The AI Investment Boom and Housing Wealth: Evidence from Data Centers

Henrik Cronqvist, Rui Dai, Mitch Warachka, and Frank Yu*

This draft: August 14, 2026

Abstract

Data centers are the largest corporate-investment boom of the current era. Using a new facility-level database of U.S. data centers in a staggered difference-in-differences design, we examine who captures the financial windfall. Data center counties expand their tax bases and return part of the gain to residents: property-tax rates fall, and the cut capitalizes into homes, households' most valuable asset, whose price rises 4.5%. Capitalization is insignificant where the tax base is riskier, as for speculative crypto facilities, and the cuts are larger where local government is less transparent, consistent with officials paying the surplus as the "price of consent." The estimates survive entropy balancing, a cancelled-project placebo, and a border county-pair design. In aggregate, data centers have created over $100 billion in housing wealth while delivering little to renters, bondholders, or schools. The paper is among the first to show how technology investment and its risks are transmitted into household portfolios.

JEL Classification: G12; G35; G51; R31; D72
Keywords: Data centers; corporate finance; municipal finance; household finance

*Cronqvist and Warachka: Chapman University, George L. Argyros College of Business and Economics (hcronqvist@chapman.edu and mwarachka@chapman.edu); Dai: Wharton Research Data Services (rdai@wharton.upenn.edu); Yu: China Europe International Business School (fyu@ceibs.edu).

I Introduction

Data centers are the largest corporate-investment boom of the current era: as the computing demands of artificial intelligence accelerate, the buildout is on course to exceed, as a share of GDP, the railroads, electrification, and the interstate highways (Van Nieuwerburgh, 2026).[1] Today's average data center carries an assessed value in the hundreds of millions of dollars but employs almost no one: a $500 million facility employs perhaps fifty workers, an automobile plant of the same cost two thousand. It is a "million-dollar plant" (Greenstone and Moretti, 2003) without the (direct) employment effects, so what it delivers to a community is not wages but a financial windfall. Take the case of Google, which is building large data centers in Columbia County, Georgia. Before the first building was even finished, county officials had written legislation and announced that the projected revenue would eliminate property taxes for all residents. That is, a corporate siting decision may end the oldest tax its community's residents pay. Who captures the fiscal surplus a data center creates, and how markets price such a gain, are the questions of this paper.

Whether and how officials distribute any surplus back to the community's residents (voters) is a question of political economy. Hosting a data center is increasingly politically costly: a vast majority of Americans say they would oppose one near their home,[2] and siting is contested through lawsuits, moratoria, and hundreds of state bills. The costs are conspicuous and concentrated near the site, from noise to grid load, while the fiscal gain arrives county-wide as an abstraction on the county's balance sheet. Officials mindful of recalls and re-election therefore have strong reason to "buy" votes, and the property tax is a Coasean exchange: concentrated local costs met with a county-wide compensating transfer (Coase, 1960). Spending the windfall is nearly invisible to voters, the flypaper effect (Hines and Thaler, 1995), while a property-tax rate cut is salient and directly affects household portfolios. The tax is the levy households resent the most (Cabral and Hoxby, 2012), so when the relief recurs, it is credited to the incumbent politicians, and comes from the very base that the data center expands. In this sense, the payout is the "price of consent," and its

[1] According to industry estimates and news reports, global data-center capital expenditure passes $1 trillion in 2026, with the hyperscalers budgeting about $750 billion, and cumulative outlays projected to reach $3–4 trillion by 2030. The U.S. hosts more than 5,000 data centers, with roughly 50 gigawatts of capacity installed, some 36 gigawatts under construction, and announced projects that would add as much as 200 gigawatts by 2032.

[2] A 2026 Gallup poll finds seven in ten Americans oppose the construction of an AI data center in their local area. In the Pew Research Center's 2026 survey of data-center attitudes, tax revenue is among the few local impacts Americans view favorably on net.

terms should reflect the bargain that produces it: how large the new tax base is, how persistent it is expected to be, and how much residents trust the local government that holds it. Those terms organize our paper's empirical tests.

Testing these hypotheses means following the cash flows from a data center siting decision to county balance sheets and on to household portfolios. We assemble a new data set for this purpose from more than a dozen sources. The core is the facility-level S&P 451 Research Data Center Database, verified facility by facility and supplemented with hand-collected records: 4,505 verified facilities, 2,881 of them operating, with locations, capacities, and operators. We follow the money with the Censuses of Governments and audited financial reports licensed from MuniSpot, which cover roughly ninety thousand local governments and yield county-year effective property-tax rates. We price the housing market with home values and rents from the American Community Survey (ACS) and Zillow, and we observe the county's other claimants through municipal bond issues and trades from SDC and the Municipal Securities Rulemaking Board (MSRB) and school staffing records from the National Center for Education Statistics (NCES) and student achievement from the Stanford Education Data Archive, plus property-tax abatement disclosures, building permits, and governance measures from government form to reporting certification. The result is a panel that observes, for nearly 3,000 U.S. counties over 15 years, the arrival of AI computing capital, the public finances it feeds, and the housing market that prices it. To our knowledge, no other data set joins the county's levy, base, and tax rate to the facilities that expand them and the homes that price them.

Our empirical strategy is a staggered difference-in-differences comparison of counties that gain a data center with counties verified never to host any facility. Because data centers cluster within and across counties, treatment is a county's cumulative installed capacity, and because staggered adoption can distort least squares, our estimates are confirmed with group-time estimators built for it (Callaway and Sant'Anna, 2021). The central identification challenge is site selection: data centers locate where power, fiber, and land are available. Data center counties are indeed larger and wealthier than others, but they are not on steeper income or home-value growth paths before entry, and entropy balancing reweights the control group to the treated counties' pre-period growth trajectories. Inference then confronts chance and interference directly. Randomization inference reassigns entire capacity histories across counties, asking how often chance would reproduce our

estimates. A border-pair design compares each data center county with its contiguous neighbors, holding local labor and housing markets fixed. And counties where a data center was announced but never built supply placebos with the same pre-trends and fundamentals but no tax base change. Alternative specifications, samples, and additional controls leave the results unchanged.

We report support for what we dub the "payout hypothesis," which starts on the county's balance sheet. After a data center arrives, the county government's property-tax base grows about 1.2 percentage points per year faster than in matched control counties. For counties receiving a fiscally material facility, the excess growth is about 3.8 percentage points per year, or about half a million dollars a year in added collections for the median such county. Effective property-tax rates fall with the capital installed, by 0.117 percentage points per gigawatt; for a campus of representative size, the cut is approximately 7% of the median county's rate. That is far from the elimination select case studies allude to, but the average is still economically very meaningful. The housing market responds as asset pricing and capitalization predict: home prices rise 4.5%, a gain of about $12,000 for the median home, and the increase is county-wide rather than concentrated near the facility.

How much is capitalized depends on whether the tax base is expected to last, and the market prices such risk factors. The 4.5% response is the durable-base estimate, and diluting the sample with speculative facilities halves it. The market also prices the risk's resolution: two-thirds of the increase arrives around the announcement, with the remainder when operations begin. Counties with speculative, exit-prone cryptocurrency mining data centers capitalize essentially nothing, and the crypto crash of 2022 produced no reversal, because little had been capitalized in the first place. Counties where an announced project is never built show no lasting gain. Our county governance results further support the price of consent: property-tax cuts are larger where county government is less transparent, consistent with lower trust raising the price, as residents who cannot verify how a retained surplus is spent demand payment in the currency they can verify, their own tax bill. In sharp contrast to the effects on household portfolios, the boom delivers little to renters, bondholders, or schools: rents do not move, and we find no significant effect on new-issue bond spreads or on school metrics, from students' lagged standardized test scores to teacher salaries.

The paper contributes to several literatures in financial economics. The first contribution bridges corporate finance and municipal finance, an intersection that is under-researched. The data-center

buildout is investment at a scale that recurs across technological revolutions (Pástor and Veronesi, 2009), in booms that have often rewarded the investing firms' own shareholders poorly (Titman, Wei and Xie, 2004; Cooper, Gulen and Schill, 2008). The capital leaves a local footprint beyond the firm itself (Dougal, Parsons and Titman, 2015). The municipal finance literature has approached local governments largely from the securities side, through borrowing costs, credit, and disclosure (Butler, Fauver and Mortal, 2009; Schwert, 2017; Adelino, Cunha and Ferreira, 2017), whereas the arrival of the capital confronts them with a question from the real side: what to do with an unexpected and recurring addition to the tax base. Firms that receive such cash windfalls tend to keep them rather than pay them out (Blanchard, Lopez-de-Silanes and Shleifer, 1994). We analyze the county's version of the choice with the same logic of free cash flow and payout policy (Jensen, 1986; Stulz, 1990; La Porta, Lopez-de Silanes, Shleifer and Vishny, 2000), treating the decision to retain or distribute the windfall as one would a firm's.[3]

The second contribution brings a new transmission channel to household finance. The home dominates most household portfolios (Campbell, 2006; Cocco, 2005), and the wealth shocks the literature studies reach it through labor income, asset returns, or credit (Mian, Rao and Sufi, 2013; Feyrer, Mansur and Sacerdote, 2017). In our setting the shock arrives through the recurring tax bill on the home, a claim whose horizon extends over the life of the asset (Giglio, Maggiori and Stroebel, 2015). The literature descending from Oates (1969) estimates how fiscal variables capitalize into house prices (Cellini, Ferreira and Rothstein, 2010; Hilber, 2017); we ask, in addition, how the strength of that response varies with the riskiness of the technology behind the tax base, connecting property markets to the pricing of long-dated risks in housing and municipal securities (Bernstein, Gustafson and Lewis, 2019; Painter, 2020). The capitalization coefficient becomes an equilibrium price of fiscal risk rather than a constant of nature.

The remainder of the paper proceeds as follows. Section II develops the conceptual framework and states our hypotheses, which the model in Appendix A formalizes. Section III describes the matched data sets and the construction of the variables. Section IV presents the empirical methodology. Section V reports the results, and Section VI concludes.

[3]A substantial economics literature studies what communities pay to attract corporate capital (Greenstone, Hornbeck and Moretti, 2010; Slattery, 2025). We ask the complementary question: what communities do with what the capital pays them.

II Conceptual Framework and Hypotheses

This section develops a framework at the intersection of corporate finance, municipal finance, and household finance: a corporation's investment decision creates a fiscal claim on a municipality, and the disposition of that claim reaches household portfolios. We view the siting of a data center as a two-sided decision, trace how the windfall it creates is divided, and state the hypotheses that follow. Appendix A contains a simple model of this section's discussion.

