supply (Heterogeneity Interaction Only)



(c) Supplying Distributors



(d) Active Ingredient Substitution Index

Figure A.5. Heterogeneity in Hospital Drug Supply among Rural Hospitals

Notes: Figure A.5 studies the dynamic heterogeneous effects of FDA regulatory enforcement on hospital drug supply across rural and urban hospitals. All estimates are from difference-in-difference regression equations with the form:

$$
\begin{aligned}
Y_{h,t,\tau} = \sum_{q=-2}^{3} & \beta_q^1 \text{HospitalExposure}_{h,\tau} \times \mathbf{1}\left\{q \text{ Quarters Since Enforcement}\right\}_{t,\tau} \\
& + \beta_q^2 \text{HospitalExposure}_{h,\tau} \times \mathbf{1}\left\{q \text{ Quarters Since Enforcement}\right\}_{t,\tau} \times \text{Rural}_h \\
& + \delta_{h,\tau} + \eta_{r,\tau} + X_{h,t,-\tau} + \varepsilon_{h,t,\tau}
\end{aligned}
$$

where β_q^1 represents the baseline treatment effect q quarters after FDA enforcement (plotted in orange with circular markers) and β_q^2 represents the heterogeneity interaction capturing the deviation in treatment effects among rural hospitals during the same time period (plotted in purple with triangular markers). Quarters are indexed such that $q = 0$ includes the month of FDA enforcement and the two following months. Panels (a), (c), and (d) plots β_q^1 and β_q^2 when the outcome variable is total hospital drug supply, number of wholesale distributor suppliers, and active ingredient drug substitution index, respectively. Symbols plotted in the figure denote the statistical significance of the test $\beta_q^1 + \beta_q^2 = 0$, indicating whether the baseline treatment effect and heterogeneity effect are collectively different from zero: + 10% * 5% ** 1% ***. Panel (b) plots β_q^2 for decomposed measures of hospital drug volume, including from the enforced manufacturing facility and other suppliers, capturing deviations from baseline treatment effects for rural hospitals.

Table A.1. Patient Health and Quality of Care (Top Quartile)

	(1)	(2)	(3)	(4)	(5)
	90-Day Mortality	Unhealthy Days	30-Day Readmission	Log(Hospital Charges)	Log(Reimbursement)
$\mathbb{E}$ [HospitalExposure] × 0–6 Mo. Post	0.0017	0.1517	0.0035	−0.0011	0.0113
	(0.0030)	(0.1094)	(0.0047)	(0.0126)	(0.0128)
$\mathbb{E}$ [HospitalExposure] × 7–12 Mo. Post	0.0023	0.0128	−0.0072	−0.0076	0.0060
	(0.0030)	(0.1135)	(0.0048)	(0.0129)	(0.0131)
N	6,841,406	6,841,406	6,841,406	6,841,406	6,768,431
R^2	0.2199	0.4178	0.0965	0.6182	0.6424
Hospital-by-Event FEs	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓
0–6 Month Coef. as % of Baseline Mean	2.6400	1.4200	2.1800	−0.0100	0.1300
7–12 Month Coef. as % of Baseline Mean	3.5000	0.1200	−4.4900	−0.0800	0.0700
0–6 Month Coef. 95% CI as % of Baseline Mean	[−6.3, 11.58]	[-.59, 3.43]	[−3.55, 7.91]	[-.26, .23]	[-.16, .42]
7–12 Month Coef. 95% CI as % of Baseline Mean	[−5.68, 12.69]	[−1.96, 2.2]	[−10.34, 1.37]	[-.33, .18]	[-.23, .37]
Dep. Var. Pre-Mean	0.0700	10.7000	0.1600	10.0700	8.6600
Ind. Var. Mean	0.1300	0.1300	0.1300	0.1300	0.1300
Ind. Var. SD	0.1600	0.1600	0.1600	0.1600	0.1600
Ind. Var. Range	[0, .91]	[0, .91]	[0, .91]	[0, .91]	[0, .91]

Notes: * 5%, ** 1%. Table A.1 reports the results of stacked difference-in-differences estimates comparing the outcomes of index hospital visits in the top quartile of predicted SI opioid utilization, across patients with local hospital environments differentially reliant on the implicated facility for baseline drug supply (see equation 3). The treatment variable is a convex combination of hospital-level exposure to each FDA enforcement event and the pre-period hospital visit distribution, calculated at patient ZIP-code-level. Parentheses denote standard errors clustered at the patient ZIP code level.

Table A.2. Effects of FDA OAI Inspections on Hospital Visited (Top Quartile)

	(1)	(2)	(3)	(4)	(5)	(6)
	Nearest Hosp	Log(Distance)	Teaching Cntr	Rural Hosp	Hosp Exposure	Pre-Period Hospital Visit Share (All)
$\mathbb{E}$ [HospitalExposure] × 0–6 Mo. Post	0.0152***	−0.0217	−0.0122**	0.0031	−0.0034*	0.0056
	(0.0054)	(0.0163)	(0.0051)	(0.0030)	(0.0020)	(0.0034)
$\mathbb{E}$ [HospitalExposure] × 7–12 Mo. Post	0.0071	0.0189	−0.0011	0.0020	−0.0070***	0.0023
	(0.0056)	(0.0168)	(0.0057)	(0.0030)	(0.0020)	(0.0036)
N	6,836,524	6,836,524	6,840,155	6,840,155	6,841,406	6,841,406
R^2	0.3486	0.3938	0.4674	0.6246	0.7847	0.3987
Hospital-by-Event FEs	✓	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓	✓
0–6 Mo Coef % of Baseline Mean	3.8600	−0.9500	−2.3600	2.4300	−2.6400	1.5900
7–12 Mo Coef % of Baseline Mean	1.8100	0.8300	−0.2200	1.6100	−5.4500	0.6600
0–6 Mo 95% CI % of Baseline Mean	[1.17, 6.56]	[−2.34, .44]	[−4.29, -.44]	[−2.13, 6.98]	[−5.72, .44]	[-.31, 3.48]
7–12 Mo 95% CI as % of Baseline Mean	[-.99, 4.62]	[-.62, 2.27]	[−2.38, 1.94]	[−3.09, 6.31]	[−8.49, −2.4]	[−1.31, 2.63]
Dep. Var. Pre-Mean	0.3900	2.3100	0.5100	0.1300	0.1000	0.3500
Ind. Var. Mean	0.1300	0.1300	0.1300	0.1300	0.1300	0.1300
Ind. Var. SD	0.1600	0.1600	0.1600	0.1600	0.1600	0.1600
Ind. Var. Range	[0, .91]	[0, .91]	[0, .91]	[0, .91]	[0, .91]	[0, .91]

Notes: * 5%, ** 1%. Table A.2 reports the results of stacked difference-in-differences estimates comparing the outcomes of index hospital visits in the top quartile of predicted SI opioid utilization, across patients with local hospital environments differentially reliant on the implicated facility for baseline drug supply (see equation 3). The treatment variable is a convex combination of hospital-level exposure to each FDA enforcement event and the pre-period hospital visit distribution, calculated at patient ZIP-code-level. Parentheses denote standard errors clustered at the patient ZIP code level.

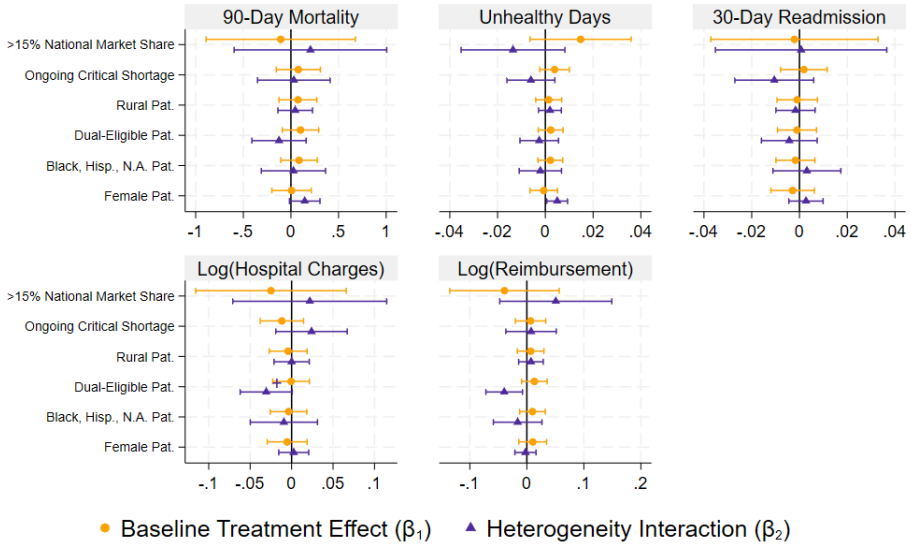


Figure A.6. Heterogeneity in Resilience of Patient Health, Quality of Care, and Healthcare Costs

Notes: Figure A.6 examines heterogeneity in the impacts of FDA regulatory enforcement on patient health and quality outcomes drug supply based on the degree of stress imposed by the enforcement action on the supply chain and distributional concerns. All estimates are from difference-in-difference regression equations with the form:

$$Y_{i,t,\tau} = \beta_1 \text{Post}_{t,\tau} \times \mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau} + \beta_2 \text{Post}_{t,\tau} \times \mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau} \times \text{Heterogeneity}_{h,\tau} + \eta_{r,\tau} + \delta_{z(i),\tau} + X_{z(i),t,-\tau} + W_i + \varepsilon_{i,t,\tau}$$

where β_1 represents the baseline post-period treatment effect of FDA enforcement (plotted in orange with circular markers) and β_2 represents the heterogeneity interaction capturing the deviation in treatment effects by each heterogeneity characteristic (plotted in purple with triangular markers). Symbols plotted in the figure denote the statistical significance of the test $\beta_1 + \beta_2 = 0$, indicating whether the baseline treatment effect and heterogeneity effect are collectively different from zero: + 10% * 5% ** 1% *** 0.1%.

