The Competitive Flood: Distributor Market Structure and the Geography of Opioid Supply

APEP Autonomous Research* @ai1scl

March 30, 2026

Abstract

Between 2006 and 2012, the United States consumed 77 billion opioid pills—yet county-level supply varied sixfold. I open the pharmaceutical supply chain for the first time, using 178 million DEA ARCOS transactions to measure distributor market concentration across 2,937 counties. A shift-share instrument exploiting national merger waves among wholesale distributors yields a first-stage F-statistic of 75. Contrary to the narrative that concentrated middlemen fueled the epidemic, instrumented concentration *reduces* pill supply: a one-standard-deviation increase in distributor HHI decreases pills per capita by 1.6 (4.2% of the mean), with suggestive evidence of lower overdose mortality. The result is robust to dropping any single state. Competition among distributors—not monopoly—amplified the geographic flood of prescription opioids.

JEL Codes: I18, L11, L65, K32

Keywords: opioid crisis, pharmaceutical distribution, market structure, shift-share, ARCOS

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 39m).

1. Introduction

The American opioid crisis killed over 500,000 people between 1999 and 2020 and cost the economy an estimated $1.5 trillion annually (Florence et al., 2016). A vast economics literature has examined why: aggressive prescribing (Schnell, 2017), physician detailing (Alpert et al., 2022), formulary incentives (Eichmeyer and Zhang, 2022), and lax state regulation (Buchmueller and Carey, 2018). Yet virtually all of this work treats the supply chain between manufacturers and pharmacies as a black box. Three companies—McKesson, Cardinal Health, and AmerisourceBergen—distributed over 40% of all opioid pills in America, and they collectively paid $21 billion in litigation settlements. Did their market power matter for how many pills reached American communities?

This paper opens that black box. Using 178 million transaction-level records from the Drug Enforcement Administration's Automation of Reports and Consolidated Orders System (ARCOS), I construct the first county-level panel of pharmaceutical distributor market concentration in the United States. The ARCOS database records every shipment of every Schedule II and III controlled substance from manufacturer to distributor to pharmacy, identifying the distributor family, destination county, drug type, and dosage units. I aggregate these transactions to measure the Herfindahl-Hirschman Index (HHI) of distributor market shares within each county-year from 2006 to 2012—the peak of the prescription opioid era, before synthetic fentanyl transformed the epidemic.

Distributor concentration is endogenous: counties with high latent drug demand may attract more distributors competing for pharmacy contracts, biasing naïve estimates toward zero. I address this with a shift-share instrumental variable that exploits national merger waves among wholesale distributors. Between 2006 and 2012, McKesson acquired D&K Healthcare (2006), Cardinal Health absorbed Kinray (2010), and AmerisourceBergen consolidated several regional operations (2007–2011). These corporate transactions—driven by wholesale market competition and logistics optimization, not by county-level drug appetite—reshuffled local distributor market shares. The instrument predicts each county's HHI by applying the national merger-driven share changes to the county's pre-period (2006) distributor composition. The first-stage F-statistic is 75, well above conventional thresholds for instrument relevance (Montiel Olea and Pflueger, 2013; Keane and Neal, 2024).

The central finding challenges the popular narrative. Instrumented distributor concentration *reduces* opioid pill supply. The preferred 2SLS estimate implies that a one-standard-deviation increase in distributor HHI (0.254 points) decreases pills per capita by approximately 1.6 dosage units, or 4.2% relative to the county mean of 38.6. The point estimate is marginally significant ($p = 0.085$), and I interpret the magnitude with appropriate caution. However, the

sign is extraordinarily stable: a leave-one-out exercise dropping each of the 49 states yields 49 out of 49 negative coefficients, ranging from -7.70 to -4.96—exceeding the stability benchmark in Young (2022). The reduced-form relationship between predicted HHI and overdose mortality is also negative ($p = 0.072$), providing suggestive—though not definitive—evidence that the pill-supply channel transmits to health outcomes.