A The Political Economy of Data Center Siting

A data center is sited where two decision-makers, each accountable to a different principal, find the same location optimal. On one side stands a corporation whose executives, pressed by shareholders for a return on billions of dollars of computing capital, compare scarce land, power, water, and network connectivity across competing counties. On the other side stands a community whose elected officials, mindful of re-election and recall, weigh the fiscal benefits, including tax concessions (Slattery and Zidar, 2020), against the local costs their residents would absorb.[4] Neither chooses alone: the corporation cannot locate without zoning, permits, and entitlements, and the community cannot demand terms from capital it can lose to a competitor (Slattery, 2025). The siting decision is thus a Coasean bargain (Coase, 1960): zoning gives residents an effective property right over entry, and interjurisdictional competition disciplines what their officials can charge for it.

The corporation's outlay is a capital-expenditure shock whose local effects differ sharply across two phases. The construction phase is brief. For eighteen to twenty-four months the project injects demand, but its specialized crews come largely from elsewhere, leaving behind mainly lodging, restaurant, and retail spending. The operational phase is long but employs almost no one. Here the data center breaks from the "million-dollar plants" of the place-based literature (Greenstone and Moretti, 2003), which reach households through job multipliers (Moretti, 2010) and agglomeration spillovers to other firms (Greenstone, Hornbeck and Moretti, 2010).[5] What a data center delivers is fiscal, and largely out of proportion to its footprint: only dozens of employees beside an assessed

[4] Even the incidence of local business taxation is divided, among firm owners, workers, and landowners (Suárez Serrato and Zidar, 2016), so what a concession costs residents is itself a bargaining variable.

[5] Kline and Moretti (2014) show that the Tennessee Valley Authority raised local manufacturing employment through agglomeration gains that persisted for decades after the subsidies ended.

value in the hundreds of millions. Against this stand real but diffuse costs, from grid and water load to noise and ecological strain (Joint Legislative Audit and Review Commission, 2024).[6]

Once operating, the facility pays property tax, and the local government collects an increment of revenue, net of whatever was abated. Because the base arrives without the residents whose schools, highways, and services ordinarily absorb such gains, the increment requires little offsetting expenditure. In corporate terms, this *delta tax cash flow* is high-margin, close to pure surplus for the local budget. Its disposition is where corporate finance meets municipal finance: the free-cash-flow problem of the firm (Jensen, 1986; Blanchard, Lopez-de-Silanes and Shleifer, 1994), transplanted onto a public balance sheet. The government may retain it, though it rarely does, having no capital market to reward accumulated value and every reason to spend. It may be diverted to projects serving the politician more than the community, the public counterpart of empire-building in the firm (Stulz, 1990). It may be spent on public goods. Or it may be returned to the residents themselves. Which use prevails is settled not in a market but through the Coasean bargain of the siting decision. The transfer that secures residents' consent prices the facility's local costs, and Coase's logic implies a bargain is struck whenever the surplus covers the harm, but it does not fix the terms. Those terms, we argue next, depend on how much new base arrives, how long it is expected to last, and how far residents trust the officials who hold it. Each of these uses also presupposes that the increment arrives at all, which abatement can prevent, so we verify this premise empirically before testing how the cash flow is divided.

The hypotheses engage an emerging literature on data centers, which has developed along three dimensions. The first studies local real effects. Alvarez, Argente, Chow and Van Patten (2026) find higher employment and establishment counts in the counties most exposed to the buildout. Yue and Zeng (2026) document similar gains, primarily in metropolitan counties. In both analyses, the employment response is modest relative to the capital installed, and neither studies the fiscal channel, which for a facility this labor-light is the natural conduit to households.[7] A second strand studies who bears the costs. Retail electricity prices rise with entry and the increase falls on incumbent

[6]Industrial openings are often priced as disamenities: Currie, Davis, Greenstone and Walker (2015) find toxic plant openings depress house prices within about a mile of the site. The sign can flip where households share the cash flow, as when shale royalties offset contamination risk (Muehlenbachs, Spiller and Timmins, 2015).

[7]On the local economic footprints of earlier investment booms, see Donaldson and Hornbeck (2016) on railroads and Forman, Goldfarb and Greenstein (2012) on the internet.

ratepayers (Benetton, Compiani and Morse, 2023), concentrated where investor-owned utilities pass costs through (Meeks, Pless, Qi and Wang, 2026). The fiscal gain, by contrast, accrues through the county's tax roll, an asymmetry our framework prices as the cost of consent. A third line studies the terms of attracting a facility: state tax incentives increase data-center construction (Gargano and Giacoletti, 2026), leaving what the county does with the base it receives, and how durable that base will prove, the open questions our risk and governance hypotheses take up. None of these papers observes the county's own levy, base, and tax rate, which is where our paper contributes.

B Hypotheses

Residents acquire a claim on this surplus because hosting a data center is politically costly: the local costs are conspicuous, the benefits are not, and siting is increasingly challenged through lawsuits, hearings, and moratoria. Officials mindful of re-election therefore have reason to pay for residents' consent, and the property tax is the natural instrument. Dispersed spending is hard for voters to observe or attribute, the obstacle behind the flypaper effect (Hines and Thaler, 1995). A rate cut is the opposite: the property tax is the levy households resent most,[8] it recurs and is credited to the incumbent, and the windfall arrives on the very base it reduces.

Hypothesis 1 (Payout). *When a data center arrives, the host community returns part of the resulting fiscal windfall to residents, most visibly as a reduction in the effective property-tax rate.*

A recurring reduction in a household's tax bill is more than an annual saving. It raises, potentially in perpetuity, the after-tax cash flow the home yields, so the price rises by the present value of that stream (Oates, 1969), discounted at a rate that reflects the stream's risk. Delivered as lower taxes, the relief becomes a stock of housing wealth concentrated in the asset that dominates the household balance sheet. It accrues to incumbent owners, not to renters, who hold no claim, nor to later buyers, who pay the capitalized price.

Hypothesis 2 (Capitalization). *The reduction in residents' property-tax burden is capitalized into higher residential property values, raising the housing wealth of incumbent homeowners.*

Whether relief is granted, and how much capitalizes, turns on the revenue's durability and risk.

[8] Cabral and Hoxby (2012) tie the property tax's salience to its lump-sum billing and link that salience to tax revolts, and Finkelstein (2009) shows that less salient taxes support higher tax levels.

First, because a rate cut is nearly irreversible within the electoral horizon, the official commits only against revenue that will persist, so uncertainty, not the expected level alone, deters relief.[9] Second, the home price is the present value of the saving, so a riskier base is discounted at a higher rate and capitalizes less. The two channels weigh risk differently: the market capitalizes only the base's systematic risk, while the official, undiversified in one dominant taxpayer, bears its total risk, so a single volatile taxpayer deters relief more than its systematic risk warrants.

Hypothesis 3 (Risk). *The riskier the incremental tax cash flow, the smaller the payout and its capitalization: the official grants less relief against a stream that may not persist, and the market discounts whatever relief is granted at a higher rate.*

A final condition is the quality of local governance. The payout may be viewed as the price of residents' consent, and that price depends on how much residents trust the government that holds the surplus. Where government is opaque and public integrity is weak, residents rationally expect the delta tax cash flow to be diverted, so a dollar retained by the county is worth less to them than a dollar returned, and consent commands a larger and more visible payout. Where government is transparent, residents can verify how the money is spent and grant consent more cheaply, so officials can retain more of it. The corporate parallel is the substitute model of dividends of La Porta et al. (2000), in which it is the weakly governed firm that must pay out, as shareholders discount cash retained inside the firm (Dittmar and Mahrt-Smith, 2007).

Hypothesis 4 (Governance). *The payout is larger where local government is less transparent and public integrity is weaker: the weaker the governance, the higher the price of residents' consent.*

III Data

The empirical analysis rests on a county-year panel that we assemble from a national census of data centers together with the local public finances and residential real estate of U.S. counties. This section describes the construction of the sample and then the three groups of variables the analysis employs: the data center characteristics, the property tax and home price variables, and the county

[9] Practitioners behave accordingly: Loudoun County, the largest host, routes part of its data-center revenue into a dedicated Revenue Stabilization Fund, hedging the stream's impermanence (Loudoun County, Virginia, 2024).

characteristics. All variables are defined in Appendix B.

A Sample Construction

We identify data centers from the S&P 451 Research Data Center Database, the most comprehensive commercial census of U.S. facilities, compiled by industry analysts through site visits, operator briefings, surveys, and public permitting records, and integrated by a national-laboratory report commissioned by the U.S. Department of Energy into its inventory of the sector. For each facility the database reports its coordinates and ZIP code, its installed capacity in megawatts, its operator and provider type, and the year it became operational. The registry is selective against speculative announcements relative to rival inventories, and we check the data further: because the raw records may still retain phantom entries and never-built projects, we verify the registry facility by facility and drop what did not materialize.[10] The result is a census of 4,505 verified facilities spanning the status spectrum from planned to operating, 2,881 of them operating, counted at the facility level rather than by operator or tenant, each spatially joined to a county.

A county enters the treated group when its data centers are fiscally material, with estimated facility value of at least 1% of the county's property-tax revenue in the 2017 Census of Governments, the last complete fiscal census before the boom accelerated, so the denominator is measured before the windfall it screens. Of the 515 counties with a verified facility, 435 have a usable 2017 denominator and 123 clear the materiality screen. The median such county's data centers are worth 0.10% of its property tax, so the screen removes counties where no fiscal effect is possible. The treated counties are set against a control pool of 2,635 counties verified to have no data center of any size, which excludes the sub-threshold counties a naive never-treated definition would misclassify as controls. Nine control counties missing an entropy-balance covariate are dropped when the reweighting is built, leaving 2,626. The resulting panel is annual and observes nearly 3,000 counties over 15 years.

[10]The caution is warranted in this industry. Inventories that count announced projects have carried more than twice this registry's capacity before waves of cancellations, and a Federal Reserve Board analysis of construction records finds that roughly one third of announced data-center projects is abandoned. Facility counts also differ across commercial sources by unit of observation, since some listings count operators or tenant deployments rather than buildings, and by status screen, since some report every recorded project and others only operating facilities. Our census is verified at the facility level, and the never-built announcements return below as a placebo group.

B Variable Construction

B.1 Data Center Characteristics

The treatment intensity in our designs is the county's cumulative operating capacity, built from dated capacity readings of operational facilities alone. Three further facility measurements matter for the tests. The provider classification separates durable from speculative tax bases, and it is the database's own provider type, recorded facility by facility and fixed before any outcome of ours was examined. Aggregating to the county, a non-crypto data center county is one whose capacity is entirely hyperscale, colocation, or enterprise, a crypto county is one whose capacity is cryptocurrency mining, and the small number of counties with capacity of both kinds is identified separately, so the risk comparisons contrast pure types. For facilities of at least 50 megawatts we hand-collect the year of the project's first public announcement from contemporaneous public records, so each date is verifiable, and the lag from announcement to operation averages 2.1 years. This announcement clock dates the market's first opportunity to price entry. Finally, we hand-collect publicly announced projects that were never built, and counties with such a project and less than 10 megawatts of other capacity form the placebo group.