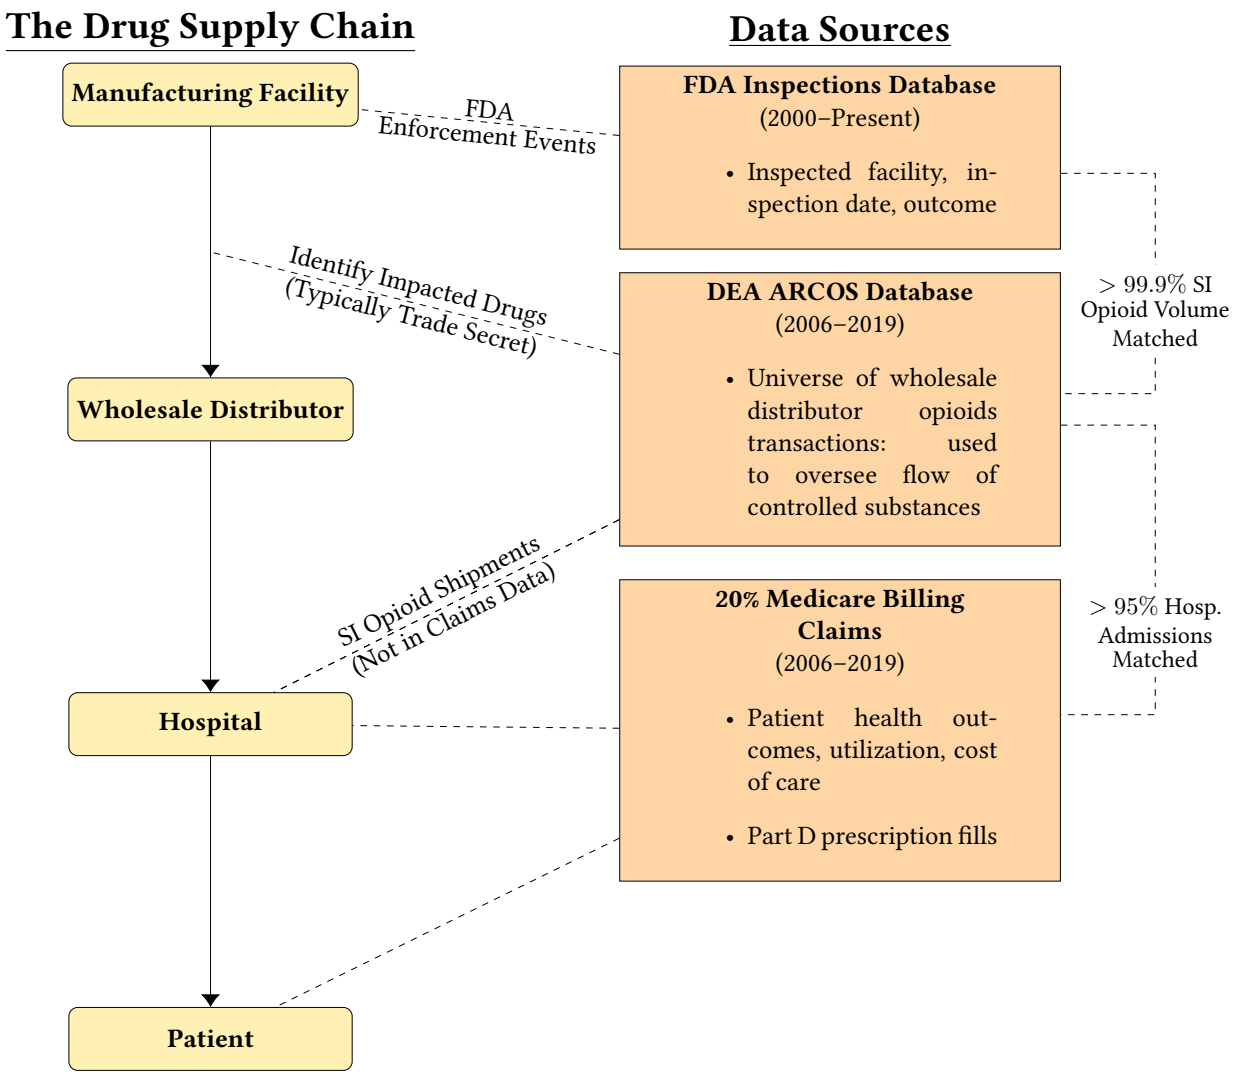
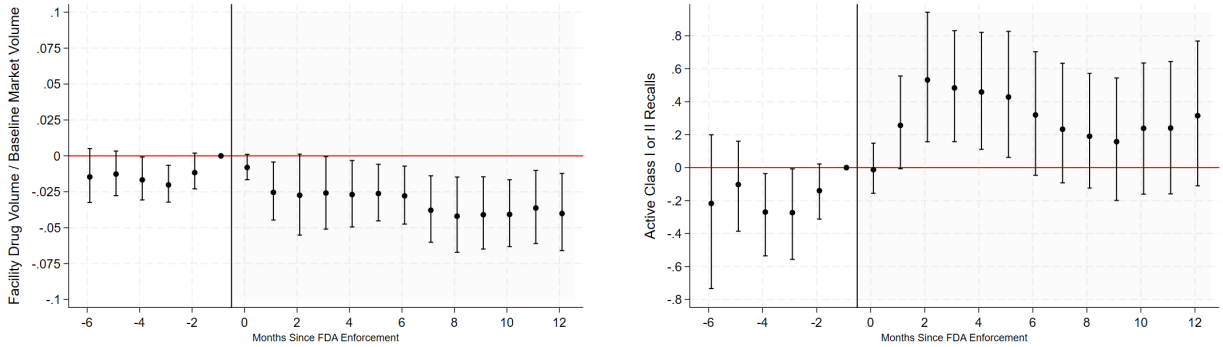


Figure A.7. Visualization of Data Sources and Drug Supply Chain

B Appendix to Upstream Manufacturing Facility Disruption



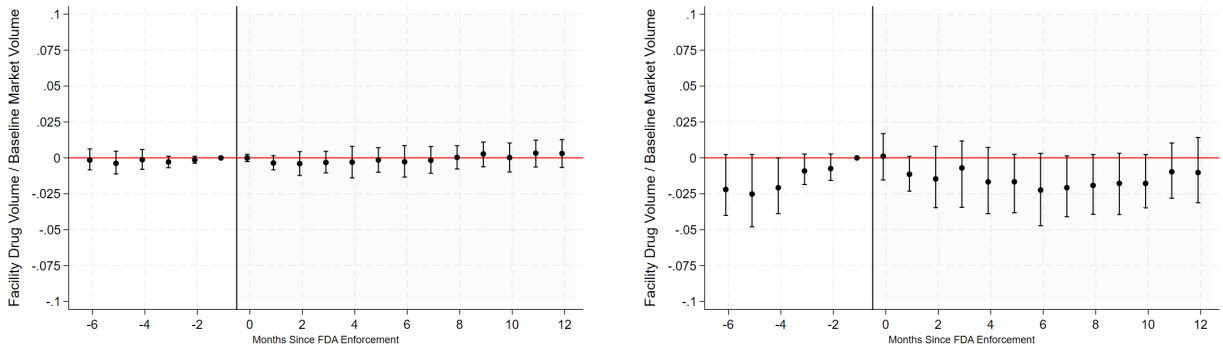
(a) Facility Drug Volume as Share of Market Baseline **(b)** Active Class 1 and 2 Recalls

Figure B.1. Impact of FDA OAI Inspections on Enforced Facility: SI Opioid Treatment Events Only

Notes: Figure B.1 plots the differential evolution of manufacturing facility-by-opioid market outcomes between enforced sterile injectable opioid manufacturing facilities and matched controls that were enforced 1–3 years later and serve strictly different opioid drug markets. These capture 15 of the 16 FDA enforcement events studied in the analysis of downstream propagation to hospitals and patients in Sections 5 and 6, respectively, where one event is dropped due to a lack of matched controls. Error bands represent 90% confidence intervals based on standard errors calculated using the wild cluster bootstrap with manufacturing facility-level clusters. The outcome variable in panel (a) is national drug volume originating from the implicated facility measured in milligrams of morphine equivalents divided by mean total market volume in the baseline period 13–18 months prior to enforcement. The results can therefore be interpreted as percentage point changes in baseline market share. Panel (b) plots the evolution of active class I and II recalls of drugs originating from each manufacturing facility. Class I and II recalls refer to recalls where use of the drug has the potential to cause adverse health consequences, and recalls are counted as active between their initiation and termination months. Since sterile injectable opioid drug prices are not reliably captured in Medicare Part D data, this outcome cannot be studied among this subset of enforcement events.

B.1 Placebo Test: Upstream Effects of VAI Inspections

To reinforce the causality of the decline in drug volume following OAI inspections, I conduct a placebo test using FDA inspections that result in VAI classifications, which are the most common FDA inspection outcome and carry comparatively less scope to cause disruptions than OAIs. Although VAI inspections do uncover some objectionable conditions, they signal an acceptable level of regulatory compliance, do not result in further enforcement actions, and leave facilities within the routine surveillance inspection schedule rather than the higher frequency "for-cause" schedule associated with OAI facilities.



(a) All Opioid Drug Manufacturing Facilities **(b)** SI Opioid Manufacturing Facilities Only

Figure B.2. Impact of FDA VAI Inspections on Directly Impacted Facility Drug Volume

Notes: Figure B.2 plots the results of a placebo test that replicates the empirical strategy in section 4 for voluntary action indicated (VAI) inspection classifications rather than OAI classifications corresponding to FDA enforcement. Specifically, estimates compare the evolution of drug volume between manufacturing facilities receiving voluntary action indicated (VAI) inspection classifications and matched controls that receive VAIs 1–3 years later and serve strictly different opioid drug markets. Error bands represent 90% confidence intervals based on standard errors calculated using the wild cluster bootstrap with manufacturing facility-level clusters. Panel (a) plots the results from a specification using the set of VAI inspections across the set of all manufacturing facilities represented in the DEA ARCOS database. Panel (b) restricts to treated units that produce sterile injectable opioid drugs.

Figure B.2 presents the dynamic difference-in-differences estimates comparing facilities treated by VAI inspections against matched controls that will receive a VAI 1–3 years in the future. Table B.1, reporting the difference-in-difference estimates, shows that VAI inspections have a statistically insignificant and near zero impact on drug volume both in the full sample of all opioid manufacturing facilities and the subset that produce SI opioids.