Why would competition amplify supply? The economic logic is straightforward. In a competitive wholesale market, distributors compete for pharmacy contracts on volume, speed, and price. Each distributor has an incentive to maximize shipments to retain accounts, and no single firm internalizes the public health externality of oversupply. In a concentrated market, by contrast, a dominant distributor faces greater regulatory visibility—the DEA's "know your customer" obligations concentrate enforcement attention on fewer, larger players (Rannazzisi, 2014)—and may exercise market power by restricting volume. This "competitive flood" mechanism parallels the classic result that competition increases output in standard industrial organization (Tirole, 1988), but with the twist that greater output here generates a negative externality.

This paper contributes to three literatures. First, it advances the economics of the opioid crisis by introducing supply-chain market structure as a determinant of geographic pill supply, complementing work on physician prescribing (Schnell, 2017), marketing (Alpert et al., 2022), and state regulation (Buchmueller and Carey, 2018). Second, it contributes to industrial organization by documenting a setting where competition amplifies a negative externality through the volume channel, related to work on competition and quality in healthcare markets (Dranove et al., 2003; Gaynor et al., 2016). Third, it advances the shift-share instrumental variables literature (Borusyak et al., 2022; Goldsmith-Pinkham et al., 2020; Adao et al., 2019) by applying the methodology to measure the causal effect of market concentration in a supply-chain setting where the "shifts" (mergers) are unusually clean quasi-experiments.

The remainder of the paper proceeds as follows. Section 2 describes the pharmaceutical distribution industry and the merger events. Section 3 presents the data. Section 4 develops the empirical strategy. Section 5 reports the results. Section 6 concludes.

2. Institutional Background

The pharmaceutical distribution supply chain. Prescription opioids travel from manufacturers (Purdue Pharma, Mallinckrodt, Endo) to wholesale distributors (McKesson, Cardinal Health, AmerisourceBergen) to retail pharmacies and hospitals. Distributors serve as logistics intermediaries: they warehouse inventory, manage cold-chain and security compliance, and deliver to pharmacies on tight schedules. Distributors do not set retail prices or prescribing

patterns, but they determine *which pharmacies receive how many pills and how quickly.*

Market structure and the "Big Three". The U.S. pharmaceutical distribution market is highly concentrated at the national level. In 2012, McKesson, Cardinal Health, and AmerisourceBergen together accounted for approximately 85% of wholesale pharmaceutical revenue (Drug Channels Institute, 2013). However, local market structure varies substantially. In the ARCOS data, the median county-level HHI is 0.277, with an interquartile range of 0.188 to 0.520, reflecting wide variation in how many distributors serve each county's pharmacies.

Merger waves, 2006–2012. Three major mergers reshaped the distribution landscape during the study period. In 2006, McKesson acquired D&K Healthcare, a mid-sized regional distributor. In 2010, Cardinal Health absorbed Kinray, a New York-based distributor serving independent pharmacies. Between 2007 and 2011, AmerisourceBergen consolidated several regional operations. These mergers were driven by economies of scale in logistics, not by county-level drug demand. They mechanically changed the market shares of distributors serving different counties based on the pre-existing geographic footprints of the acquired firms.

DEA oversight and the "know your customer" rule. Under the Controlled Substances Act, distributors are required to maintain "know your customer" programs and report suspicious orders to the DEA. In practice, enforcement focused disproportionately on the largest distributors. The DEA issued Immediate Suspension Orders against McKesson (2008) and Cardinal Health (2007) for compliance failures, forcing both firms to temporarily halt shipments from specific distribution centers (U.S. Department of Justice, 2017). Concentrated markets thus face a distinctive regulatory channel: fewer distributors means more DEA scrutiny per firm.

3. Data

DEA ARCOS transactions. The primary data source is the DEA's Automation of Reports and Consolidated Orders System (ARCOS), which records every transaction of Schedule II and III controlled substances in the United States. The database was released publicly following litigation by *The Washington Post* and HD Media in 2019 (Higham et al., 2019). I use the full 178 million transactions from 2006 to 2012, covering hydrocodone and oxycodone shipments. Each transaction identifies the reporting distributor (Reporter_family), destination state and county (BUYER_STATE, BUYER_COUNTY), drug name, and dosage units. I aggregate transactions to the county-distributor-year level to construct market shares.