B.2 Property Tax and Home Price Variables

The effective property-tax rate, the outcome of the payout hypothesis, is median real-estate taxes paid (ACS table B25103) divided by median home value (B25077), computed annually at the county level, and it reflects what homeowners actually pay regardless of how a state labels its levies.[11] We complement it with the millage rate, the nominal rate per dollar of assessed value that the county's governing body adopts each year with its budget, so each observation is a public legal document rather than a survey response. Millage histories are not centrally compiled, so we hand-collect them from county budget ordinances, building a panel of 1,054 counties across twenty-four states from 1996 through 2026, including 57 counties whose data center arrives inside the observed series. Because statutes shape how a rate can move, we also classify every state's regime from its statutes, anchoring

[11] Because the numerator and denominator come from the same household survey, the measure is constructed identically in every county and is comparable across states whose assessment ratios, exemptions, and levy structures differ. A statutory rate is not: two counties with the same posted rate can bill very different shares of market value.

each classification to the statute's text: states where rates change only by an act of the county's governing body, states that cap growth in assessed values, and states whose statutes recalibrate rates upon revaluation.

Home values are measured four ways, with the ACS one- and five-year medians and with the Zillow Home Value Index at the county and ZIP-code levels, so that no single index is decisive. The measures rest on independent primary sources, transaction records for the Zillow index and household survey responses for the ACS medians, so agreement across them is informative rather than mechanical. Median gross rent (B25064) prices the rental margin of the capitalization hypothesis with the same survey instrument as the home values, so the owner-renter contrast is not an artifact of differing sources. For the within-county tests, ZIP-code centroids are linked to counties through a ZIP-code-tabulation-area crosswalk, which yields each ZIP code's distance to the nearest facility, and the Census county-adjacency file identifies the contiguous neighbors of the border-pair design.

B.3 County Characteristics

County finances come from the Census of Governments and from audited Annual Comprehensive Financial Reports, supplemented by the MuniSpot academic license for general- and all-fund line items, which together yield revenue, expenditure, debt, and fund balances. A county shares its property-tax base with the school districts, municipalities, and special districts layered atop it, so we assemble the property-tax revenue of each layer and aggregate it to the county, which prevents attributing to the county government a windfall that in fact accrues to the districts, and we add the state intergovernmental revenue the county receives for the state-aid test. The Census property-tax series carries scope inconsistencies across states, and where they arise we cross-check the data center counties' figures against their audited statements and correct them. Property-tax abatements are observed through Governmental Accounting Standards Board Statement 77 disclosures.

County governance is measured by the form of government, hand-collected to distinguish counties with a professional appointed manager from commission counties, and by whether the county holds the Government Finance Officers Association reporting certificate. Housing-supply elasticity is measured four ways: residential building permits relative to the housing stock from the Census Building Permits Survey, the metropolitan supply elasticity of Saiz (2010), the Wharton Residential

Land Use Regulatory Index (Gyourko, Hartley and Krimmel, 2021), and the tract-level elasticities of Baum-Snow and Han (2024). The remaining characteristics of Table 1, population, income, education, homeownership, and the pre-period growth trajectories that serve as balancing covariates, come from the ACS, and the count of power plants comes from EIA Form 860.

C Summary Statistics

Table 1 reports summary statistics for the 340 counties with at least one data center and compares those to the 2,626 counties with no data center at all. The two groups differ on several economic characteristics. Five findings are particularly relevant for our empirical analysis.

First, data centers locate close to population centers rather than away from them. Put differently, and contrary to a common perception, data centers do not generally locate in rural communities. The typical data center county is populous, high-income, and well educated, with each of these differences from other counties significant at the 1% level. Such counties also hold more of the input infrastructure that data center operations require, from power plants to the long-haul fiber backbone that carries internet traffic.

Second, the data centers themselves are economically consequential. The average data center county has 8.0 data centers with a combined size of 122 megawatts, and the facilities' attributable property tax averages 7.7% of the county's property-tax base. The treatment we study is therefore meaningful enough to move the county finances at the center of our hypotheses.

Third, counties with data centers have higher median home values (consistent with being wealthier counties) and therefore collect more property-tax revenue, so a windfall arrives where a meaningful base and a tax collection apparatus already exist, which is what gives a community something to hand back to its homeowners. The effective property-tax rate is slightly above 1% in both groups, modestly higher in data center counties. It is worth noting that county fixed effects absorb this level difference, and the payout hypothesis concerns how the rate changes at entry.

Fourth, the homeownership rate in data center counties is 67%, lower than the 73% elsewhere but still two of every three households. A payout delivered through the property tax therefore reaches the electoral majority of the community that grants consent. The remaining third of households rent, and what the windfall does for renters is a question our data can answer directly, which we

take up when we examine rents alongside home prices.

Fifth, the pre-period growth trajectories of the two groups run against a selection-on-momentum story. Income growth is identical (5.35 versus 5.31%, $p = 0.910$), and home-value growth is slower in the data center counties (3.81 versus 5.99%, $p < 0.001$). Data center siting follows people, wealth, and resources, not counties whose incomes and home values are already accelerating.

IV Empirical Methodology

This section presents the model specification, the balancing of the control group, and the inference and interference challenges of our setting together with how we resolve them.

A Model Specification

Our design is a staggered difference-in-differences (DiD) comparison of counties where a data center enters with counties where none does:

$$y_{c,t} = \beta \, \text{Data Center}_{c,t} + \mu_c + \delta_{s,t} + \varepsilon_{c,t}, \tag{1}$$

where $y_{c,t}$ is the outcome for county c in year t and μ_c and $\delta_{s,t}$ are county and state-by-year fixed effects. The county fixed effects absorb permanent differences across counties, in size, wealth, and amenities, and the state-by-year fixed effects absorb statewide shocks such as housing cycles and changes in state tax policy, so β is identified from changes within a county relative to same-state counties in the same year. The treatment variable $\text{Data Center}_{c,t}$ takes two forms: an indicator turning on when the county's first facility enters, or the continuous *Data Center Size*, the county's cumulative installed capacity in gigawatts. Both are zero before entry and in counties without a facility, so the familiar interaction of treatment and post-period indicators is embedded in the variable itself. The control pool consists of counties with no facility in the database at any time.

Our target is the average effect of treatment on the treated. Under staggered adoption, least squares applied to Equation (1) averages many two-group, two-period comparisons, some using already-treated counties as controls for later cohorts (Goodman-Bacon, 2021), and when effects vary across cohorts or over time some cohort-by-period effects can receive negative weight, with analogous

problems for the continuous form. We therefore re-estimate the main specifications with the group-time estimator of Callaway and Sant'Anna (2021), which draws controls only from not-yet-treated and never-treated counties and aggregates cohort-by-period effects with non-negative weights, and with the estimators of de Chaisemartin and D'Haultfœuille (2020) and Sun and Abraham (2021). Agreement across estimators shows this weighting does not drive our estimates of β.

B Balancing the Control Group

Data centers are placed where power, fiber, land, and permitting are favorable, so a raw comparison risks crediting the facility with advantages that predate it. These same fundamentals are features of a county's industrial character itself, so no credible exclusion restriction is available for an instrumental variable, and we rely on research design rather than instruments. Unweighted, Equation (1) requires parallel trends within state-by-year cells, and reweighting the controls toward the treated distribution weakens this to a conditional version given the balanced characteristics. Entropy balancing is our workhorse because, when weights exist, it matches the specified control moments to the treated moments exactly (Hainmueller, 2012). Propensity-score matching and doubly robust estimation serve as alternatives, the latter consistent if either its outcome or its treatment model is correct.

Our primary specification is a single entropy-balancing weight, held fixed across all outcomes. It matches the controls' pre-treatment income and home-value growth paths to those of the treated, targeting selection on trajectories directly. A more demanding balance additionally equalizes pre-treatment levels, the siting covariates, and census division at a smaller effective sample, which we report so the balance-precision trade-off is explicit.

C Inference and Interference

This subsection confronts four remaining issues: the role of chance, spillovers across county lines, selection on unobserved characteristics, and the computation of standard errors.

C.1 Randomization Inference

Treatment is concentrated in a relatively modest number of counties, particularly at the start of the build-out period, so a few lucky draws could produce estimates of our size, where asymptotic

approximations give limited guidance. We therefore compare the estimate of β with its distribution under re-randomized treatment, reassigning entire capacity histories at random across counties with timing preserved and re-estimating Equation (1) on 1,000 such draws. Up to simulation error, the resulting p-value is exact for the sharp null that treatment affected no county in any year, under the assumption that any county was equally likely to receive any capacity history. Actual siting is deliberate, so we use the test as a demanding benchmark rather than as literal design-based inference: it tells us how often estimates of our size arise when capacity histories are scattered across counties without regard to fundamentals, free of any distributional assumption on $\varepsilon_{c,t}$.

C.2 Border-Pair Design

Data centers cluster where their inputs are abundant, so a data center county's neighbors are rarely bystanders: they share its labor and housing markets, and its effects need not stop at the county line. A data center county can then move its neighbors' outcomes, violating the assumption implicit in Equation (1) that county c's potential outcomes depend only on $\text{Data Center}_{c,t}$. We purge the control group of each county with a facility and gauge residual spillovers by also dropping contiguous counties. Surviving spillovers attenuate the estimate only if they share the sign of the direct effect, since benefits leaking to controls narrow the treated-control gap while displacement would widen it, so the contiguity exclusion guards against both directions. The border-pair design trades a measured dose of spillover for comparability, pairing each data center county with contiguous neighbors across a state line, where the housing market is shared but the tax treatment differs.

C.3 Never-Built Placebo

Reweighting matches only observable trajectories, and counties that attract a data center may differ in characteristics no covariate captures. The concern has a classic answer in the plant-location literature, which compares the counties that win a large plant not with average counties but with the near-winners that almost received it (Greenstone and Moretti, 2003; Greenstone, Hornbeck and Moretti, 2010). Our analog is the set of counties where a data center was announced but never built. These counties cleared the same siting screens, attracted the same developer interest, and display the same rising pre-period paths, but no new tax base ever arrived. If our estimates reflected selection

on prospects rather than the arrival of taxable capital, home prices in these counties should break at the pseudo-event the way prices in data center counties break at entry. Instead they continue along their prior path, which is the pattern selection alone should produce.

C.4 Standard Errors

Baseline inference clusters standard errors by county, the level at which treatment is assigned and the property tax is set, allowing $\varepsilon_{c,t}$ to correlate serially within a county. Some designs leave few treated clusters, where cluster-robust estimators over-reject even with many clusters overall (MacKinnon and Webb, 2017), and we rely on the wild cluster bootstrap (Cameron, Gelbach and Miller, 2008). Because $\varepsilon_{c,t}$ may correlate across nearby counties, we confirm the results survive state-level clustering and spatial standard errors that let dependence decay with distance (Conley, 1999).