Table B.1. Impact of FDA VAI Inspections on Implicated Facility Drug Volume

	(1)	(2)	(3)	(4)	(5)	(6)
	All Opioid Manufacturing Facilities				**SI Facilities Only**	
	Market Share	Active Recalls	Log(Enforced $/MME)	Log(Rest of Market $/MME)	Market Share	Active Recalls
Enforced× 0–12 Months Post	0.00	−0.16	−0.02	0.02	−0.00	−0.09
N	11,134	11,134	1,900	1,900	3,971	3,971
R^2	0.97	0.87	0.97	0.99	0.97	0.80
Dep. Var. Pre-Mean	0.05	0.23	−2.55	−3.05	0.11	0.47
Facility-by-Event FEs	✓	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓	✓
Overlapping Event Controls	✓	✓	✓	✓	✓	✓
Wild Cluster Bootstrap 95% CI	[-.005989, .008499]	[-.439, .0676]	[-.202, .186]	[-.017, .045]	[-.01188, .01049]	[-.3341, .153]
Wild Cluster Bootstrap p-value	0.79	0.30	0.90	0.41	0.99	0.57
# of Events	124	124	30	30	25	25
# of Facilities	112	112	41	41	59	59

Notes: * 10%, ** 5%, *** 1%. Table B.1 reports the results of placebo difference-in-differences estimates of the effects of FDA voluntary action indicated (VAI) inspections, rather than official action indicated (OAI) inspections that mark regulatory enforcement action. Estimates compare VAI facilities with matched controls receiving VAI classifications 1–3 years later and serving non-overlapping opioid drug markets (see equation 1). Outcomes include drug volume (divided by baseline total market volume), active class I and II drug recalls, and the logarithms of facility drug prices and rest of market drug prices. Columns 5 and 6 report the results when restricting to FDA enforcement events impacting SI opioid manufacturing facilities, which are further studied in the analysis of downstream propagation. All inference is conducted using the wild cluster bootstrap with manufacturing facility-level clusters.

B.2 Dose-Response Impacts of OAI Inspections on Aggregate Outcomes

To study the dose-response relationship between the size of the upstream manufacturing facility implicated by FDA regulatory enforcement and aggregate opioid market volume, I run a modified version of the matched controls analysis presented in section 4 where the binary FDA enforcement variable is replaced with a continuous measure of the national market share held by the enforced facility 13–18 months before enforcement.

$$Y_{f,m,t,\tau} = \sum_{r=-6, r\neq -1}^{12} \beta_r \times \text{EnforcedMarketShare}_{f,m,\tau} \times \mathbf{1}\left\{r \text{ Months Since Enforcement}\right\}_{t,\tau} + \eta_{r,\tau} + \delta_{f,m,\tau} + X_{f,m,t,-\tau} + \varepsilon_{f,m,t,\tau} \quad (4)$$

Though the results of this analysis are best interpreted as suggestive due to the statistical imprecision of estimates based on a single national unit of observation for each opioid drug market, it is intended to complement the more precise hospital cross-sectional analysis presented in section 5 in two ways. First, it tests whether the premise of the cross-sectional analysis—that the potential disruption created by FDA enforcement is proportional to the share of baseline drug supply accounted for by the enforced facility—is borne out in aggregate data. Second, the analysis allows me to directly estimate competing facility ramp-up and the general equilibrium "missing intercept" absorbed by fixed effects in cross-sectional analysis.

To closely link to the analysis of downstream hospital drug supply, I restrict to enforcement events at SI opioid manufacturing facilities only. For the competing facility and total market drug volume outcomes, standard errors are clustered at the market-by-event level while those for facility-specific drug volume are clustered at the manufacturing facility level. All inference is conducted using the wild cluster bootstrap.

Estimated effects on drug volume from the enforced facility, all other manufacturing facilities producing for the impacted market, and total market drug volume are presented in Table B.2, with the corresponding dynamic estimates plotted in Figure B.3. The results bear out the prediction that FDA enforcement at larger drug manufacturing facilities result in larger volume declines from the enforced facility. For every percentage point of national market share held by the enforced facility at baseline, its drug volume falls by the equivalent of 0.2pp of market share in the 12 months after enforcement. Importantly, the results also hint at supply-side responses to FDA enforcement in the form of increased drug volume from non-enforced competing facilities, which increase by 0.35pp of market share for every percentage point of drug baseline market share held by the implicated facility.

This finding provides important context to the rapid substitution to alternative manufacturing facilities documented in section 5, indicating ramp-up by other manufacturing facilities rather

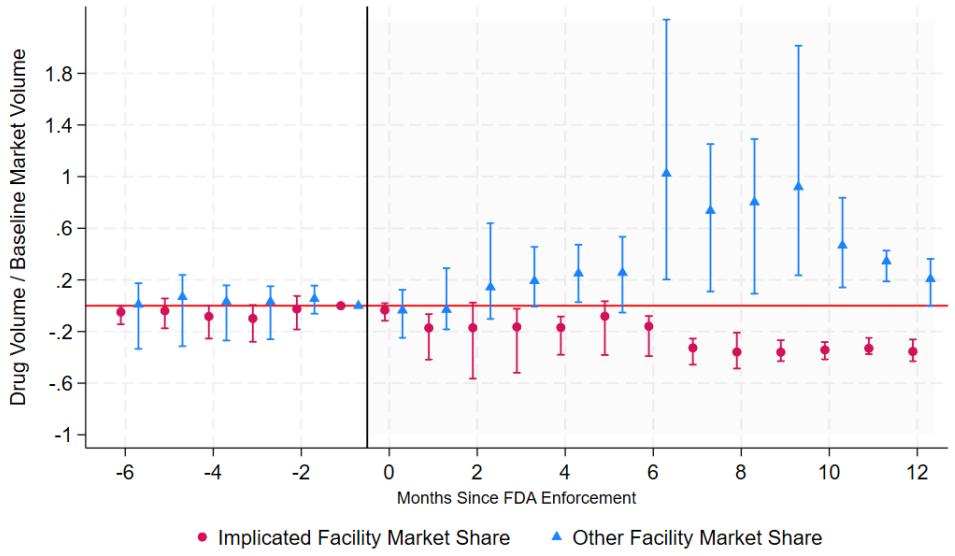


Figure B.3. Impact of FDA OAI Inspections on Implicated and Other Facility Drug Volume

Notes: Figure B.3 plots the evolution of aggregate national drug volume by the enforced manufacturing facility and by all other manufacturing facilities serving the same opioid drug market around FDA OAI inspections. Each panel plots $\hat{\beta}_r$ dynamic difference-in-differences estimates from equation 4, which compare the trajectory of aggregate drug volume in opioid markets served by facilities implicated by FDA OAI inspections against aggregate drug volume in markets served by matched control facilities that will be enforced 1–3 years later. The treatment variable for these regressions is the national market share held by the OAI facilities at baseline, thereby capturing the size of the potential disruption to drug volume created by FDA enforcement. By design, treatment and control facilities matched together serve distinct opioid drug market. Drug volume originating from the implicated facility and from all competing manufacturing facilities are divided by total opioid drug market volume at baseline ($r \in [-18, -13]$) so that the results have a percentage point change in market share interpretation. Error bands represent 90% confidence intervals based on standard errors calculated using the wild cluster bootstrap with market-by-event clusters.

Table B.2. Impact of FDA OAI Inspections on Aggregate Volume

	(1)	(2)	(3)
	Implicated Facility	Other Facilities	Total Market Volume
Treat × Baseline Market Share× 0–12 Mo. Post	−0.19***	0.35**	0.15
N	1,330	1,330	1,330
R^2	0.98	0.93	0.94
Facility-by-Event FEs	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓
Overlapping Event Controls	✓	✓	✓
Wild Cluster Bootstrap 95% CI	[-.2768, -.1146]	[.05717, .7686]	[-.07689, .4274]
Wild Cluster Bootstrap p-value	0	0.01	0.38
Wild Cluster Bootstrap t-stat	−6.13	3.05	1.70
Dep. Var. Pre-Mean	0.11	0.87	0.97
# of Events	15	15	15
# of Facilities	31	31	31

Notes: * 10%, ** 5%, *** 1%. Table B.2 studies the role of competitor ramp-up in response to OAI inspections by reporting difference-in-differences estimates of the dose-response relationship between the size of the enforced manufacturing facility, as measured by its baseline national market share, and aggregate drug volume outcomes. To complement the cross-sectional analysis of downstream hospital drug supply, the regression sample is limited to enforcement events impacting manufacturing facilities that produce SI opioid drugs. The estimates compare enforced facilities with matched controls enforced 1–3 years later and serving non-overlapping opioid drug markets (see equation 4). Outcomes include national facility drug volume, drug volume by all other facilities serving the same opioid drug market, and total opioid drug market volume (each divided by baseline total market volume). Standard errors are clustered at the facility-level for implicated facility volume and at the market-by-event level for other facilities and total market volume. All inference is conducted using the wild cluster bootstrap with manufacturing facility-level clusters.

than pure redistribution across more- and less-exposed hospitals in the supply chain. Added together into total market drug volume, the estimates point to noisy but statistically insignificant effects on total drug volume.

B.3 Future-Treated Matched Controls Algorithm

My analysis considers the set of all FDA OAI inspections with drug volume recorded in the DEA ARCOS database starting at least 18 months prior to the month of inspection, two or more months of positive drug volume in the $r \in [-6, -1]$ pre-period, and in which the enforced facility held at least 0.01% market share in the $r \in [-18, -13]$ baseline period. These restrictions, which are also applied to matched controls, are intended to ensure that facility drug volume can be reliably measured before and after each inspection.[56]

Additionally, OAI inspections are only considered if the facility did not receive another OAI in the $r \in [-12, -1]$ pre-period to prevent contamination from the ongoing effects of past enforcement. The treatment group therefore consists of "initial" OAI inspection events, where the effects of subsequent OAI inspections within 12 months appear in the $r \in [0, 12]$ post-period and

[56] Although this criteria excludes manufacturing facilities with drug volume initially beginning in $[-18, -1]$, the results should be interpreted as applying to established rather than to new drug volume.

are treated as follow-ups to the initial enforcement action. Indeed, once facilities receive an OAI, they tend to be inspected at a higher frequency within the FDA's "for-cause" inspection regime, which can result in repeated OAI classifications in quick succession.[57]

The final restriction is to drop inspections of drug compounding facilities that are not registered as 503B outsourcers at the time of inspection, since these facilities fall under the primary regulatory jurisdiction of state pharmacy boards rather than the FDA. Since there are documented instances of such facilities resisting FDA inspections and subsequent enforcement action, OAIs at these facilities are less likely to cause disruptions.