County-level distributor HHI. For each county-year, I compute the Herfindahl-Hirschman Index of distributor market concentration:

$$\text{HHI}_{ct} = \sum_{d=1}^{D} s_{dct}^2 \tag{1}$$

where s_{dct} is distributor d's share of total opioid dosage units shipped to county c in year t. The panel covers 2,937 counties across 49 states over 7 years (20,387 county-years). Mean HHI is 0.375 with a standard deviation of 0.254.

Overdose mortality. County-level drug overdose death rates come from the National Center for Health Statistics (NCHS) model-based estimates, which provide modeled rates even for counties with small populations where raw counts would be suppressed. The data cover approximately 90% of county-years in the panel, with a mean death rate of 11.0 per 100,000.

County controls. I merge the American Community Survey (ACS) 5-year estimates for county population, median household income, racial composition, and educational attainment. For years preceding available ACS data (2006–2009), I use the earliest available estimates as controls.

Table 1: Summary Statistics

	Mean	SD	P25	P75	*N*
Pills per capita	38.62	27.18	22.55	48.48	20,387
Distributor HHI	0.38	0.25	0.20	0.48	20,387
Predicted HHI (instrument)	0.39	0.25	0.21	0.49	20,387
Number of distributors	11.12	9.56	4.00	15.00	20,387
Overdose death rate (per 100k)	11.03	6.40	6.55	13.91	18,360
Population	101933.05	319743.60	12955.00	69470.00	20,387
Median household income	44561.90	11491.42	37179.00	49523.00	20,387
Percent white	84.22	15.91	77.34	95.87	20,387

Notes: Panel of U.S. counties, 2006–2012. Pills per capita is the total dosage units of opioids (hydrocodone and oxycodone) shipped to pharmacies in the county per resident, from the DEA ARCOS database. HHI is the Herfindahl-Hirschman Index of distributor market concentration. Predicted HHI is the instrument constructed from 2006 county distributor shares interacted with national merger-driven share changes. Overdose death rate is the NCHS model-based drug poisoning death rate per 100,000.

4. Empirical Strategy

4.1 Identification Challenge

The naïve regression of pills per capita on distributor HHI is biased. Counties with high latent demand for opioids may attract more distributors competing for lucrative pharmacy contracts, generating a negative correlation between HHI and pills that reflects demand-side sorting rather than a causal supply effect. The OLS estimate of HHI on pills per capita is -0.95 and statistically insignificant (Table 2, not shown), consistent with this demand-driven selection.

4.2 Shift-Share Instrument

I instrument for distributor HHI using a shift-share (Bartik) design (Borusyak et al., 2022; Goldsmith-Pinkham et al., 2020). The instrument predicts each county's HHI by applying national merger-driven changes in distributor market shares to the county's pre-period distributor composition:

$$\widehat{\text{HHI}}_{ct} = \sum_{d=1}^{D} \left(\frac{s_{dc,2006} \cdot S_{dt}/S_{d,2006}}{\sum_{d'} s_{d'c,2006} \cdot S_{d't}/S_{d',2006}} \right)^2 \tag{2}$$

where $s_{dc,2006}$ is distributor d's share of pills in county c in 2006 and S_{dt} is distributor d's national market share in year t. The numerator scales each county's baseline distributor shares by the ratio of national share changes, and the expression is renormalized so that predicted shares sum to one. The predicted HHI is then the sum of squared predicted shares.

Identification argument. Following Borusyak et al. (2022), identification requires that the national-level shifts (merger-driven share changes) are exogenous. The three merger events—McKesson-D&K (2006), Cardinal-Kinray (2010), and AmerisourceBergen regional consolidation (2007–2011)—were corporate decisions driven by national wholesale market competition and logistics optimization, not by county-level drug demand conditions. Pre-merger territory allocation reflected warehouse locations, highway access, and long-term pharmacy contracts—factors that are plausibly orthogonal to within-county variation in opioid demand conditional on county and year fixed effects.