V Empirical Results

A financial windfall that ends up with a community's homeowners must first arrive on the county's balance sheet, survive the state's aid formulas and clawbacks, pass through as a lower effective tax rate, and capitalize into property values at a price that reflects its risk. In this section, we test each mechanism in this chain.

A Effect on County's Property-Tax Base

We start by analyzing whether a county that admits a data center into the community in fact collects more property tax. Each data center county is compared with its matched controls, and we measure the growth of the property-tax base using the Censuses of Governments. Table 2 shows that hosting a data center delivers a windfall of incremental tax revenue. The fiscal gains documented in case studies of prominent data center counties, such as the counties of Northern Virginia (Prince William County, Virginia, 2022), are not artifacts of a few cases but a general feature of data-center entry.

The community's property-tax increase is statistically significant and economically meaningful. Panel A of Table 2 shows that the consolidated local property-tax base grows about 1.1 percentage points per year more in data center counties than in their matched controls, a difference significant at

the 1% level, which compounds to a base roughly 5% larger over the five-year window. Decomposing the gain by layer of government, the county's own base grows 1.2 percentage points faster per year ($p = 0.003$), and the city/township layer grows 0.7 points faster (Wilcoxon $p = 0.006$).[12]

Because the unscreened sample mixes counties where no fiscal effect is arithmetically possible with those where one is, we also constrain the sample to the counties where data centers are fiscally meaningful, with attributable tax of at least 5% of property-tax revenue. The screen approximately triples the county effect, from 1.2 to 3.8 percentage points per year ($p = 0.020$) across the matched pairs that survive the screen, and this scaling of the effect with the facility's fiscal weight is itself further evidence that the data center, and not some correlate of siting, drives the gain. The median county in this restricted sample has a property-tax base of $14.1 million and collects about $0.5 million a year of delta tax cash flow. In homes, the magnitude is starker: the median data center county's facilities pay property tax equivalent to roughly 3,400 average homes, 14% of what its entire owner-occupied stock contributes, measured gross of abatements. The comparison understates the fiscal advantage, because residential development costs local governments $1.10 to $1.30 for each dollar of tax revenue it brings, most of it schooling, while a data center adds almost no service load.

The windfall's anatomy is visible in the counties' own audited financial reports. The disclosure exists for a subsample of data center counties. Within this coverage, a data center now ranks among the ten largest taxpayers in 21 of 40 counties and is the single largest in nine, and none appeared in its county's disclosure a decade ago. For example, in Loudoun County, all ten current largest taxpayers are data-center operators, against two a decade ago. Yet the typical covered county is no more concentrated after entry than before: the largest taxpayer's median share of the base edges from 2.9% to 2.8% between the two disclosure dates (sign test $p = 0.143$), and a Herfindahl index across the ten largest taxpayers is likewise flat. The base grows without turning the typical county into a one-taxpayer county, and single-taxpayer dependence arises only in the most exposed tail.

Panel B analyzes whether the state claws back part of the county's gain. If state-aid formulas offset local tax revenue growth, the increase would be redistributed away from the data center community rather than retained by it. In contrast, we find that state aid to the county shows no change. The

[12]Many states cap the growth of a county's property-tax levy, but thirty-six states carve newly constructed property out of the cap, adding a new facility to the county's levy capacity at full rates, and 89% of the counties with data centers lie in these states. Levy limits therefore place no constraint on the tax a new facility generates.

point estimate, 0.53 percentage points per year ($p = 0.533$), is statistically indistinguishable from zero, and the distribution-free Wilcoxon test agrees ($p = 0.137$). The estimated delta tax cash flow to the data center county therefore augments the community's net resources rather than substituting for aid the county would otherwise have received.

The reported estimates are lower bounds on what a fully taxed data center delivers, for three reasons. First, the headline comparison imposes no size screen, so it averages in many counties whose facilities are too small to move a budget, and a precisely estimated positive differential on such a sample is more difficult to dismiss than one built from large facilities alone. Second, the census thins out below a megawatt, so small enterprise installations it does not record sit in the control counties and attenuate the contrast. Third, abatements are a standard instrument in data-center location negotiations, and where a community grants one, or defers or rebates the facility's taxes, the tax cash flows never reach the local balance sheet or arrive only with delay (Joint Legislative Audit and Review Commission, 2024), yet such counties remain in the treated group. Each feature biases the measured increase toward zero.

B Effect on Homeowners' Property-Tax Rates

The payout hypothesis predicts that part of the tax gain from a data center is distributed to residents as a lower property-tax rate. Table 3 reports the evidence, difference-in-differences estimates across all data center counties, with the effective rate as the dependent variable and *Data Center Size*, the county's cumulative facility power in gigawatts, as the treatment variable. In the baseline specification the rate falls by 0.117 percentage points per gigawatt of data center size ($p < 0.001$). For a representative new campus of 500 megawatts our estimate implies a cut of approximately 0.06 percentage points, or about 7% of the median county's effective tax rate of 0.82 percentage points. That is a substantial concession: absent a major expansion of the tax base, few county governments can sustain a permanent cut of such a significant size in the property-tax rate.

The estimated rate cut cannot easily be attributed to the construction of the outcome measure. First, a county that merely grows sees tax bills and home values rise in tandem, leaving the effective rate unchanged, so the measure is insulated by construction from the price trends that complicate housing analyses. Reassessment lags may create a wedge, but it is bounded by the price response,

transitory, and too small to produce the cut we estimate. Second, the estimate does not depend on the survey: across the samplings of the American Community Survey it stays between -0.111 and -0.092 percentage points per gigawatt, negative and significant in each cell.

Complementary evidence that the cut is enacted policy comes from the millage rate, the nominal rate of tax per dollar of assessed value that the county's governing body adopts each year with its budget, and the rate the public sees. Millage histories are not centrally compiled, so we hand-collect them from county budget ordinances, building a panel of 1,054 counties across twenty-four states from 1996 through 2026. Across the 57 counties whose data center arrives inside their observed rate series, the adopted millage falls by 3.4% after entry (randomization $p = 0.068$). The adopted and effective rates also agree where both are measured: on the identical 34-county switcher sample, the adopted rate falls 3.0% and the effective rate 1.9%, and at the sample's mean capacity of 0.4 gigawatts the per-gigawatt estimate of Table 3 implies a 4.8% decline. Three routes to the rate run in the same direction at comparable magnitudes: the ordinances officials vote, the bills homeowners report, and the capacity design.

A remaining question is whether any statute, rather than any official, sets the rate cut. Table 4 re-estimates the rate response within each statutory regime. In the discretionary states, the county's posted rate changes only by an act of its governing body, and the cut is largest there, 0.143 percentage points per gigawatt. In the assessment-cap states, statutes constrain assessed values rather than rates, leaving the rate itself to the county, and the cut persists. The statutory-adjustment states recalibrate rates when existing property is revalued, and these statutes exclude newly constructed property from the calculation, so a data center, which enters the tax roll as new construction at full rates, lies outside the adjustment: the cut estimated there is likewise one the county adopted. In no regime is the rate reduction that follows a data center automatic or required by statute.

The cuts could still reflect the counties data centers select rather than the facilities themselves. Entropy balancing addresses this by reweighting the control group to the treated counties' pre-entry growth trajectories. The results hold: under this balancing, our primary specification, the tax rate is cut by 0.092 percentage points per gigawatt ($p < 0.001$), and under the more demanding full-covariate match by 0.051 ($p = 0.040$), negative and significant throughout. Because data center capacity is geographically concentrated, we complement clustered standard errors with randomization

inference, reassigning entire capacity histories across counties and re-estimating the specification on each reassignment. An effect as large as ours arises by chance in fewer than five percent of the reassignments in either specification ($p = 0.003$ baseline, $p = 0.047$ entropy), so the estimates hold under conventional and under the most conservative inference alike. The design also passes the pre-trend check: in an event-study version, the pre-entry coefficients are individually and jointly indistinguishable from zero ($\chi^2(3) = 5.83$, $p = 0.120$), so effective rates were not drifting down before data center capacity arrived.

An alternative empirical analysis is the border-pair design: each data center county is compared with its contiguous neighbors, which share its housing and labor markets, so local conditions are held fixed while data-center exposure differs. As the empirical methodology notes, the design accepts a measured dose of spillover in exchange for this comparability, and because any spillover to the neighbor biases the comparison toward zero, the estimates it produces are conservative. Table 5 reports the design under three definitions of a pair, and the property-tax cut appears in each: 0.093 percentage points per gigawatt of data center size in the cross-state pairs, the cleanest benchmark, 0.056 across all contiguous pairs, and 0.064 in pairs within 50 miles, definitions that add power. All three estimates also survive inference by wild-cluster bootstrap, with $p = 0.006$ in the cross-state pairs. The tax cut is estimated as larger than in Table 3 because the national estimates average over all data center counties, while the border design measures each against its most comparable neighbor.

C Effect on Residential Property Values

The capitalization hypothesis predicts that the payout, a recurring reduction in the household's tax bill, is capitalized into the price of the homes that receive it: a buyer acquires the home together with its future tax bills, so a recurring cut raises the price by the present value of the stream. Because housing is forward looking, the treatment is an entry: the market prices the expected payouts when a data center arrives. Table 6 reports difference-in-differences estimates of the average effect of a data center's entry on residential property values, employing four measures. The estimate is positive and statistically and economically significant for each home-price index, from 2.2% for the overlapping

five-year ACS median to 3.8% for the Zillow index.[13]

The baseline estimate is not sensitive to alternative model specifications. Table 7 re-estimates the effect with the Zillow index (the index available annually for the largest number of counties), and each column confronts a distinct concern. First, states that attract data centers may be experiencing statewide housing booms, which the baseline's state-by-year fixed effects absorb. Column (1) removes them in favor of year fixed effects. Second, home prices and current tax bills are in practice jointly determined, so column (2) adds the county's effective property-tax rate as a control. Third, the housing index's county coverage expands over time, so column (3) restricts to a balanced panel. Every estimate remains close to the 3.8% Zillow baseline of Table 6, and each is significant at the 5% level or better.[14]

The capitalization hypothesis also carries a geographic prediction. The rate cut applies to every home on the county's tax roll, so the price response should be county-wide rather than concentrated near the facility, the gradient a localized story would produce. Table 8 tests this prediction within counties, placing each ZIP code of the data center counties in a ring by its distance to the nearest data center. Within twenty miles of a facility, the distance a county typically spans, the response is statistically uniform, with ring estimates from 3.0% to 4.7% that are indistinguishable from one another. The gain reaches the whole county, as relief delivered through the tax rate should, and the common concern that a data center depresses the value of nearby homes finds no support.