Among all OAI inspections meeting these inclusion criteria, a separate OAI event is created to capture the trajectory of drug volume for each distinct opioid drug category served by the implicated facility. The 80 opioid manufacturing facilities receiving an FDA OAI inspection served an average of 1.8 distinct opioid drug categories. To reduce redundancy and contamination risk, OAI events impacting the same opioid market and taking place in the same month are combined into a single composite event considering their combined impact. For each of these events, a sub-experimental balanced panel dataset spanning event months $[-6, 12]$ around the OAI event is constructed to reflect drug volume originating from the treated facilities and possibly multiple future-treated matched control facilities. Other than having no OAI in event months $r \in [-12, -1]$, no further restrictions are placed on the treatment group's inspection outcomes: NAI and VAI classifications are permitted both in the pre- and post-periods and follow-up OAIs are permitted in $r \in [1, 12]$.

Matched controls cannot receive any OAIs in event months $r \in [-12, 12]$, must meet the same criteria for minimum drug volume in the pre-period, and must be subject to an OAI inspection in event months $r \in [13, 36]$. Facility-by-opioid market volume trajectories from matched controls are only included if they are from the subset of opioid categories that are *not* served by the treatment group OAI facility. This ensures that matched controls are not subject to spillover effects that would bias estimates, such as from competitive ramp-up responses to a production disruption at the treated facility.[58] If one or more potential matched control facility-by-opioid category trajectories meets all of the requirements listed above, their balanced time series of drug volume spanning $r \in [-6, 12]$ is appended to the sub-experimental dataset for the OAI inspection. If no matches are found, the entire OAI inspection sub-experiment is dropped from the analysis. Appending all sub-experimental datasets together yields the final regression sample.

[57] Falsification tests using VAI inspections are based on analogous but more restrictive criteria to prevent contamination from OAIs. To be considered an "initial" VAI inspection, the manufacturing facility must not have received any VAI or OAI in the 12 months leading up to inspection. Additionally, the facility should not receive an OAI inspection in the 12-month post-inspection period to ensure that outcomes are not contaminated.

[58] Since all markets considered in my analysis are for opioid drugs, using different markets captures secular trends related to opioids and the prescription opioid crisis but without scope for immediate substitution.

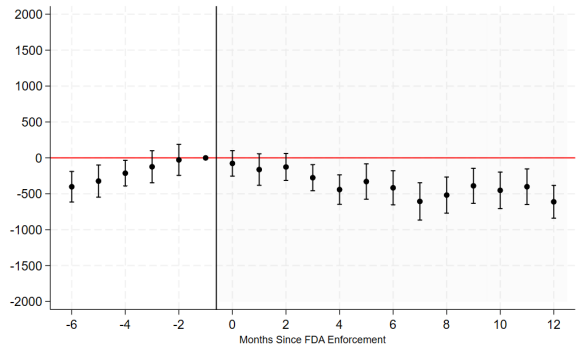
C Appendix to Downstream Hospital Drug Supply Analysis

C.1 Addressing Hospital Drug Volume Outliers in DEA ARCOS Data

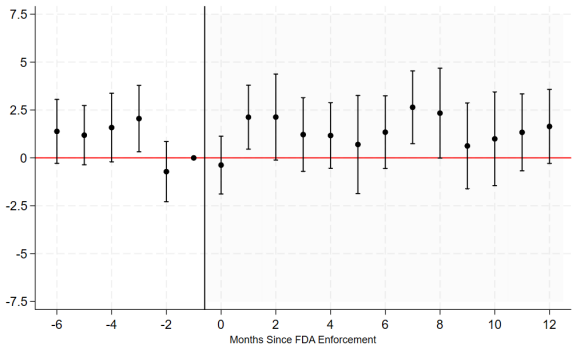
To minimize the influence of spurious observations, I implement an outlier detection procedure that conservatively flags implausibly high drug shipment volumes. This approach is motivated by instances in which a hospital's drug volume exhibits an isolated spike far exceeding surrounding values, which may be due to data entry errors (e.g., recording quantities in micrograms instead of milligrams). For each hospital (168 monthly observations over 14 years for those observed throughout the study period), I calculate the 95th percentile of monthly drug volume in non-zero months. This serves as an empirical upper bound characterizing typical behavior, excluding the influence of potentially extreme values. An observation is then classified as an outlier if it exceeds this 95th percentile threshold by 20 times or more.

This threshold reflects a conservative approach. While hospitals may experience legitimate one-off increases in volume—for example, replenishing inventory or stockpiling after a recall, or experiencing a temporary surge in patient demand—such events are unlikely to entail orders 20 times larger than the 95th percentile, which already represents a high level of ordering.

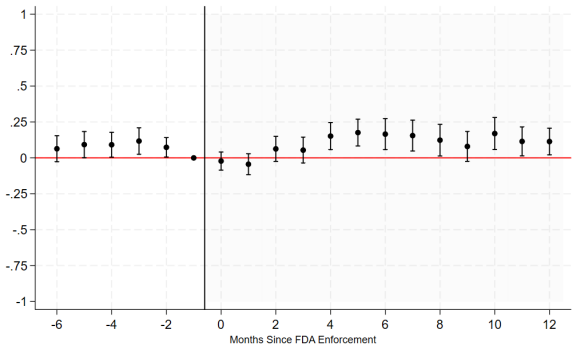
For each FDA enforcement stack in my stacked regression sample, a hospital's entire time series of outcomes is dropped from the regression sample if one or more months is flagged as an outlier. Overall, 68, or 0.01%, hospital-months are flagged as outliers. These represent only 0.25% of hospital-months with volume in excess of the 95th percentile, reinforcing the rarity of outlier observations. Of these, 57% occur between November 2011 and February 2012 and 19% between April 2014 and July 2014. In both of these periods, shipment volume from a single wholesale distributor appears consistent with misreporting of drug units.



(a) HHI Across Manufacturing Facilities



(b) Active Ingredient Substitution Index

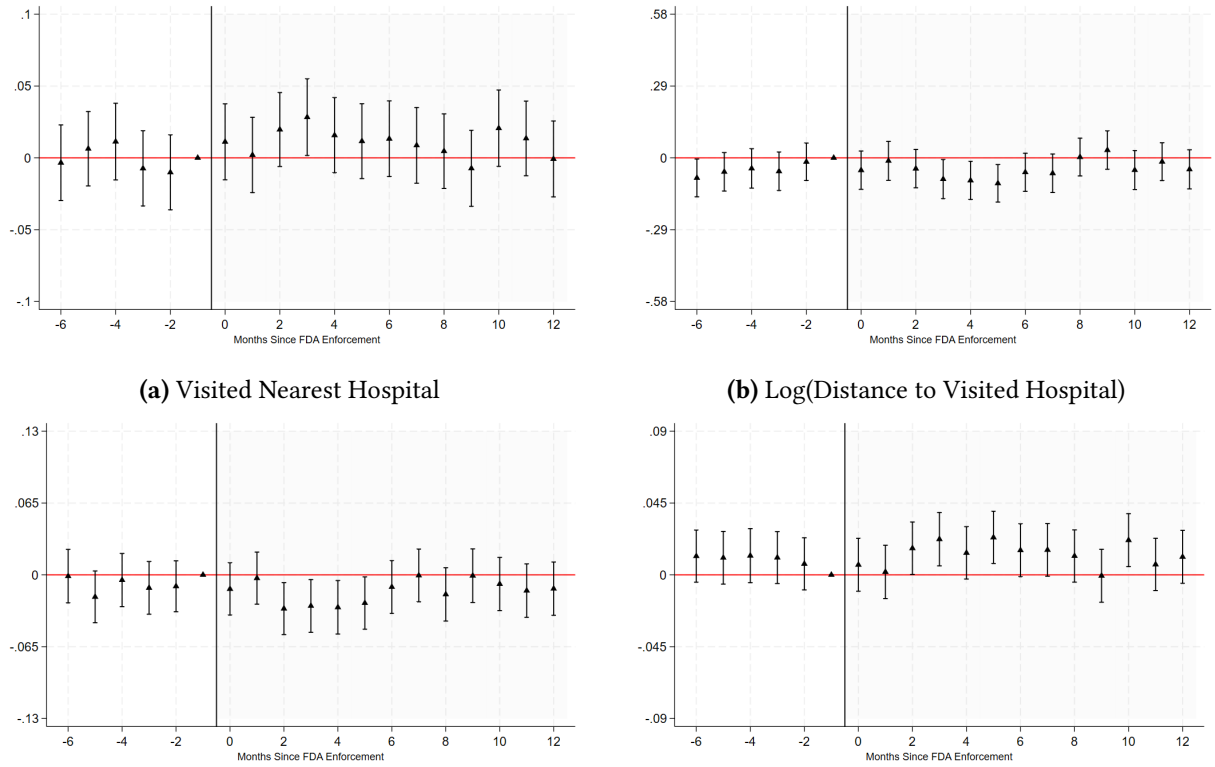


(c) # of Supplying Distributors

Figure C.1. Hospital Adaptation Responses to FDA Enforcement

Notes: Figure C.1 plots the evolution of hospital drug supply characteristics around FDA OAI inspections. The dynamic difference-in-differences estimates corresponding to $\hat{\beta}_r$ from equation 2 compares the trajectories of hospitals differentially reliant on the enforced upstream manufacturing facility at baseline for drug supply. Panels (a), (b), and (c) plot the results when using hospital Herfindahl-Hirschman index (HHI) across supplying manufacturing facilities, active ingredient substitution index, and number of supplying distributors as outcome variables, respectively. HHI is calculated by summing the squared shares of monthly hospital drug supply accounted for by each supplying manufacturing facility, measuring concentration in hospital supply networks. The active ingredient substitution index is constructed by calculating the absolute value of the change in share of hospital drug supply accounted for by each opioid active ingredient within the class of sterile injectable opioid drugs between the current month and the same month in the previous year, and then summing these changes together across active ingredients. Stability in the active ingredient composition will lead to an index value close to zero, with increases corresponding to changes in drug composition. The number of supplying distributors includes both wholesale distributors and the direct-from-manufacturer purchases. Vertical axes are set to cover the range spanning $\pm 50\%$ of the outcome variable's pre-period baseline mean. Error bands represent 95% confidence intervals with standard errors clustered at the hospital level.