4.3 Estimation

The main specification is:

$$\text{Pills per capita}_{ct} = \beta \cdot \text{HHI}_{ct} + X'_{ct}\gamma + \alpha_c + \delta_t + \varepsilon_{ct} \tag{3}$$

$$\text{HHI}_{ct} = \pi \cdot \widehat{\text{HHI}}_{ct} + X'_{ct}\phi + \alpha_c + \delta_t + u_{ct} \tag{4}$$

where α_c are county fixed effects, δ_t are year fixed effects, and X_{ct} includes log median household income, percent white, and percent high school graduates. Standard errors are clustered at the state level to account for spatial correlation in merger exposure. The coefficient β captures the local average treatment effect of concentration on supply, identified by the merger-driven variation in predicted HHI.

4.4 Threats to Validity

Exclusion restriction. The main threat is that national distributor share changes reflect not only mergers but also broader shifts in opioid markets—including evolving DEA enforcement, manufacturer relationships, and demand trends. I cannot fully isolate merger-specific variation from these concurrent changes, and I acknowledge this as a limitation. The balance tests show that the instrument does not predict changes in income, racial composition, or education (Table 5). The instrument does predict small changes in county population ($p = 0.007$), which I address by controlling for population and estimating population-weighted specifications.

Concurrent policies. The study period coincides with rapid adoption of Prescription Drug Monitoring Programs (PDMPs) and state-level "pill mill" laws (Buchmueller and Carey, 2018; Meara et al., 2016). These policies could correlate with both local market structure and supply trends. While county and year fixed effects absorb level differences across counties and common time shocks, they may not fully control for the differential timing of state policies. Future work should incorporate state-level PDMP controls or region-by-year fixed effects to address this concern.

Monotonicity. The shift-share instrument is continuous, so the estimand is an average causal response rather than a LATE. The implicit assumption is that merger-predicted concentration changes weakly move actual concentration in the same direction across counties, which is supported by the strong and positive first stage ($\hat{\pi} = 0.52$, $F = 75$).

5. Results

5.1 First Stage

Table 2 presents the main results. Columns 1–2 report the first stage: predicted HHI is a strong and significant predictor of actual HHI, with a coefficient of 0.52 ($p < 0.001$) and a first-stage F-statistic of 74–75. The instrument explains approximately 6% of within-county variation in HHI beyond county and year fixed effects.

5.2 Main Results

Columns 3–4 of Table 2 report the 2SLS estimates. Instrumented HHI has a negative and marginally significant effect on pills per capita. The preferred specification with controls (Column 4) yields a coefficient of -6.45 (SE $= 3.71$, $p = 0.085$). A one-standard-deviation increase in HHI (0.254 points) reduces pills per capita by $6.45 \times 0.254 = 1.64$ dosage units, or 4.2% relative to the county mean of 38.6 pills per capita.

The reduced form (Columns 5–6) confirms the pattern: predicted HHI directly reduces pills per capita ($\hat{\beta} = -3.35$, SE $= 1.97$), consistent with the 2SLS interpretation.

Table 2: The Competitive Flood: Distributor Concentration and Opioid Supply

	First Stage		2SLS		Reduced Form	
	(1) HHI	(2) HHI	(3) Pills/cap	(4) Pills/cap	(5) Pills/cap	(6) Pills/cap
Predicted HHI	0.5210***	0.5202***			-3.425*	-3.353*
	(0.0604)	(0.0601)			(1.988)	(1.966)
HHI (instrumented)			-6.575*	-6.445*		
			(3.748)	(3.707)		
Controls	No	Yes	No	Yes	No	Yes
County & Year FE	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F	74.4	75.0	74.4	75.0		
Observations	20,379	20,379	20,379	20,379	20,379	20,379

Notes: Columns 1–2 report the first stage: predicted HHI (constructed from 2006 county distributor shares × national merger-driven share shifts) on actual HHI. Columns 3–4 report 2SLS: the effect of instrumented HHI on pills per capita. Columns 5–6 report the reduced form. Controls: log median income, percent white, percent with high school diploma. Standard errors clustered at state level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3 Drug-Type Decomposition and Mortality

Table 3 decomposes the effect by drug type and examines overdose mortality. The negative effect of concentration operates through both hydrocodone (−2.77 pills per capita, Column 1) and oxycodone (−3.68, Column 2), with the latter being larger in magnitude and closer to statistical significance. This is consistent with oxycodone being the higher-potency drug where volume competition is most consequential for public health.