Housing supply elasticity distinguishes two explanations for why prices rise. If entry raises prices because the facility raises the demand for living in the county, from jobs, spending, or migration toward the lower tax bill, then the price response should depend on how easily the county builds: where supply is elastic, new construction absorbs the demand and prices barely move, and where supply is inelastic, the same demand becomes price (Hilber, 2017). A capitalized tax cut obeys no such geography. The cut attaches to the parcel, lowering the tax on every existing home and every future one, so its present value belongs to the land, whose supply is fixed even where building is easy. The demand story therefore predicts a price response concentrated in inelastic counties, and

[13]The continuous data center size measure confirms the capitalization. The Zillow index rises 8.26 log points per gigawatt of data center size ($p = 0.044$), or about 4.1% for a 500-megawatt campus.

[14]The estimate also holds under estimators designed for staggered adoption: the group-time estimator of Callaway and Sant'Anna (2021), which draws controls only from not-yet-treated and never-treated counties, yields 5.1% ($p = 0.005$), and the synthetic difference-in-differences estimator of Arkhangelsky, Athey, Hirshberg, Imbens and Wager (2021), which chooses its own control-unit and time weights, yields 6.2% ($p = 0.006$).

the capitalization story predicts the same response everywhere. Table 9 runs the race, interacting entry with four elasticity measures: building permit intensity, the geographic supply elasticity of Saiz (2010), the Wharton Residential Land Use Regulatory Index (WRLURI) of Gyourko, Hartley and Krimmel (2021), and the tract-level supply elasticities of Baum-Snow and Han (2024). The interactions are uniformly insignificant: the price response is no larger where building is hard. Two of the measures, those of Saiz and of Baum-Snow and Han, capture the same construct yet enter with opposite signs, so the flatness does not reflect a weak proxy. Prices rise alike where construction is easy and where it is difficult, which is the mark of a tax stream being capitalized, not of demand pressing against supply.

D Effect of Tax-Base Risk

The risk hypothesis holds that the property-tax gain is capitalized only to the extent that it is expected to persist. Data centers may differ on this margin. A hyperscaler or colocation operator is commonly an investment-grade taxpayer with a multi-decade horizon. In contrast, a cryptocurrency-mining operator is speculative and exit-prone, as the recent failures of several large miners confirm. If the market prices this difference, the capitalization effect should concentrate in the counties whose data center property-tax base is expected to be more persistent. Table 10 reports difference-in-differences estimates for non-crypto data center counties only. The Baseline column is unweighted, and its estimate is approximately twice the pooled estimate of Table 6: restricting the sample to more durable tax bases alone doubles the measured home-price response, and both weighted estimates remain significant at the 1% level. The market pays for persistence: the response is strongest exactly where the tax base behind it is expected to endure.

A corollary of the risk hypothesis concerns when the home-price response arrives in the market. A data center is announced years before it operates, and a forward-looking market may capitalize part of the expected tax cut at announcement, with the remainder arriving once operations begin and the completion risk is resolved. Table 11 re-anchors the event clock at the county's first announcement. Dating treatment from announcement raises the estimate to 4.9% ($p = 0.002$), and the response decomposes as 3.7% arriving at announcement ($p = 0.019$) and 1.7% once operations begin ($p = 0.147$). About two-thirds of the value is priced when the project becomes public. The

announcement clock inherits a rising pre-period of its own (joint $p = 0.025$), so these estimates date the arrival of the response rather than eliminate the run-up that precedes it.

The average size of a data center is a potential confounder of the crypto effect. Crypto facilities are generally smaller, with a mean of 134 megawatts against 367 for their non-crypto counterparts, so their smaller home-price response may simply reflect a smaller tax base rather than risk. Table 12 separates the two by interacting entry with data center size and with the facility's share of the tax base. The gap between the two kinds of counties is essentially unchanged by either control: home prices rise between 6.9 and 9.1 percentage points more in non-crypto counties than in crypto counties across specifications, significant at the 1% level throughout. The premium attaches to the persistence of the tax base and not to its size. The cryptocurrency-mining bankruptcies of 2022 provide an out-of-sample check. From the county's perspective the collapse was an exogenous shock to the durability of its tax base, yet crypto-county home values show no reversal: little had been capitalized, so there was little to reverse, and most failed facilities converted to other computing uses rather than closing.

Counties where a data center was announced but never built constitute an informative placebo analysis. These counties share the siting fundamentals of the data center counties, so the market can price the expected tax cut at announcement, and when the project is never completed, the price reaction should decay with the probability of completion. Figure 1 shows exactly this pattern: home prices climb steeply into the pseudo-event, show no break when it arrives, and drift back in the years that follow, while the counties whose facilities were built break at entry and keep rising. The never-built counties are a placebo rather than a control group. They are larger, metropolitan counties on visibly steeper pre-event price paths, and that steepness shows the contrast: if shared siting fundamentals rather than the facility drove the gains, prices in the never-built counties should have kept rising, and instead the anticipation unwinds.

E Effects of County Governance

The governance hypothesis predicts a larger payout where local government is less transparent and public integrity is weaker, because lack of trust raises the price of residents' consent. Panel A of Table 13 re-estimates the rate response in the two subsamples with weaker governance: commission

counties, which operate without a professional appointed manager, and counties that do not hold the Government Finance Officers Association (GFOA) certificate for financial reporting. The estimates are 0.085 and 0.105 percentage points per gigawatt, comparable to the full-sample estimate of Table 3, with the larger cut in the counties without certified reporting. Both subsamples are small, so neither estimate is precise ($p = 0.453$ and $p = 0.282$), but the pattern is consistent with the hypothesis.

Panel B tests the mechanism more directly. If residents discount retained money, a dollar paid out should raise home values more than a dollar retained. Each treated county contributes one observation, the change in log home value from the year before entry to three years after, regressed on the annual tax saving per home implied by the county's adopted rate cut and on the annualized change in the county government's cash and security holdings per capita. Column (1) prices the payout alone, at 0.026 log points per dollar of annual saving. Column (2) adds the retained flow, and the coefficients separate, 0.034 against 0.009, an ordering unchanged in leave-one-out re-estimates. Column (3) excludes the two Nevada counties, whose posted-rate concept and assessment ratio make their dollar conversions unreliable, and interacts the retained flow with thin local news coverage. The interaction is negative (-0.015, $p = 0.040$): retention is discounted more where residents have fewer means to verify its use. The sample is 36 counties and the estimates are not precise, but they line up with the substitute-dividends logic: the market values a dollar retained by the government below a dollar paid out.

F Effects on Other Claimants

The more of the windfall that is paid out to homeowners, the less remains for the county's other claimants. Table 14 follows the accounting through the three most natural ones, renters, bondholders, and schools. Renters come first because the capitalization hypothesis itself makes a prediction about them: a generic local demand shock would raise home values and rents together, while a property-tax cut that attaches to ownership raises values alone. The data show the latter. The median gross rent is flat after entry (-0.5%, $p = 0.536$),[15] so the gain lands on the voters who hold the asset the policy

[15]The rent series averages over a trailing five-year window of the American Community Survey, which can only mute a true rent response toward zero, so the comparison with home prices, measured contemporaneously, is conservative.

revalues, while renters bear no offsetting burden.[16]

Bondholders may gain, as local fiscal conditions are priced in the municipal market (Painter, 2020), and an enlarged tax base could compress the spread on the county's bonds. We observe this market through the offering spreads of newly issued municipal bonds from SDC Public Finance, spanning 2008 through 2025. In a sample of 71,225 new issues, the estimated change in the offering spread is 5.0 basis points ($p = 0.113$), statistically indistinguishable from zero, and its positive sign indicates that spreads, if anything, widen rather than compress.

Schools hold the largest statutory claim on the property tax, so a distributed surplus should surface in their inputs or outcomes. We measure both margins, achievement with the Stanford Education Data Archive and staffing and salaries with the NCES Common Core of Data, and neither responds to entry. Test scores are unchanged three years after entry, with an estimate of +0.03 standard deviations ($p = 0.727$) across 36,771 county-years, and the inputs are similarly unaffected: teacher staffing per thousand students does not increase (-2.9%, $p = 0.15$), and teacher salary per student does not change (0.6%, $p = 0.60$). Among the claimants on the county's budget, renters, bondholders, and schools included, homeowners are the ones the windfall reaches.

VI Conclusion

This paper is the first to study the intersection of corporate, municipal, and household finance and trace the chain end to end, from the siting decision, through the county's finances, to the housing wealth of its residents. The wealth effects are economically significant and compound from the individual household to the entire country. For the median home, entry raises its value by about 4.5%, or $12,000, the market's pricing of the property-tax reduction net of any burdens residents expect the facility to bring. Summed over the county's owner-occupied homes, the gain amounts to almost $500 million for the median data center county. Aggregated across the United States, it reaches over $100 billion, about $2.5 billion for every gigawatt of operating capacity. This is the housing wealth created so far, and with the capacity now under construction approaching what

[16]Turnout rises most reliably with residential stability and education (Wolfinger and Rosenstone, 1980), and homeowners are rooted and better educated on average. Ownership also aligns votes with home values (DiPasquale and Glaeser, 1999; Fischel, 2001).

already operates, much of it may still lie ahead. These estimates are confined to the direct fiscal channel: a facility may raise home values indirectly as well, by lifting local growth and housing demand over the longer term, so the total effect on household property values may exceed what the tax channel alone delivers.

The findings carry policy implications for both communities and corporations. For the officials who set the terms of siting, the question is why residents' resistance to data centers remains strong when the pass-through is economically substantial. The resistance need not be irrational, because the payout is unevenly shared. It accrues to homeowners, the larger the home the larger the relief, and homeowners have more education, with human capital potentially more resilient to artificial intelligence's impact. Renters appear to gain little from the AI investment boom, and their human capital may in fact be more exposed to what the data center produces. Those with the least to gain from the facility may have the most to fear from its products, and theirs may be the loudest voices in local debates. The evidence implies that broadening the payout beyond the tax roll, through rebates or relief on rented housing, could win the consent that a property-tax rate cut alone does not. For the corporate executives on the other side of the bargain, the implication is that a durable, unabated tax presence lowers the price of consent. The full levy is a larger expense than an abated one, but the difference funds the transfer that secures residents' approval, and with it the probability that the investment proceeds at all.

Several gaps remain for future research, and two stand out to us. One is financing and corporate leverage. The build-out is increasingly funded with debt, and the markets that price it, from credit ratings to credit default swaps, hold forward-looking assessments of the firms whose taxes constitute the delta tax cash flow. Linking those assessments to capitalization would tie household finance to corporate credit, a natural next step. The other is the identity of the taxpayer. We analyze facilities by category, but persistence may also vary with the corporation behind the campus. A legacy corporation such as Google or Amazon carries decades of operating history, the frontier labs of artificial intelligence may or may not exist in a decade, and whether the market capitalizes the two differently is an extended version of our risk hypothesis.