D Appendix to Downstream Patient Outcomes Analysis



(a) Visited Nearest Hospital

(b) Log(Distance to Visited Hospital)

(c) Visited Teaching Hospital

(d) Predicted Visit Probability of Visited Hospital

Figure D.1. Impact of FDA OAI Inspections on Hospital Care is Received From

Notes: Figure D.1 plots the evolution of the characteristics of hospitals that patients receive care from around FDA OAI inspections. The dynamic difference-in-differences estimates corresponding to $\hat{\beta}_r$ from equation 3 compares the trajectories of hospital visits by patients residing in ZIP codes with local hospital environments with baseline drug supply differentially reliant on the enforced manufacturing facility. The treatment variable used to make this comparison, $\mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau}$, is a convex combination of hospital-level exposure to FDA enforcement event τ and patient ZIP code-level hospital visit probabilities, which are based on the fixed empirical distribution of hospital visits in the pre-study period. Error bands represent 95% confidence intervals with standard errors clustered at the patient ZIP code level.

Table D.1. Patient Health and Quality of Care (Top Quartile, Quartile-Specific Hospital Shares)

	(1)	(2)	(3)	(4)	(5)
	90-Day Mortality	Unhealthy Days	30-Day Readmission	Log(Hospital Charges)	Log(Reimbursement)
$\mathbb{E}$ [QtileHospitalExposure] × 0–6 Mo. Post	0.0020	0.1870	0.0055	0.0058	0.0164
	(0.0032)	(0.1174)	(0.0050)	(0.0134)	(0.0137)
$\mathbb{E}$ [QtileHospitalExposure] × 7–12 Mo. Post	0.0024	0.0040	−0.0072	−0.0039	0.0096
	(0.0033)	(0.1213)	(0.0051)	(0.0138)	(0.0141)
N	6,836,807	6,836,807	6,836,807	6,836,807	6,763,894
R^2	0.2199	0.4178	0.0964	0.6182	0.6424
Hospital-by-Event FEs	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓
0–6 Month Coef. as % of Baseline Mean	3.0500	1.7500	3.4400	0.0600	0.1900
7–12 Month Coef. as % of Baseline Mean	3.6400	0.0400	−4.4900	−0.0400	0.1100
0–6 Month Coef. 95% CI as % of Baseline Mean	[−6.5, 12.6]	[-.4, 3.9]	[−2.66, 9.53]	[-.2, .32]	[-.12, .5]
7–12 Month Coef. 95% CI as % of Baseline Mean	[−6.18, 13.46]	[−2.19, 2.26]	[−10.71, 1.74]	[-.31, .23]	[-.21, .43]
Dep. Var. Pre-Mean	0.0700	10.7000	0.1600	10.0700	8.6600
Ind. Var. Mean	0.1200	0.1200	0.1200	0.1200	0.1200
Ind. Var. SD	0.1600	0.1600	0.1600	0.1600	0.1600
Ind. Var. Range	[0, .89]	[0, .89]	[0, .89]	[0, .89]	[0, .89]

Notes: * 5%, ** 1%. Table D.1 reports the results of stacked difference-in-differences estimates comparing the outcomes of index hospital visits in the top quartile of predicted SI opioid utilization, across patients with local hospital environments differentially reliant on the implicated facility for baseline drug supply (see equation 3). The treatment variable is a convex combination of hospital-level exposure to each FDA enforcement event and the top quartile's pre-period hospital visit distribution, calculated at patient ZIP-code-level and shrunk towards the HSA-level rates using empirical bayes. Parentheses denote standard errors clustered at the patient ZIP code level.

Table D.2. Patient Health and Quality of Care (Top Quartile, Visited Hospital Exposure)

	(1)	(2)	(3)	(4)	(5)
	90-Day Mortality	Unhealthy Days	30-Day Readmission	Log(Hospital Charges)	Log(Reimbursement)
HospitalExposure× 0–6 Mo. Post	−0.0008	0.0653	0.0036	−0.0091	−0.0032
	(0.0020)	(0.0725)	(0.0032)	(0.0079)	(0.0083)
HospitalExposure× 7–12 Mo. Post	−0.0006	−0.0220	−0.0027	−0.0171**	−0.0042
	(0.0021)	(0.0738)	(0.0033)	(0.0083)	(0.0088)
N	6,944,168	6,944,168	6,944,168	6,944,168	6,871,191
R^2	0.1775	0.3896	0.0495	0.6406	0.6387
Hospital-by-Event FEs	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓
0–6 Month Coef. as % of Baseline Mean	−1.2300	0.6100	2.2700	−0.0900	−0.0400
7–12 Month Coef. as % of Baseline Mean	−0.9800	−0.2000	−1.6600	−0.1700	−0.0500
0–6 Month Coef. 95% CI as % of Baseline Mean	[−7.10, 4.64]	[-.71, 1.92]	[−1.64, 6.19]	[-.24, .06]	[-.22, .15]
7–12 Month Coef. 95% CI as % of Baseline Mean	[−7.16, 5.2]	[−1.55, 1.14]	[−5.73, 2.41]	[-.33, -.01]	[-.25, .15]
Dep. Var. Pre-Mean	0.0700	10.6900	0.1600	10.0600	8.6600
Ind. Var. Mean	0.2400	0.2400	0.2400	0.2400	0.2400
Ind. Var. SD	0.2000	0.2000	0.2000	0.2000	0.2000
Ind. Var. Range	[0, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]

Notes: * 5%, ** 1%. Table D.2 reports the results of stacked difference-in-differences estimates comparing the outcomes of index hospital visits in the top quartile of predicted SI opioid utilization, across patients who visited hospitals differentially reliant on the implicated facility for baseline drug supply (see equation 3). In other words, the treatment variable is based on the hospital actually visited to receive care. Parentheses denote standard errors clustered at the hospital level.

Table D.3. Patient Health and Quality of Care (All Index Visits)

	(1)	(2)	(3)	(4)	(5)
	90-Day Mortality	Unhealthy Days	30-Day Readmission	Log(Hospital Charges)	Log(Reimbursement)
$\mathbb{E}$ [HospitalExposure] × 0–6 Mo. Post	0.0003	0.0222	−0.0017	−0.0058	0.0060
	(0.0014)	(0.0497)	(0.0023)	(0.0063)	(0.0064)
$\mathbb{E}$ [HospitalExposure] × 7–12 Mo. Post	0.0010	0.0052	−0.0024	−0.0135**	0.0056
	(0.0015)	(0.0510)	(0.0024)	(0.0065)	(0.0067)
N	27,019,426	27,019,426	27,019,426	27,019,426	26,838,123
R^2	0.1816	0.3965	0.0496	0.6559	0.6758
Hospital-by-Event FEs	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓
0–6 Month Coef. as % of Baseline Mean	0.5000	0.3000	−1.0500	−0.0600	0.0800
7–12 Month Coef. as % of Baseline Mean	1.6100	0.0700	−1.5400	−0.1500	0.0800
0–6 Month Coef. 95% CI as % of Baseline Mean	[−3.95, 4.94]	[−1.02, 1.62]	[−3.93, 1.84]	[-.2, .07]	[-.09, .25]
7–12 Month Coef. 95% CI as % of Baseline Mean	[−2.99, 6.22]	[−1.28, 1.42]	[−4.53, 1.45]	[-.29, -.01]	[-.1, .25]
Dep. Var. Pre-Mean	0.0600	7.3900	0.1600	9.0300	7.4700
Ind. Var. Mean	0.1300	0.1300	0.1300	0.1300	0.1300
Ind. Var. SD	0.1600	0.1600	0.1600	0.1600	0.1600
Ind. Var. Range	[0, .98]	[0, .98]	[0, .98]	[0, .98]	[0, .98]

Notes: * 5%, ** 1%. Table D.3 reports the results of stacked difference-in-differences estimates comparing the outcomes of all index hospital visits across patients with local hospital environments differentially reliant on the implicated facility for baseline drug supply (see equation 3). The treatment variable is a convex combination of hospital-level exposure to each FDA enforcement event and the pre-period hospital visit distribution, calculated at patient ZIP-code-level. Parentheses denote standard errors clustered at the patient ZIP code level.

Table D.4. Patient Health and Quality of Care (All Index Visits, Visited Hospital Exposure)

	(1)	(2)	(3)	(4)	(5)
	90-Day Mortality	Unhealthy Days	30-Day Readmission	Log(Hospital Charges)	Log(Reimbursement)
HospitalExposure× 0–6 Mo. Post	0.0002	0.0088	−0.0006	−0.0072	−0.0019
	(0.0010)	(0.0355)	(0.0016)	(0.0048)	(0.0049)
HospitalExposure× 7–12 Mo. Post	−0.0008	0.0140	−0.0008	−0.0100*	0.0032
	(0.0010)	(0.0377)	(0.0017)	(0.0053)	(0.0054)
N	27,140,717	27,140,717	27,140,717	27,140,717	26,959,020
R^2	0.1668	0.3869	0.0335	0.6824	0.6780
Hospital-by-Event FEs	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓
0–6 Month Coef. as % of Baseline Mean	0.2700	0.1200	−0.3600	−0.0800	−0.0300
7–12 Month Coef. as % of Baseline Mean	−1.2900	0.1900	−0.5000	−0.1100	0.0400
0–6 Month Coef. 95% CI as % of Baseline Mean	[−2.79, 3.32]	[-.81, 1.04]	[−2.39, 1.67]	[-.18, .02]	[-.15, .1]
7–12 Month Coef. 95% CI as % of Baseline Mean	[−4.45, 1.87]	[-.79, 1.16]	[−2.58, 1.57]	[-.22, .01]	[-.1, .18]
Dep. Var. Pre-Mean	0.0600	7.3900	0.1600	9.0300	7.4700
Ind. Var. Mean	0.2400	0.2400	0.2400	0.2400	0.2400
Ind. Var. SD	0.1900	0.1900	0.1900	0.1900	0.1900
Ind. Var. Range	[0, 1]	[0, 1]	[0, 1]	[0, 1]	[0, 1]

Notes: * 5%, ** 1%. Table D.4 reports the results of stacked difference-in-differences estimates comparing the outcomes of all index hospital visits across patients who visited hospitals differentially reliant on the implicated facility for baseline drug supply (see equation 3). In other words, the treatment variable is based on the hospital actually visited to receive care. Parentheses denote standard errors clustered at the patient ZIP code level.