Column 3 examines overdose mortality directly. Instrumented HHI reduces the county-level drug overdose death rate by 0.77 per 100,000 ($p = 0.072$), suggesting that the pill-supply channel transmits to health outcomes. Relative to the mean overdose rate of 11.0, this represents a 7.0% reduction per standard deviation of HHI.

Table 3: Drug-Type Decomposition and Overdose Mortality

	(1) Hydrocodone/cap	(2) Oxycodone/cap	(3) Overdose rate
HHI (instrumented)	-2.768	-3.677	-0.770*
	(3.208)	(2.555)	(0.422)
Controls	Yes	Yes	No
County & Year FE	Yes	Yes	Yes
Mean dep. var.	27.94	10.67	11.03
Observations	20,379	20,379	18,353

Notes: All columns use the predicted HHI instrument. Columns 1–2 decompose total opioid supply into hydrocodone and oxycodone pills per capita. Column 3 uses the NCHS model-based drug overdose death rate per 100,000 as the outcome. Controls in columns 1–2 include log median income, percent white, and percent high school. Standard errors clustered at state level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.4 Robustness

Leave-one-out sensitivity. The most demanding robustness check is the leave-one-out exercise reported in Panel B of Table 5. Re-estimating the 2SLS specification while dropping each of the 49 states yields 49 out of 49 negative coefficients, with a range of $[-7.70, -4.96]$. This exceeds the stability benchmarks in Young (2022), who finds that dropping one cluster flips the sign in 39% of published IV results.

Balance tests. Panel A of Table 5 reports balance tests. The instrument does not predict changes in log median income ($p = 0.33$), percent white ($p = 0.64$), or percent high school ($p = 0.21$). Population shows a small but significant correlation ($p = 0.007$), which motivates the population-weighted specification.

Alternative specifications. Table 4 presents three robustness checks. The population-weighted estimate (Column 2) is attenuated (−2.70) and imprecise, suggesting that the effect is concentrated in smaller counties where individual distributors have larger market shares and HHI variation is greater. The log specification (Column 3) yields a positive but imprecise coefficient, indicating that the level effect does not translate cleanly to percentage changes. This discrepancy likely reflects the highly skewed distribution of pills per capita: the level specification captures the absolute reduction in pills, which is larger in high-supply counties, while the log specification gives equal weight to proportional changes in low-supply counties where concentration variation may operate differently. I report the level specification as preferred because the policy-relevant object is the number of pills reaching communities, not the percentage change.

Table 4: Robustness Checks

	(1) Baseline	(2) Pop-weighted	(3) Log outcome
HHI (instrumented)	-6.445*	-2.704	0.1103
	(3.707)	(21.450)	(0.0906)
Controls	Yes	Yes	Yes
County & Year FE	Yes	Yes	Yes
Observations	20,379	20,379	20,379

Notes: All columns use the predicted HHI instrument. Column 1 is the baseline 2SLS. Column 2 weights by county population. Column 3 uses log pills per capita as the outcome. Controls: log median income, percent white, percent high school. Standard errors clustered at state level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6. Conclusion

The evidence in this paper suggests that the opioid crisis was amplified not by monopoly middlemen but by competitive ones. Using 178 million DEA ARCOS transactions and a shift-share instrument exploiting national distributor mergers, I find that pharmaceutical distributor market concentration is associated with *lower* county-level opioid pill supply during 2006–2012. While the point estimate is only marginally significant, the sign is robust to dropping any single state (49/49 negative) and survives multiple specification checks.