With an investment boom come questions of a possible bubble and subsequent bust. Dot-com fiber traded at a fraction of its construction cost within a decade, and if AI disappoints, data-center

overcapacity may follow. County exposure is concrete: tax collections fall with assessed values, the base concentrates in a few facilities, and tax-rate cuts are politically challenging to reverse. Yet a bust does not erase the physical infrastructure, which is built to sustain decades of use. Bankruptcy law may resolve a failed operator in an orderly Chapter 11 reorganization, passing the facility to new owners or new uses, as converted cryptocurrency mines recently demonstrate. Redeployment value depends on asset specificity, and because capacity is presumably in its highest-value use today, a forced conversion is likely to realize less value. Distress can also reopen the tax bargain itself, as new owners negotiate abatements or rebates that shrink the tax base.

Our findings contribute to two literatures in financial economics. The first is the intersection of corporate and municipal finance. The classic corporate evidence on cash windfalls is that firms retain them rather than pay them out. The counties we study pay out, and the difference has a specific source: the terms are set before the cash arrives, as the price of residents' consent, and the residents vote, so consent does the work that governance does in firms. The second literature is household finance. Corporate investment reaches household wealth here through the property tax on the home. Capital expenditure at this scale rewrites local tax bases years before its products reach any market, and the housing market prices the rewrite in real time. The paper is among the first to show how technology risk is transmitted into household portfolios: the value of the payout rests on the durability of the taxpayer, so the household's largest asset carries an exposure it cannot easily hedge, one likely to matter more as technology's share of the economy grows even further.

References

Adelino, M., Cunha, I., Ferreira, M.A., 2017. The economic effects of public financing: Evidence from municipal bond ratings recalibration. Review of Financial Studies 30, 3223–3268.

Alvarez, F.E., Argente, D., Chow, J., Van Patten, D., 2026. Data centers and local economies in the age of AI: A shift-share approach. Working Paper 35194. National Bureau of Economic Research.

Arkhangelsky, D., Athey, S., Hirshberg, D.A., Imbens, G.W., Wager, S., 2021. Synthetic difference-in-differences. American Economic Review 111, 4088–4118.

Baum-Snow, N., Han, L., 2024. The microgeography of housing supply. Journal of Political Economy 132, 1897–1946.

Benetton, M., Compiani, G., Morse, A., 2023. When cryptomining comes to town: High electricity-use spillovers to the local economy. Working Paper 31312. National Bureau of Economic Research.

Bernstein, A., Gustafson, M.T., Lewis, R., 2019. Disaster on the horizon: The price effect of sea level rise. Journal of Financial Economics 134, 253–272.

Blanchard, O.J., Lopez-de-Silanes, F., Shleifer, A., 1994. What do firms do with cash windfalls? Journal of Financial Economics 36, 337–360.

Butler, A.W., Fauver, L., Mortal, S., 2009. Corruption, political connections, and municipal finance. Review of Financial Studies 22, 2873–2905.

Cabral, M., Hoxby, C., 2012. The hated property tax: Salience, tax rates, and tax revolts. Working Paper 18514. National Bureau of Economic Research.

Callaway, B., Sant'Anna, P.H.C., 2021. Difference-in-differences with multiple time periods. Journal of Econometrics 225, 200–230.

Cameron, A.C., Gelbach, J.B., Miller, D.L., 2008. Bootstrap-based improvements for inference with clustered errors. Review of Economics and Statistics 90, 414–427.

Campbell, J.Y., 2006. Household finance. Journal of Finance 61, 1553–1604.

Cellini, S.R., Ferreira, F., Rothstein, J., 2010. The value of school facility investments: Evidence from a dynamic regression discontinuity design. Quarterly Journal of Economics 125, 215–261.

Coase, R.H., 1960. The problem of social cost. Journal of Law and Economics 3, 1–44.

Cocco, J.F., 2005. Portfolio choice in the presence of housing. Review of Financial Studies 18, 535–567.

Conley, T.G., 1999. GMM estimation with cross sectional dependence. Journal of Econometrics 92, 1–45.

Cooper, M.J., Gulen, H., Schill, M.J., 2008. Asset growth and the cross-section of stock returns. Journal of Finance 63, 1609–1651.

Currie, J., Davis, L., Greenstone, M., Walker, R., 2015. Environmental health risks and housing values: Evidence from 1,600 toxic plant openings and closings. American Economic Review 105, 678–709.

de Chaisemartin, C., D'Haultfœuille, X., 2020. Two-way fixed effects estimators with heterogeneous treatment effects. American Economic Review 110, 2964–2996.

DiPasquale, D., Glaeser, E.L., 1999. Incentives and social capital: Are homeowners better citizens? Journal of Urban Economics 45, 354–384.

Dittmar, A., Mahrt-Smith, J., 2007. Corporate governance and the value of cash holdings. Journal of Financial Economics 83, 599–634.

Donaldson, D., Hornbeck, R., 2016. Railroads and American economic growth: A "market access" approach. Quarterly Journal of Economics 131, 799–858.

Dougal, C., Parsons, C.A., Titman, S., 2015. Urban vibrancy and corporate growth. Journal of Finance 70, 163–210.

Feyrer, J., Mansur, E.T., Sacerdote, B., 2017. Geographic dispersion of economic shocks: Evidence from the fracking revolution. American Economic Review 107, 1313–1334.

Finkelstein, A., 2009. E-ZTax: Tax salience and tax rates. Quarterly Journal of Economics 124, 969–1010.

Fischel, W.A., 2001. The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies. Harvard University Press, Cambridge, MA.

Forman, C., Goldfarb, A., Greenstein, S., 2012. The Internet and local wages: A puzzle. American Economic Review 102, 556–575.

Gargano, A., Giacoletti, M., 2026. Subsidizing the cloud: U.S. state incentives to data centers. SSRN Working Paper No. 5881105.

Giglio, S., Maggiori, M., Stroebel, J., 2015. Very long-run discount rates. Quarterly Journal of Economics 130, 1–53.

Goodman-Bacon, A., 2021. Difference-in-differences with variation in treatment timing. Journal of Econometrics 225, 254–277.

Greenstone, M., Hornbeck, R., Moretti, E., 2010. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. Journal of Political Economy 118, 536–598.

Greenstone, M., Moretti, E., 2003. Bidding for industrial plants: Does winning a 'million dollar plant' increase welfare? Working Paper 9844. National Bureau of Economic Research.

Gyourko, J., Hartley, J.S., Krimmel, J., 2021. The local residential land use regulatory environment across U.S. housing markets: Evidence from a new Wharton index. Journal of Urban Economics 124, 103337.

Hainmueller, J., 2012. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. Political Analysis 20, 25–46.

Hilber, C.A.L., 2017. The economic implications of house price capitalization: A synthesis. Real Estate Economics 45, 301–339.

Hines, J.R., Thaler, R.H., 1995. Anomalies: The flypaper effect. Journal of Economic Perspectives 9, 217–226.

Jensen, M.C., 1986. Agency costs of free cash flow, corporate finance, and takeovers. American Economic Review 76, 323–329.

Joint Legislative Audit and Review Commission, 2024. Data Centers in Virginia. JLARC Report 598. Commonwealth of Virginia. URL: https://jlarc.virginia.gov/pdfs/reports/Rpt598.pdf.

Kline, P., Moretti, E., 2014. Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority. Quarterly Journal of Economics 129, 275–331.

La Porta, R., Lopez-de Silanes, F., Shleifer, A., Vishny, R.W., 2000. Agency problems and dividend policies around the world. Journal of Finance 55, 1–33.

Loudoun County, Virginia, 2024. FY 2025 adopted budget, volume 1: Executive summary. Office of Management and Budget, Loudoun County, Virginia.

MacKinnon, J.G., Webb, M.D., 2017. Wild bootstrap inference for wildly different cluster sizes. Journal of Applied Econometrics 32, 233–254.

Meeks, R., Pless, J., Qi, Z., Wang, Z., 2026. Who pays for growth? Evidence from data centers and the grid. SSRN Working Paper No. 6967558.

Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. Quarterly Journal of Economics 128, 1687–1726.

Moretti, E., 2010. Local multipliers. American Economic Review 100, 373–377.

Muehlenbachs, L., Spiller, E., Timmins, C., 2015. The housing market impacts of shale gas development. American Economic Review 105, 3633–3659.

Oates, W.E., 1969. The effects of property taxes and local public spending on property values: An empirical study of tax capitalization and the Tiebout hypothesis. Journal of Political Economy 77, 957–971.

Painter, M., 2020. An inconvenient cost: The effects of climate change on municipal bonds. Journal of Financial Economics 135, 468–482.

Pástor, Ľ., Veronesi, P., 2009. Technological revolutions and stock prices. American Economic Review 99, 1451–1483.

Prince William County, Virginia, 2022. Data center fiscal impact analysis. Prepared by PFM Group Consulting LLC.

Saiz, A., 2010. The geographic determinants of housing supply. Quarterly Journal of Economics 125, 1253–1296.

Schwert, M., 2017. Municipal bond liquidity and default risk. Journal of Finance 72, 1683–1722.

Slattery, C., 2025. Bidding for firms: Subsidy competition in the United States. Journal of Political Economy 133, 2563–2614.

Slattery, C., Zidar, O., 2020. Evaluating state and local business incentives. Journal of Economic Perspectives 34, 90–118.

Stulz, R.M., 1990. Managerial discretion and optimal financing policies. Journal of Financial Economics 26, 3–27.

Suárez Serrato, J.C., Zidar, O., 2016. Who benefits from state corporate tax cuts? A local labor markets approach with heterogeneous firms. American Economic Review 106, 2582–2624.

Sun, L., Abraham, S., 2021. Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics 225, 175–199.

Titman, S., Wei, K.C.J., Xie, F., 2004. Capital investments and stock returns. Journal of Financial and Quantitative Analysis 39, 677–700.

Van Nieuwerburgh, S., 2026. Financing the AI buildout. SSRN Working Paper No. 6444363. Journal of Economic Perspectives, forthcoming.

Wolfinger, R.E., Rosenstone, S.J., 1980. Who Votes? Yale University Press, New Haven, CT.

Yue, D., Zeng, Y., 2026. The local economic effects of data center entry. SSRN Working Paper No. 6497238.

Appendix A: The Price of Consent

This appendix formalizes the framework of Section II: an official allocates a data-center windfall between retention and a payout that secures residents' consent, the payout capitalizes into incumbent home values, and the four hypotheses emerge as comparative statics.