Table D.5. Patient Health and Quality of Care (Bottom Quartile)

	(1)	(2)	(3)	(4)	(5)
	90-Day Mortality	Unhealthy Days	30-Day Readmission	Log(Hospital Charges)	Log(Reimbursement)
$\mathbb{E}$ [HospitalExposure] × 0–6 Mo. Post	−0.0005	−0.0025	−0.0123***	0.0081	0.0210*
	(0.0022)	(0.0766)	(0.0043)	(0.0115)	(0.0115)
$\mathbb{E}$ [HospitalExposure] × 7–12 Mo. Post	0.0011	0.0425	−0.0033	−0.0192	0.0034
	(0.0022)	(0.0801)	(0.0045)	(0.0118)	(0.0118)
N	7,032,631	7,032,631	7,032,631	7,032,631	7,002,095
R^2	0.1584	0.2705	0.0765	0.5318	0.4648
Hospital-by-Event FEs	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓
0–6 Month Coef. as % of Baseline Mean	−1.6300	−0.0600	−8.7900	0.1000	0.3300
7–12 Month Coef. as % of Baseline Mean	3.2900	1.1000	−2.3300	−0.2300	0.0500
0–6 Month Coef. 95% CI as % of Baseline Mean	[−14.83, 11.58]	[−3.96, 3.83]	[−14.85, −2.74]	[-.18, .37]	[-.02, .67]
7–12 Month Coef. 95% CI as % of Baseline Mean	[−10.14, 16.72]	[−2.97, 5.18]	[−8.69, 4.03]	[-.52, .05]	[-.3, .41]
Dep. Var. Pre-Mean	0.0300	3.8300	0.1400	8.2000	6.4700
Ind. Var. Mean	0.1300	0.1300	0.1300	0.1300	0.1300
Ind. Var. SD	0.1600	0.1600	0.1600	0.1600	0.1600
Ind. Var. Range	[0, .98]	[0, .98]	[0, .98]	[0, .98]	[0, .98]

Notes: * 5%, ** 1%. Table D.5 reports the results of placebo stacked difference-in-differences estimates comparing the outcomes of index hospital visits in the *bottom* quartile of predicted SI opioid utilization, across patients with local hospital environments differentially reliant on the implicated facility for baseline drug supply (see equation 3). The treatment variable is a convex combination of hospital-level exposure to each FDA enforcement event and the pre-period hospital visit distribution, calculated at patient ZIP-code-level. Parentheses denote standard errors clustered at the patient ZIP code level.

Table D.6. Placebo Effects of FDA OAI Inspections on Hospital Visited (Bottom Quartile)

	(1)	(2)	(3)	(4)	(5)	(6)
	Nearest Hosp	Log(Distance)	Teaching Cntr	Rural Hosp	Hosp Exposure	Pre-Period Hospital Visit Share (All)
$\mathbb{E}$ [HospitalExposure] × 0–6 Mo. Post	−0.0007	−0.0334**	0.0012	−0.0012	−0.0009	0.0030
	(0.0050)	(0.0159)	(0.0045)	(0.0023)	(0.0017)	(0.0029)
$\mathbb{E}$ [HospitalExposure] × 7–12 Mo. Post	0.0057	−0.0196	−0.0046	0.0004	−0.0007	−0.0004
	(0.0053)	(0.0162)	(0.0053)	(0.0024)	(0.0017)	(0.0031)
N	7,027,913	7,027,913	7,031,112	7,031,112	7,032,631	7,032,631
R^2	0.4084	0.3778	0.5949	0.7883	0.8326	0.4905
Hospital-by-Event FEs	✓	✓	✓	✓	✓	✓
Event-by-Event Time FEs	✓	✓	✓	✓	✓	✓
Primary Diagnosis FEs	✓	✓	✓	✓	✓	✓
0–6 Mo Coef % of Baseline Mean	−0.1300	−1.7000	0.2700	−0.6900	−0.7300	0.6900
7–12 Mo Coef % of Baseline Mean	1.1100	−1	−1.0300	0.2300	−0.5500	−0.0900
0–6 Mo 95% CI % of Baseline Mean	[−2.04, 1.79]	[−3.3, -.11]	[−1.7, 2.24]	[−3.32, 1.93]	[−3.43, 1.97]	[-.66, 2.03]
7–12 Mo 95% CI as % of Baseline Mean	[-.9, 3.12]	[−2.62, .62]	[−3.34, 1.28]	[−2.49, 2.95]	[−3.28, 2.18]	[−1.49, 1.31]
Dep. Var. Pre-Mean	0.5200	1.9800	0.4400	0.1800	0.1000	0.4300
Ind. Var. Mean	0.1300	0.1300	0.1300	0.1300	0.1300	0.1300
Ind. Var. SD	0.1600	0.1600	0.1600	0.1600	0.1600	0.1600
Ind. Var. Range	[0, .98]	[0, .98]	[0, .98]	[0, .98]	[0, .98]	[0, .98]

Notes: * 5%, ** 1%. Table D.6 reports the results of placebo stacked difference-in-differences estimates comparing the outcomes of index hospital visits in the *bottom* quartile of predicted SI opioid utilization, across patients with local hospital environments differentially reliant on the implicated facility for baseline drug supply (see equation 3). The treatment variable is a convex combination of hospital-level exposure to each FDA enforcement event and the pre-period hospital visit distribution, calculated at patient ZIP-code-level. Parentheses denote standard errors clustered at the patient ZIP code level.

D.1 Sterile Injectable Opioid Utilization Prediction

To increase the sensitivity of my analysis and mitigate the dilution of estimated treatment effects from index hospital visits unlikely to require the disrupted drugs, I restrict attention to a subsample of patients with the highest ex-ante likelihood of treatment with SI opioid medications within standard care.[59] This high-risk sample is defined using predictions from a cross-validated linear LASSO model trained on pre-period data and using anesthesia or operating room services as a proxy for likely SI opioid utilization. This proxy is informed by clinical input from an emergency physician, who emphasized that SI opioids are routinely administered for surgical and post-operative pain management. A proxy is necessary because SI drugs are reimbursed in inpatient settings through bundled payments, and so cannot reliably be directly observed.

The LASSO model is trained using patient characteristics that are immutable at the time of hospital arrival, including age (and age squared), sex, race, and primary diagnosis. To decrease dimensionality in generating the prediction, I map granular International Classification of Diseases (ICD) primary diagnosis codes to 285 mutually exclusive and clinically meaningful categories using the Clinical Classification Software (CCS) developed by the Agency for Healthcare Research and Quality (AHRQ).[60] While disruptions to care can hypothetical influence a patient's diagnoses, such as by causing adverse events related to their care, treating primary diagnosis as immutable is reasonable because such complications rising to the level of the patient's *primary* diagnosis should be rare. Complications can be noted as one of the patient's several possible secondary diagnoses, which are studied as an outcome variable.

Since disruptions caused by FDA regulatory enforcement have scope to impact utilization of anesthesia and operating room services, the LASSO model is trained on pre-period hospital visits, which span January 2006 to June 2007. I consider the set of all inpatient and emergency department claims to acute or critical access hospitals among Medicare beneficiaries who are 65 years old and over. Unlike my primary regression sample, I do not restrict to "index" visits (i.e., a beneficiary's first hospital visit in a rolling 12-month period of continuous enrollment in Parts A and B) due to the lack of look-back early in the data. Predictions arising from training on a sample of each patient's first visit in the pre-period are almost perfectly correlated with those from using the full pre-period sample ($\rho > 0.99$).

Model-generated parameters are used to predict SI opioid use among regression sample index visits spanning May 2007 to May 2019 based on beneficiary race, age, sex, and primary diagnosis. The 25th and 75th percentiles of the pre-period predicted probability distribution are used as

[59] While the patient-level predictions used to define this subsample are themselves estimated, I do not explicitly adjust for the associated estimation error because these errors are expected to be orthogonal to the predicted hospital exposure treatment variable. This is empirically supported by the absence of differential pre-trends in patient predicted SI opioid utilization.

[60] The model is trained using 5-fold cross-validation.

cutoffs to identify visits with the highest and lowest ex-ante probabilities of SI opioid use.[61]

Using anesthesia or operating room services as a proxy yields a top quartile of hospital visits that are highly likely to utilize SI opioids within standard care, but it may not be all-inclusive in capturing all clinical scenarios in which these drugs are used. For example, patients who experience severe pain outside of the context of surgical operations are unlikely to have high predicted values. Though this may result in a small number of patients in the bottom quartile whose treatment involves sterile injectable opioids, it nevertheless provides a useful comparison group with much less scope for impacts from disruptions to SI opioid drug supply.

Table D.7 compares the characteristics of regression sample index visits in the top and bottom quartiles (statistics for the pre-period training sample are qualitatively similar), illustrating clear differences. Index visits in the top quartile have greater clinical severity and treatment intensity relative to those in the bottom quartile. For example, top quartile visits are 58 percentage points more likely to involve anesthesia or operating room use and have a 2 percentage point higher rate of in-hospital mortality. These visits are also 63 percentage points more likely to result in inpatient admission (rather than outpatient emergency care), 36 percentage points more likely to be elective admissions, and are associated with an average hospital stay three days longer. Reflecting these differences in care intensity, hospital reimbursements are ten times higher for top-quartile visits, and total hospital charges are nearly seven times higher.

The hospitals that patients receive care from also differ systematically across the top and bottom quartiles. Top quartile visits are 22 percentage points less likely to take place at rural hospitals, 18 percentage points more likely to be at teaching hospitals, and 25 percentage points less to be at the patient's nearest hospital (with respect to their ZIP code centroid). The distribution of hospital ownership and 340B Drug Pricing Program participation is similar across quartiles.