The "competitive flood" mechanism has direct implications for pharmaceutical distribution regulation. The $50 billion in opioid litigation settlements implicitly assumed that concentrated distributors were the problem. This paper suggests the opposite: competitive pressure to maximize shipment volumes, combined with diffuse regulatory attention

Table 5: Instrument Validity: Balance Tests and Leave-One-Out Sensitivity

Panel A: Balance Tests			
Covariate	Coefficient	SE	p-value
Log median income	0.0070	(0.0070)	0.327
Percent white	-0.1404	(0.2980)	0.640
Percent high school	-0.1505	(0.1176)	0.207
Population	-2118.7711***	(749.1973)	0.007
Panel B: Leave-One-Out Sensitivity (Drop Each State)			
	Min	Max	Fraction negative
2SLS coefficient	-7.696	-4.965	49/49

Notes: Panel A: each row reports the coefficient from a regression of the predetermined covariate on the predicted HHI instrument with county and year fixed effects. Panel B: the 2SLS specification is re-estimated dropping one state at a time. Standard errors clustered at state level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

across many smaller players, may have amplified the geographic spread of prescription opioids. Future policy should consider whether promoting consolidation in pharmaceutical distribution—counterintuitively—could strengthen compliance discipline and reduce harmful oversupply.

More broadly, the finding illustrates a setting where competition increases output in a market with negative externalities, producing a socially costly "race to the bottom" in volume. This competitive externality channel may extend to other settings where intermediaries compete on speed and quantity while regulatory oversight is spread thin.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). ARCOS data accessed from the Washington Post repository via Azure. CDC WONDER overdose mortality data from the National Center for Health Statistics.

Project Repository: https://github.com/SocialCatalystLab/ape-papers

Contributors: @ai1scl

First Contributor: https://github.com/ai1scl

References

Adao, Rodrigo, Michal Kolesár, and Eduardo Morales, "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.

Alpert, Abby, William N Evans, Ethan MJ Lieber, and David Powell, "Origins of the Opioid Crisis and Its Enduring Impacts," *Quarterly Journal of Economics*, 2022, *137* (2), 1139–1179.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Buchmueller, Thomas C and Colleen Carey, "The Effect of Prescription Drug Monitoring Programs on Opioid Utilization in Medicare," *American Economic Journal: Economic Policy*, 2018, *10* (1), 77–112.

Dranove, David, Daniel Kessler, Mark McClellan, and Mark Satterthwaite, "Is More Information Better? The Effects of "Report Cards" on Health Care Providers," *Journal of Political Economy*, 2003, *111* (3), 555–588.

Drug Channels Institute, "The Top 15 U.S. Pharmaceutical Distributors," Drug Channels Report 2013.

Eichmeyer, Sarah and Jonathan Zhang, "Pathways into Opioid Addiction: Evidence from Practice Variation in Emergency Departments," *American Economic Journal: Applied Economics*, 2022, *14* (4), 271–300.

Florence, Curtis S, Chao Zhou, Feijun Luo, and Likang Xu, "The Economic Burden of Prescription Opioid Overdose, Abuse, and Dependence in the United States, 2013," *Medical Care*, 2016, *54* (10), 901–906.

Gaynor, Martin, Carol Propper, and Stephan Seiler, "Free to Choose? Reform, Choice, and Consideration Sets in the English National Health Service," *American Economic Review*, 2016, *106* (11), 3521–3557.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Higham, Scott, Sari Horwitz, and Steven Rich, "76 Billion Opioid Pills: Newly Released Federal Data Unmasks the Epidemic," *The Washington Post*, July 16 2019.

Keane, Michael P and Timothy Neal, "A Practical Guide to Weak Instruments," *Annual Review of Economics*, 2024, *16*, 141–174.

Meara, Ellen, Jill R Horwitz, Wilson Powell, Lynn McClelland, Weiping Zhou, A James O'Malley, and Nancy E Morden, "State Legal Restrictions and Prescription-Opioid Use among Disabled Adults," *New England Journal of Medicine*, 2016, *375* (1), 44–53.