Setup. A data center is a contested land use: hosting requires residents' consent, and the payout we model is the price officials pay for it. Siting creates a risky incremental property-tax stream, the delta tax cash flow. Let $W > 0$ be its annual value, with mean $\mathbb{E}[W]$ held fixed and a mean-preserving spread $\sigma \geq 0$ indexing the risk that the base does not persist, and let $\kappa = 1 - \text{abatement} \in (0, 1]$ be the fraction the county collects, fixed in the entry bargain with the firm, so $\kappa\,\mathbb{E}[W]$ is the expected windfall. An incumbent official divides it between a payout $d \geq 0$ to residents, delivered as a property-tax rate cut, and retention $r \geq 0$, worth a private benefit $\delta \in (0, 1)$ per dollar to the official. Both are drawn from the windfall alone, leaving the pre-existing budget and the claims of schools and bondholders untouched. Residents are risk neutral but cannot observe how retained dollars are spent, so they expect a retained dollar to return only $\theta \in [0, 1]$ of its value in services, with θ the quality of local governance: nothing leaks at $\theta = 1$, nothing returns at $\theta = 0$. The problem is a government's version of corporate payout policy, the consent constraint its Coasean price of local costs (Coase, 1960), and the model an organizing device for the empirical analysis rather than a claim to generality.

The price of consent. The facility imposes a perceived local cost C, with $0 < C < \kappa\,\mathbb{E}[W]$, so the windfall covers the nuisance. Residents consent only if what they expect to receive offsets that cost, and they do not value the windfall's two flows alike. The payout is enacted when consent is granted, handed over rather than promised, so residents take the committed d as given and count it at face value. The retained remainder is an unsecured claim on future services, worth θ on the dollar to residents who cannot observe how it is spent. Consent therefore requires $d + \theta\,(\kappa\,\mathbb{E}[W] - d) \geq C$, and for $\theta < 1$ this defines the price of consent,

$$d \;\geq\; \bar{d}(\theta) = \frac{C - \theta\,\kappa\,\mathbb{E}[W]}{1 - \theta}, \qquad \bar{d}'(\theta) = \frac{C - \kappa\,\mathbb{E}[W]}{(1-\theta)^2} < 0. \tag{A.1}$$

At $\theta = 1$ the constraint reads $\kappa\,\mathbb{E}[W] \geq C$ and holds by assumption. The floor falls as trust rises, from $\bar{d}(0) = C$ under full opacity, where the entire cost must be paid in the currency residents can verify, to zero at $\theta = C/(\kappa\,\mathbb{E}[W])$, beyond which consent is free. Securing consent is always feasible, since $\bar{d}(\theta) \leq C < \kappa\,\mathbb{E}[W]$: residents never demand more than the harm, and the windfall exceeds it by assumption. Residents charge for consent whatever the retained surplus cannot credibly deliver.

The electoral margin. The official's own incentive to pay determines when the consent price binds. A dollar paid out earns re-election value $\alpha\,\pi(d)$, with $\alpha > 0$ the political weight of homeowners and $\pi' > 0$, $\pi'' < 0$: visible relief wins votes, at a diminishing rate. A dollar retained keeps its private benefit δ. Because the nearly irreversible cut is committed before W is realized, a committed dollar carries an irreversibility premium $\phi(\sigma)$, $\phi(0) = 0$, $\phi' > 0$, the reduced-form price of the shortfall borne if the base decays, so risk enters the ex ante problem, written against $\mathbb{E}[W]$, through ϕ. The official solves

$$\max_{d \geq 0,\, r \geq 0} \;\alpha\,\pi(d) + \delta\,r - \phi(\sigma)\,d \quad \text{s.t.} \quad d + r \leq \kappa\,\mathbb{E}[W], \quad d \geq \max\{0,\; \bar{d}(\theta)\}. \tag{A.2}$$

Since $\delta > 0$ the budget binds, $r^\star = \kappa\,\mathbb{E}[W] - d^\star$, leaving the objective $V(d) = \alpha\,\pi(d) - [\delta + \phi(\sigma)]\,d + \delta\,\kappa\,\mathbb{E}[W]$, strictly concave because $\alpha > 0$ and $\pi'' < 0$. We assume hosting is worthwhile at the consent price, so the sample, conditioning on entry, lies in this region. Let $\hat{d}$ maximize V on $[0, \infty)$: if $\alpha\,\pi'(0) > \delta + \phi(\sigma)$, then $\hat{d} > 0$ solves $\alpha\,\pi'(\hat{d}) = \delta + \phi(\sigma)$, equating the marginal vote value of relief with its all-in cost, and otherwise $\hat{d} = 0$. By concavity the optimum projects $\hat{d}$ onto the feasible interval,

$$d^\star = \min\big\{\kappa\,\mathbb{E}[W],\ \max\{\hat{d},\ \bar{d}(\theta),\ 0\}\big\}, \tag{A.3}$$

so the payout is the larger of what the official wants to pay and what residents charge, where the cap can bind only through $\hat{d}$, since the floor never exceeds C. When $\hat{d} < C$, the strictly decreasing floor crosses $\hat{d}$ at a unique $\hat{\theta} \in (0, 1)$ with $\bar{d}(\hat{\theta}) = \hat{d}$: residents' price sets the payout where trust is low, $\theta < \hat{\theta}$, and the electoral optimum sets it above. When $\hat{d} \geq C$ the floor never binds. The empirically relevant case is $\hat{d} < C$: votes alone pay less than the most distrustful residents charge.

Comparative statics. The payout is a recurring reduction in the home's tax bill and capitalizes into incumbent home values at the discount rate $\rho(\sigma) > 0$, so $\Delta P = d^\star/\rho(\sigma)$. The rate rises with risk, $\rho'(\sigma) > 0$, because a less durable base shortens the expected life of the cut, not because residents are risk averse. The two margins divide the predictions: the consent floor carries the governance prediction and the electoral optimum carries the risk prediction. Hypotheses 1–4 of Section II follow.

Payout. $d^\star > 0$ on two independent grounds: $d^\star \geq \hat{d} > 0$ when $\alpha\,\pi'(0) > \delta + \phi(\sigma)$, and $d^\star \geq \bar{d}(\theta) > 0$ whenever $\theta < C/(\kappa\,\mathbb{E}[W])$. The official grants relief either because a visible cut wins votes or because residents will not consent to the facility without being paid, and relief is absent only where both motives fail at once.

Capitalization. $\Delta P > 0$ whenever $d^\star > 0$: a flow of relief becomes a stock of housing wealth, and it accrues to incumbents alone, since renters hold no claim and later buyers pay the capitalized price. Because the cut is committed at entry, the price moves when the commitment becomes public rather than when the lower bill arrives.

Risk. Risk strikes twice, once in each part of the capitalization: at an interior $\hat{d}$, the first-order condition gives $\partial\hat{d}/\partial\sigma = \phi'(\sigma)/(\alpha\,\pi''(\hat{d})) < 0$, so less is granted, and the rising discount rate shrinks the value of whatever is paid. The payout channel runs through the electoral margin alone, since risk-neutral residents value the expected surplus and the floor does not move with σ. An official does not commit irreversible relief against a base that may leave, and the market pays less for a saving it may lose.

Governance. Where the floor binds, $d^\star = \bar{d}(\theta)$ and $\partial d^\star/\partial\theta = (C - \kappa\,\mathbb{E}[W])/(1-\theta)^2 < 0$: the weaker the governance, the higher the price of consent and the larger the payout. Opacity makes consent expensive because residents discount every dollar the government keeps and must be paid in the one currency they can verify, the visible rate cut. The prediction is one-sided in the empirical range: $d^\star$ falls strictly in θ below $\hat{\theta}$ and is flat at $\hat{d}$ above, so the payout is larger where public integrity is weaker, the substitute-dividends logic transplanted to the county.

Appendix B: Variable Definitions

This table defines the variables used in the analysis, one variable per row. Panel A defines the outcome variables in the order the results examine them. Panel B defines the county and facility characteristics, including the balancing covariates and the governance measures. Panel C defines the sample classifications. The treatment variables are defined in Section IV. ACS denotes the American Community Survey; ZHVI denotes the Zillow Home Value Index; S&P 451 denotes the S&P Global Market Intelligence / 451 Research Data Center Database; NCES denotes the National Center for Education Statistics; SEDA denotes the Stanford Education Data Archive.

Panel A: Outcome Variables

Variable	Definition	Source
Fiscal Outcomes		
County Property-Tax Base	Property-tax revenue of the county government, measured as an annualized growth rate between the 2017 and 2022 Censuses of Governments.	Census of Governments
All-Local Property-Tax Base	Property-tax revenue of all governments layered on the county's tax base (county, school district, city/township, and special districts), aggregated to the county.	Census of Governments
State Aid	Total intergovernmental revenue received by the county from the state.	Census of Governments
Property-Tax Rate		
Effective Property-Tax Rate	Median real-estate taxes paid divided by median home value, in percentage points, at the county level.	ACS B25103 / B25077
Millage Rate	Nominal property-tax rate per dollar of assessed value, adopted annually by the county's governing body; hand-collected panel of 1,054 counties in twenty-four states, 1996–2026.	county budget ordinances
Home Prices and Rents		
ACS Median Home Value (Five-Year, Overlapping)	Median owner-occupied home value from overlapping five-year ACS windows.	ACS B25077
ACS Median Home Value (One-Year)	Median owner-occupied home value from one-year ACS estimates, available for larger counties.	ACS B25077
ACS Median Home Value (Five-Year, Non-Overlapping)	Median owner-occupied home value from non-overlapping five-year ACS windows, survey years 2013, 2018, and 2023.	ACS B25077
Zillow Home Value Index	The ZHVI at the county and ZIP-code level, 2008–2025.	Zillow
Median Gross Rent	Median gross rent of renter-occupied units.	ACS B25064
Housing Wealth	The home-price effect multiplied by the aggregate value of owner-occupied housing in the county.	ACS; Zillow
Other Claimants		
New-Issue Spread	Offering spread of newly issued municipal bonds over a maturity-matched Treasury, in basis points.	SDC Public Finance
Student Test Scores	Standardized student achievement, in standard deviations, aggregated from school-district scores to the county-year and measured three years after entry.	SEDA
Teacher FTE per 1,000 Students	Full-time-equivalent teachers per one thousand enrolled students.	NCES Common Core of Data
Teacher Salary per Student	Teacher salary expenditure per enrolled student.	NCES Common Core of Data