Prediction Quality To evaluate the model's performance in using patient demographics and primary diagnosis to predict SI opioid utilization, I estimate a series of first-stage regressions in which the dependent variable is an indicator for whether the hospital visit involved use of anesthesia or operating room services (i.e., the proxy outcome used to train the model). Since the hospital visits used to train the model (all pre-period hospital visits) and the regression sample (index hospital visits in the study period) are distinct, I estimate the first stages separately for these two samples to quantify the impacts of these compositional differences. To examine the stability of the prediction over time, I also estimate the first-stage regressions separately by calendar in the regression sample. This may be important as the predictive relationship may degrade over time due to changes in clinical practices and billing conventions.

[61] Though two months worth of regression sample index hospital visits were used for the prediction (May and June 2007), no FDA enforcement action in the study period occurred prior to November 2007, so these early index visits are also unaffected by treatment.

Table D.7. Characteristics of Top and Bottom Quartile of Predicted SI Utilization Hospital Visits

	Top Quartile Index Visits		Bottom Quartile Index Visits	
	Mean	(SD)	Mean	(SD)
Visit Characteristics				
Any Anesthesia or OR Utilization	0.59	(0.49)	0.01	(0.12)
In-Hospital Mortality	0.02	(0.13)	0.00	(0.06)
Inpatient Hospitalization	0.77	(0.42)	0.13	(0.34)
Stay Length (Days)	4.40	(4.11)	1.46	(1.58)
Hospital Characteristics				
Rural Hospital	0.14	(0.34)	0.18	(0.39)
Teaching Hospital	0.51	(0.50)	0.43	(0.50)
Nearest Hospital	0.39	(0.49)	0.52	(0.50)
Log(Distance in Miles)	2.32	(1.48)	1.98	(1.52)
340B Program Hospital	0.35	(0.48)	0.34	(0.47)
Non-Profit Hospital	0.73	(0.45)	0.72	(0.45)
Financial Characteristics				
Total Reimbursement (Dollars)	11,846.13	(14,962.42)	1,418.02	(3,792.14)
Total Hospital Charges (Dollars)	47,829.24	(67,297.23)	7,325.29	(14,513.08)

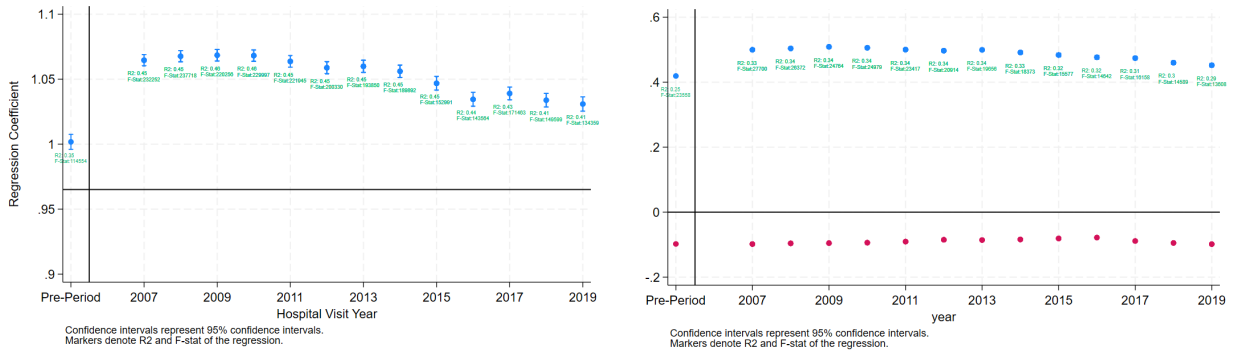
Notes: Table D.7 reports summary statistics of index hospital visits in the top and bottom quartiles of predicted sterile injectable opioid utilization.

Separate specifications include the continuous predicted probability and categorical indicators for whether the visit is in the top or bottom quartile of the prediction distribution (with the middle 50% as the reference group) as independent variables. Coefficients are plotted by year with 95% confidence intervals and R^2 values displayed to assess the strength and stability of the predictive relationship. Standard errors are clustered at the patient ZIP code level, which is the same level of clustering as in the main analysis of treatment effects on patient outcomes.

First-stage regression results displayed in Figure D.2 show that the predicted probability generated by the cross-validated LASSO machine learning algorithm is a strong and stable predictor of actual anesthesia or operating utilization in the regression sample. In the pre-period training sample, the continuous predicted probability measure yields a coefficient of 1.00, an R^2 of 0.35, and an F-statistic of 114,554, indicating strong explanatory power (see panel A of Figure D.2). When applied to the main regression sample of index hospital visits, the predictive relationship strengthens further, with a coefficient of 1.07, R^2 of 0.45, and F-statistic of 232,252 in the first year of the regression sample. This suggests that the prediction algorithm retains and even improves its explanatory power when applied to the more restrictive regression sample consisting of index hospital visits. Importantly, the predictive relationship remains robust: by the final year, 2019, the coefficient remains at 1.03 with an R^2 of 0.41 and F-statistic of 134,359, indicating only a modest

attenuation in predictive power over time.

Consistent with Table D.7, results using top- and bottom-quartile of predicted SI utilization quartiles as independent variables show drastically different patterns of anesthesia and operating room utilization across the two groups (see panel B of Figure D.2). Relative to the omitted middle 50% of the predicted likelihood distribution, hospital visits in the top quartile are 50 percentage points more likely to involve anesthesia or operating room use, while those in the bottom quartile are approximately 10 percentage points less likely. These differences are both large in magnitude and persist across calendar years in the regression sample, indicating that the prediction algorithm effectively isolates hospital visits with high versus low likelihood of SI opioid utilization.



(a) Predictiveness of $\widehat{AnesthesiaOR}$

(b) Predictiveness of Top and Bottom Quartile Indicators

Figure D.2. First Stage of SI Opioid Prediction on Actual Anesthesia or Operating Room Utilization

Notes: Figure D.2 plots the results of several first-stage regressions assessing the performance of the cross-validated linear LASSO model that uses patient demographics and primary diagnosis to predict whether a hospital visit included anesthesia or operating room use. Blue markers in panel A report reports estimated β coefficients from regressions of the form $\text{AnesthesiaOR}_i = \alpha + \beta\widehat{\text{AnesthesiaOR}}_i + u_i$. A separate β is calculated for the pre-period and for each calendar year of regression sample index hospital visits to assess variation in predictiveness over time. In panel B, blue and red markers report the estimated β_1 (blue) and β_2 (red) coefficients from regressions of the form $\text{AnesthesiaOR}_i = \alpha + \beta_1\text{TopQuartile}_i + \beta_2\text{BottomQuartile}_i + u_i$, which are also run separately for the pre-period and each calendar year of regression sample index hospital visits. For consistency with the patient-level analysis, inference is conducted with standard errors clustered at the patient ZIP code level.

Transition from ICD-10 to ICD-9 Primary Diagnosis Codes A nuance of applying the prediction to the full set of regression sample index hospital visits is that the prediction sample from January 2006 to June 2007 uses a mapping of patients' primary diagnosis from the Ninth Revision of International Classification of Diseases (ICD-9), and the healthcare industry transitioned the Tenth Revision of International Classification of Diseases (ICD-10) codes in October 2015. AHRQ created a revised set of Clinical Classifications Software Refined (CCSR) codes for ICD-10, but no crosswalk exists between CCS and CCSR. I therefore construct a hybrid mapping procedure to apply the LASSO model's predictions to index visits in October 2015 and later.

The primary mapping step uses the General Equivalence Mappings (GEMs) crosswalk between ICD-9 and ICD-10 diagnosis codes maintained by CMS. For each of the 141,752 distinct ICD-10 codes, GEMs identifies linked ICD-9s and I then identify their corresponding CCS code. A unique CCS code is associated with each of 15,073 ICD-9 codes and I confirm that the data I use for this procedure covers full set of ICD-9s in the pre-period training sample (except for one that is added to the "no known CCS" category $CCS = 0$). When the non-one-to-one nature of GEMs results in multiple CCS categories associated with a single ICD-10 code, I calculate predicted SI opioid utilization as the maximum predicted probability across all linked CCS codes. Specifically, 94.5% of ICD-10s are mapped to exactly one CCS, 5.1% are mapped to 2 CCSs, 0.2% are mapped to 3 CCSs, and the rest map to 4 or 5. [62]

After applying the GEMs mapping, approximately 93,000 (0.6%) hospital visits in the regression sample are missing a CCS code. These occur exclusively in the ICD-10 era and account for as much as 3% of hospital visits near the end of the study period including prominent diagnoses (e.g., hypertension, pancreatic disease, etc.). Manual review reveals that this missingness arises from a lack of GEMs mappings due to conflicts between ICD-10 to ICD-9 at a granular level of detail, such as because new ICD-10s introduce new concepts or levels of specificity missing from ICD-9s. Since these conflicts do not exist when zooming out to the broader CCS and CCSR codes level, I complement the CMS GEMS-based mapping with fuzzy matching between CCS and CCSR codes among the smaller set of ICD-10 with no GEMS-based CCS. After this step, only about 6,000 (0.05%) of hospital visits are missing a CCS, and the maximum share of hospital visits with a missing prediction in a given month is 0.42%. This is deemed acceptably low.

D.2 Construction of Patient ZIP Code Hospital Visit Probabilities

To reserve the hospital that patients actually received care from as an outcome of the supply disruption caused by FDA regulatory enforcement, I construct patient-level treatment variables using the pre-period (January 2006 to June 2007) empirical hospital visit distribution as a prediction of hospital visited in normal times. For robustness, I also construct pre-period hospital visit distributions specific to the top and bottom quartiles of predicted SI opioid utilization, which may rely on systematically different hospitals according to differences in medical needs. For example, table D.7 shows that patients in the top quartile are more likely to visit a teaching hospital and less likely to visit a rural hospital.