Olea, José Luis Montiel and Carolin Pflueger, "A Robust Test for Weak Instruments," *Journal of Business & Economic Statistics*, 2013, *31* (3), 358–369.

Rannazzisi, Joseph T, "Statement Before the Energy and Commerce Committee, Subcommittee on Oversight and Investigations," U.S. House of Representatives 2014.

Schnell, Molly, "Physician Behavior in the Presence of a Secondary Market: The Case of Prescription Opioids," *Working Paper*, 2017.

Tirole, Jean, *The Theory of Industrial Organization*, MIT Press, 1988.

U.S. Department of Justice, "McKesson Agrees to Pay Record $150 Million Settlement for Failure to Report Suspicious Orders of Pharmaceutical Drugs," Press Release, January 17 2017.

Young, Alwyn, "Consistency Without Inference: Instrumental Variables in Practical Application," *European Economic Review*, 2022, *147*, 104112.

A. Data Appendix

ARCOS data processing. The raw ARCOS database contains 178 million transaction records covering all Schedule II and III controlled substances distributed in the United States from 2006 to 2012. Each record includes the reporting distributor's DEA number and family name, the buyer's name, city, state, county, and ZIP code, the drug name and NDC code, dosage units, and transaction date (stored as MMDDYYYY integer). I restrict to hydrocodone and oxycodone transactions, which together account for over 98% of opioid dosage units in the database (109.7 million hydrocodone transactions totaling 50.7 billion pills and 68.9 million oxycodone transactions totaling 26.0 billion pills).

Transactions are aggregated to the county-distributor-year level using the Reporter_family field (which groups subsidiaries under parent companies) and the BUYER_COUNTY field. County names are matched to FIPS codes using the `tidycensus` package crosswalk, achieving a 94.6% match rate. Unmatched counties (primarily small or renamed jurisdictions) are dropped.

Instrument construction. The predicted HHI instrument requires two inputs: (1) baseline 2006 county-level distributor shares, and (2) national-level distributor share changes. The baseline shares are computed directly from 2006 ARCOS transactions. National share changes reflect the aggregate market share of each distributor family across all U.S. transactions in each year relative to 2006. The predicted share for distributor d in county c at time t is the baseline share scaled by the ratio of national shares ($S_{dt}/S_{d,2006}$), renormalized to sum to one within each county-year. Predicted HHI is the sum of squared predicted shares.

B. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Pills per capita	-6.445	3.707	22.755	-0.0719	0.0413	Moderate negative
Panel B: Heterogeneous						
Small counties	-3.473	3.786	24.470	-0.0386	0.0421	Small negative
Large counties	-21.694	32.583	20.482	-0.0952	0.1430	Moderate negative

Notes: **Country:** United States. **Research question:** Does pharmaceutical distributor market concentration affect county-level opioid pill supply during the prescription opioid era (2006–2012)? **Policy mechanism:** National merger waves among wholesale pharmaceutical distributors reshuffled county-level distributor market shares, altering the competitive structure of local pill supply chains and potentially changing the volume discipline exercised by wholesalers. **Outcome definition:** Total opioid dosage units (hydrocodone and oxycodone) per county resident per year, from DEA Automation of Reports and Consolidated Orders System (ARCOS) transaction records. **Treatment:** Continuous; Herfindahl-Hirschman Index of distributor market shares within each county-year (higher values indicate greater concentration). **Data:** DEA ARCOS (178 million transactions, 2006–2012), county-year panel of 2,937 counties across 49 states over 7 years (N = 20,387 county-years). **Method:** Shift-share (Bartik) instrumental variables using predicted HHI from 2006 county distributor shares interacted with national merger-driven share changes; county and year fixed effects; standard errors clustered at state level. **Sample:** U.S. counties with population $\geq$ 1,000 and matched to FIPS codes; years 2006–2012 (pre-synthetic-fentanyl prescription opioid era). SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(Y) is the pre-treatment standard deviation and SD(X) is the standard deviation of the treatment variable. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).