Since quartile-specific hospital visit distributions may rely on small samples in some ZIP

[62] Mapping ICD-10s to ICD-9s using GEMS yields a sample of 13,433 ICD-9s, which is close to the the full set of 15,073 ICD-9s with a valid CCS mapping. There are 45,671 ICD-10s for which the ICD-9 it maps back to has no CCS mapping entirely, but none of these ICD-9s occur in the Medicare claims data as the primary diagnosis, consistent with lower medical importance. Overall, there are 284 CCS codes in the ICD-9 era diagnosis data and 283 CCS codes in the ICD-10 diagnosis era.

codes, I employ an Empirical Bayes (EB) approach that shrinks ZIP code quartile-specific hospital shares toward the distribution of patients in the same risk quartile and residing in the same Hospital Service Area (HSA). The Dartmouth Atlas designed HSAs to group ZIP codes into local healthcare markets where residents receive most of their hospitalizations from the same sets of hospitals. As such, an EB approach using HSA-based priors preserves the spirit of quartile-specific hospital visit patterns while stabilizing estimates in sparse ZIP codes with data from patients in the same local healthcare market.[63]

The EB calculation uses a Dirichlet-multinomial prior smoothing approach. Letting n_z represent the number of pre-period hospital visits by patients residing in ZIP code z and $x_{z,h}$ represent indicator variables for whether a hospital visit occurred at hospital h, I calculate the sample mean visit rate to hospital h at the ZIP code level ($s^{pre}_{z,h} = \frac{1}{n_z}\sum_{n_z} x_{z,j}$) and at the HSA level ($s^{pre}_{hsa,h} = \frac{1}{n_{hsa}}\sum_{n_{hsa}} x_{hsa,j}$). The EB-adjusted mean is then calculated as a convex combination of the ZIP code and HSA-level means, where the influence of the prior depends on the ZIP code sample size and the HSA-level Kish effective sample size. The Kish effective sample size $c_{hsa} = \frac{(\sum_{z\in hsa} n_z)^2}{\sum_{z\in hsa} n_z^2}$ quantifies the extent to which information in the HSA-level mean is evenly distributed across the ZIP codes contained within it. If information is concentrated in only a single ZIP code, the HSA-level mean is treated as less reliable:[64]

$$s^{pre,EB}_{z,h} = \frac{c_{hsa} \times s^{pre}_{hsa,h} + \sum_{n_z} x_{z,h}}{c_{hsa} + n_z}$$

For each ZIP code, the vector of $s^{pre,EB}_{z,h}$ hospital visit shares is normalized such that they sum to 1. The influence of the prior $w_z = \frac{c}{c+n_z}$ is approximately 0.5 on average in the regression sample, and the correlation between the quartile-specific EB hospital and the hospital shares using the full set of hospital visits is 0.87.

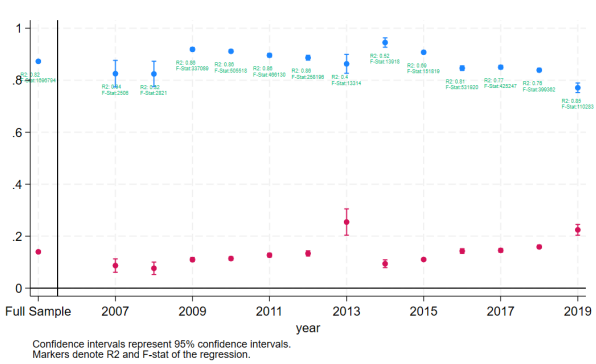
The EB mean $s^{pre,EB}_{z,h}$ is missing in a small number of instances in which the HSA-wide risk quartile hospital visit rate is missing. This occurs both when the patient ZIP code is not mapped to an HSA (ZIP codes appearing only in the Medicare data and not in the Dartmouth Atlas ZIP code to HRR crosswalk) and for HSAs with zero risk quartile hospital visits throughout the entire pre-period across. In such cases, setting the measure as missing is most reasonable as HSA-level

[63] Dartmouth Atlas defines 3,436 HSAs, most of which contain only one hospital. The broader hospital referral region (HRR) concept represents regional health care markets for tertiary medical care, where each HRR contains at least one hospital that performs major cardiovascular procedures and neurosurgery. According to Dartmouth Atlas, HRRs are collections of HSAs based on the region where the greatest proportion of major cardiovascular procedures were performed, with minor modifications to achieve geographic contiguity, a minimum population size of 120,000, and a high localization index. There are 306 hospital referral regions.

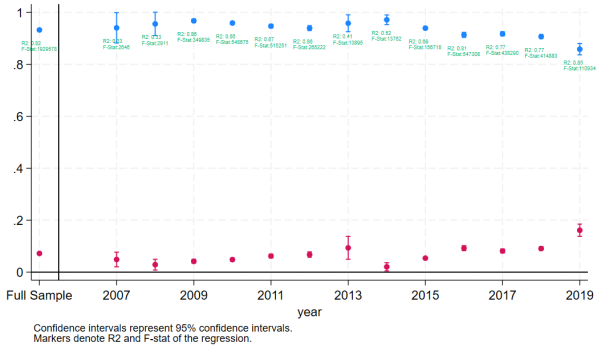
[64] c_{hsa} is greatest when all ZIP codes have the same sample size and shrinks toward 1 if a single ZIP code dominates almost all visits.

data does not provide any information with which to form a prediction.[65]

Prediction Quality Since pre-period hospital shares are calculated from data spanning January 2006 to June 2007 and then applied to a regression sample spanning July 2007 to May 2019, one might worry about degradation in predictiveness as hospital market shares fluctuate due to new hospitals opening, existing hospitals closing, and other natural sources of variation.[66] As such, the predicted hospital exposure measure $\mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau} = \sum_H s^{pre}_{z(i),h} \times \text{HospitalExposure}_{h,\tau}$ may serve as a poor measure of a patient's local hospital environment to a given FDA enforcement shock, especially near the end of the study period.



(a) Predicted Hospital Visit Shares Using All Visits



(b) EB-Adjusted Predicted Hospital Visit Shares Using Top Quartile Visits

Figure D.3. First Stage of Predicted Hospital Exposure on Actual Hospital Exposure

Notes: Figure D.3 plots the results of several first-stage regressions assessing the performance of the patient ZIP code-level predicted hospital exposure, which captures the exposure of drug supply in a patient's local hospital environment to FDA enforcement, in predicting the exposure of the hospital that patients actually visited. Regressions use the stacked sample of top predicted SI opioid utilization quartile index hospital visits used in the patient-level analysis. In panel A, blue and red markers report the estimated β_1 (blue) and β_2 (red) coefficients from distinct regressions of the form $\text{HospitalExposure}_{h(i),\tau} = \alpha + \beta_1 \mathbb{E}\left[\text{HospitalExposure}\right]_{z(i),\tau} + \beta_2 \text{EnforcedMarketShare}_\tau + u_i$, which are run for the full sample and separately for each year of the regression sample to assess the degree to which the predictive relationship persists over time. Comparing β_1 and β_2 allows the assessment of the predictive power of $\mathbb{E}\left[\text{HospitalExposure}\right]$ beyond the overall national exposure of hospital drug volume to the enforcement event. Standard errors are clustered at the patient ZIP code level.

However, first-stage regressions demonstrate that predicted hospital exposure remains highly predictive of the exposure of the hospital actually visited by the patient throughout (see Figure D.3). Regressions use the stacked regression sample of top predicted SI opioid utilization quartile index hospital visits used in the analysis of patient health, quality of care, and healthcare costs.

[65] In practice, missingness due to a lack of risk quartile visits does not take place at all for the top quartile and minimally for the bottom quartile. Instead, the vast majority of missingness is driven by a lack of an HSA map corresponding to a patient ZIP code.

[66] Closings are not a major concern because I re-calculate hospital visit shares based on the set of hospitals that have a valid treatment variable for a particular FDA enforcement event.

The outcome is the exposure the hospital actually visited, and predicted hospital exposure and the national market share of the enforced facility are used as independent variables. Including both predicted exposure and national market share creates a "horse-race" between the two measures, allowing me to assess the performance of predicted hospital exposure beyond the overall exposure of hospital drug volume to the enforcement event.[67] The regression is run separately for each year of the regression sample to assess persistence in the predictive relationship over time.

The results show that the predicted exposure measure is a vastly stronger predictor of actual exposure when using both pre-period hospital visit shares calculated using all hospital visits (panel A) and using the quartile-specific EB-adjusted shares (panel B). Every percentage point of predicted exposure is associated with a highly statistically significant 0.87pp increase in actual exposure. By comparison, a 1pp increase in the national market share of the implicated facility is associated with only a 0.13pp increase in actual exposure.

D.3 Determining ZIP Code Locations and Distances to Hospitals

Hospital travel outcomes are determined by assigning each beneficiary ZIP code, which reflects their mailing addresses in each calendar year, to a latitude-longitude using ZIP code centroids from NBER's ZIP Code Distance database (link). ZIP codes missing from that database are augmented using population-weighted ZIP code centroids from the Department of Housing and Urban Development (HUD, link).[68] This procedure provides a match for 33,193 out of 40,866 (81%) of the ZIP codes in the Dartmouth Atlas Hospital Referral Regions file.

Using the USDA's 2010 Rural-Urban Commuting Area Codes ZIP code file (link) reveals that 96% of the ZIP codes with missing latitude and longitude are "Post Office or large volume customers" while 92% of non-missing ZIP codes are "Zip Code Areas," suggesting that ZIP codes with missing latitude and longitude are primarily single points rather than a larger geographical area with a well-defined population from which a centroid can be calculated. Thus, remaining missing ZIP code latitude and longitudes are imputed as the average across all ZIP codes with the same first 3 digits, which captures all ZIP codes sharing the same sectional center facility that processes the region's mail.[69] Since 3-digit ZIP codes are sufficiently geographically proximate to share the same regional processing facility, taking the average of the locations with the same 3-digit ZIP

[67] The regression does not include a constant term because there are some years with no variation in the national market share of the implicated facility, leading the variable to be omitted when a constant term is used. Using the full regression sample, I find that the constant term is approximately zero in any case, which is equivalent to omitting the constant from the regression.

[68] The median distance between the centroids from the two files when both are available is 1 mile, and the mean is 1.88 miles.

[69] The Post Office glossary defines a sectional center facility as "a postal facility that serves as the processing and distribution center (P&DC) for Post Offices in a designated geographic area as defined by the first three digits of the ZIP Codes of those offices. Some SCFs serve more than one 3–digit ZIP Code range."

codes provides a reasonable approximation of latitude and longitude. This approach results in missing patient latitude and longitudes for only 0.05% of index visits per month on average with a single-month high of 0.1%. This is deemed acceptably low.

Patient-to-hospital distances are calculated as the geodesic distance between the patient ZIP code latitude-longitude centroid and the exact latitude-longitude of the hospital visited, obtained from the Google Maps API.