Continues on next page

Continued from last page

Panel B: County and Facility Characteristics

Variable	Definition	Source
County Economic Characteristics (Table 1, measured pre-entry)		
Population	County population, measured in 2012; a balancing covariate in the full-covariate match.	ACS B01003
Median Household Income	Median household income, measured in 2012; a balancing covariate in the full-covariate match.	ACS B19013
College Degree or Higher	Share of adult residents holding a bachelor's degree or higher, 2008–2012 five-year estimates, in percent.	ACS
Power Plants	Number of power plants located in the county, measured in 2012; a balancing covariate in the full-covariate match.	EIA Form 860
Census Division	The county's Census division, entering the full-covariate match as indicator variables.	Census Bureau
Data Center Characteristics (Table 1)		
Number of Data Centers	Count of operating data centers in the county.	S&P 451
Data Center Share of Tax Base	Estimated facility-attributable property tax divided by the county's property-tax revenue, in percent.	S&P 451; Census of Governments
Announcement Year	Year of the first public announcement of the county's first facility of at least 50 megawatts.	hand-collected
Residential Property Characteristics (Table 1)		
Median Home Value	Median owner-occupied home value, measured in 2012; a balancing covariate in the full-covariate match.	ACS B25077
Zillow Home Value Index	The county-level ZHVI, measured in 2016.	Zillow
County Property-Tax Revenue	Property-tax revenue of the county government, fiscal year 2017, in millions of dollars.	Census of Governments
Effective Property-Tax Rate	The effective property-tax rate of Panel A, measured in 2016.	ACS
Homeownership Rate	Share of occupied housing units that are owner-occupied, measured in 2012.	ACS
Pre-Period Growth Trajectories (Table 1 and balancing covariates)		
Pre-Period Income Growth	Growth in median household income over 2012–2016, in percent; a balancing covariate in the primary entropy-balancing specification.	ACS
Pre-Period Home-Value Growth	Growth in median home value over 2012–2016, in percent; the other primary balancing covariate.	ACS; Zillow

Continues on next page

Continued from last page

Panel B: County and Facility Characteristics (Continued)

Variable	Definition	Source
Geography and Housing Supply		
Distance to Facility	Distance from the ZIP-code centroid to the nearest data center, in miles.	S&P 451; ZCTA crosswalk
Building-Permit Intensity	Residential building permits issued, relative to the county's housing stock; a balancing covariate in the full-covariate match.	Census Building Permits Survey
Saiz Elasticity	The metropolitan housing-supply elasticity of Saiz (2010).	Saiz (2010)
WRLURI	The Wharton Residential Land Use Regulatory Index, 2018 survey.	Gyourko, Hartley and Krimmel (2021)
Baum-Snow–Han Elasticity	Tract-level housing-supply elasticities, aggregated to the county.	Baum-Snow and Han (2024)
County Governance		
Professionalized County	Indicator equal to one if the county government has a professional appointed manager; commission counties are the complement.	hand-collected
GFOA Certificate	Indicator equal to one if the county holds the Government Finance Officers Association reporting certificate, among counties whose status is verified.	GFOA
Paid Out	The annual property-tax saving per home implied by the county's adopted rate cut after entry, in dollars per year.	county budget ordinances; ACS
Retained	The annualized change in the county government's cash and security holdings per capita after entry, in dollars per year.	Census of Governments
Opacity	Indicator equal to one if the county has thin local news coverage.	Medill Local News Initiative

Panel C: Sample Classifications

Variable	Definition	Source
Data Center County	County with at least one verified operating data center.	S&P 451; hand-verified
No Data Center County	County with no listed data center of any size at any time in the facility database, verified at the cleaning stage.	S&P 451; hand-verified
Non-Crypto County	Data center county whose facilities are hyperscale, colocation, or enterprise data centers.	S&P 451
Crypto County	Data center county whose capacity is cryptocurrency mining.	S&P 451
Statutory Adjustment State	State whose statutes recalibrate county property-tax rates upon the revaluation of existing property, with newly constructed property excluded from the calculation.	state statutes; hand-verified
Discretionary State	State where county property-tax rates change only by an act of the county's governing body.	state statutes; hand-verified
Assessment Cap State	State that caps growth in assessed values rather than tax rates.	state statutes; hand-verified
Never-Built County	County with a publicly announced project that was never built and less than 10 MW of other capacity; the placebo group.	hand-collected

Table 1: Summary Statistics

This table reports summary statistics for counties with and without a data center. The Data Center County group comprises every county with at least one operating data center, and the No Data Center County group comprises counties that never receive a data center. For each group the table reports the mean and standard deviation, followed by the difference in group means (Diff.). Variables are grouped into county economic characteristics, data center characteristics, residential property characteristics, and pre-period growth trajectories, with the data center characteristics reported for the data center counties only. The data are from the American Community Survey, Zillow, the Census of Governments, the S&P 451 database, and EIA Form 860. Significance of the differences is assessed with Welch two-sample t-tests. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables, including their measurement dates, are defined in Appendix B.

	Data Center County		No Data Center County		
Variable	Mean	St. Dev.	Mean	St. Dev.	Diff.
County Economic Characteristics					
Population	493,839	808,409	42,664	80,214	451,175***
Median Household Income ($)	54,654	15,342	44,043	10,474	10,611***
College Degree or Higher (%)	28.62	10.97	18.04	7.49	10.58***
Power Plants (Count)	6.60	12.89	1.65	2.86	4.95***
Data Center Characteristics					
Number of Data Centers (Count)	8.04	17.43	–	–	–
Data Center Size (MW)	122.01	340.47	–	–	–
Data Center Share of Tax Base (%)	7.68	24.91	–	–	–
Residential Property Characteristics					
Median Home Value ($)	196,302	112,549	121,680	69,925	74,622***
Zillow Home Value Index ($)	227,747	144,932	145,223	83,503	82,524***
County Property-Tax Revenue ($M)	226	508	16	35	210***
Effective Property-Tax Rate (%)	1.24	0.59	1.02	0.49	0.22***
Homeownership Rate (%)	67.35	9.19	73.45	7.16	−6.09***
Pre-Period Growth Trajectories					
Income Growth (%)	5.35	5.47	5.31	8.73	0.04
Home-Value Growth (%)	3.81	10.39	5.99	12.66	−2.18***
N	340		2,626		

Table 2: Effect of Data Centers on County's Property-Tax Base

This table reports the effect of hosting a data center on the growth of the county's property-tax base and on the state aid it receives. Each estimate is a matched-pair difference-in-differences, the data center county's annualized growth between the two most recent quinquennial Censuses of Governments less the mean of its matched control counties. Each data center county is matched to three control counties by caliper matching on the pre-period log property-tax base. The matched sample comprises every county with at least one operating data center, and pair counts vary by layer with census coverage. Panel A reports the property-tax base by layer of local government. Panel B reports total state aid received by the county. The data are from the Census of Governments and audited Annual Comprehensive Financial Reports. The reported p-values are from paired t-tests and from Wilcoxon signed-rank tests over the per-pair effects. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

Panel A: Property-Tax Base

Government Layer	N	Δ Property-Tax Base (%/yr)	p (t-test)	p (Wilcoxon)
All Local Governments	309	1.05^{***}	<0.001	<0.001
County	308	1.21^{***}	0.003	0.001
City / Township	284	0.72^{*}	0.086	0.006
Special Districts	202	-0.28	0.818	0.957

Panel B: State Aid

Government Layer	N	Δ State Aid (%/yr)	p (t-test)	p (Wilcoxon)
County	298	0.53	0.533	0.137

Table 3: Effect on Homeowners' Property-Tax Rates

This table reports difference-in-differences estimates of the effect of data centers on homeowners' effective property-tax rate, in percentage points, the margin through which the payout hypothesis operates. The treatment variable is Data Center Size, the cumulative installed power of the county's operating facilities measured in gigawatts, which rises as facilities come online and is zero for counties with no data center, so each coefficient is the change in the rate per gigawatt. The Baseline column is the two-way fixed-effects specification with county and state-by-year fixed effects. The Entropy (Pre-Trend) column reweights the control counties by entropy balancing on the data center counties' pre-period growth trajectories, our primary weighting specification, and the Entropy (Full) column additionally matches the pre-period levels of home value, income, and population, pre-period building permits, the county's power-plant count, and census-division indicators. The estimation sample comprises every county with a data center, with no materiality screen, and the never-treated control pool. The data are from the American Community Survey. Standard errors clustered at the county level are reported in parentheses, and exact p-values are reported in the text. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	Baseline	Entropy (Pre-Trend)	Entropy (Full)
Data Center Size	-0.117^{***}	-0.092^{***}	-0.051^{**}
	(0.029)	(0.018)	(0.025)
County FE	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Effective Control Counties	2,635	2,529	76
N	44,424	44,424	44,424

Table 4: Property-Tax Rate Cuts by Statutory Regime

This table reports the effective-property-tax-rate response to data center capacity within each statutory rate-limit regime. Every state is classified by whether county property-tax rates change only by a discretionary act of the county's governing body (Discretionary), are constrained through caps on assessed values rather than rates (Assessment Cap), or are recalibrated under statute when existing property is revalued (Statutory Adjustment). The adjustment statutes apply to the revaluation of existing property and exclude newly constructed property, so they do not apply to the arrival of a data center. Each row reports the coefficient on *Data Center Size* interacted with an indicator for the named regime, with the effective property-tax rate as the dependent variable. Estimates apply the Entropy (Pre-Trend) weights of the preceding table, and counties in the few states with no rate-limiting statute are included but not displayed. Standard errors, clustered by county, are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	Effective Tax Rate
Data Center Size × Discretionary	-0.143^{***}
	(0.038)
Data Center Size × Assessment Cap	-0.093^{*}
	(0.055)
Data Center Size × Statutory Adjustment	-0.089^{***}
	(0.019)
County FE	Yes
State × Year FE	Yes
N	44,424

Table 5: Evidence from Neighboring Counties

This table reports estimates from a border county-pair difference-in-differences design in which each data center county is paired with nearby control counties that have no data center, and pair and pair-by-year fixed effects replace the county and state-by-year fixed effects. The dependent variable is homeowners' effective property-tax rate, in percentage points, and the treatment variable is Data Center Size, so each coefficient is the change in the rate per gigawatt. The columns report three definitions of a pair: contiguous pairs across a state line, all contiguous pairs, and pairs whose centroids lie within 50 miles, whether or not contiguous. Each data center county is paired with every county without a data center that meets the column's definition, with distance measured between county centroids, and the treated counties exclude cryptocurrency-mining and oil-extraction counties. The data are from the American Community Survey. Standard errors clustered at the treated-county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	Cross-State Pairs	Contiguous Pairs	Distance $\leq$50 mi Pairs
Data Center Size	-0.093^{**}	-0.056^{**}	-0.064^{***}
	(0.042)	(0.028)	(0.017)
Pairs	55	304	609
N	1,540	8,508	17,042

Table 6: Effect of Data Centers on Home Prices

This table reports difference-in-differences estimates of the effect of a data center on home prices in the data center counties. The dependent variable is the log median home value, measured four ways. The treated group comprises every data center county. The ACS5 Overlap and ACS5 Non-Overlap columns use the overlapping and non-overlapping five-year median home values from the American Community Survey, whose successive windows either share or exclude common survey years. The ACS1 column uses the survey's one-year median, available only for larger counties. The ZHVI column uses the county-level Zillow Home Value Index. Estimates are expressed in percent. All specifications include county and state-by-year fixed effects. The data are from the American Community Survey and Zillow. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	ACS5 Overlap	ACS1	ACS5 Non-Overlap	ZHVI
Treated × Post	2.21***	3.81***	2.87**	3.78***
	(0.85)	(1.44)	(1.30)	(1.37)
County FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
N	38,404	6,509	8,219	41,081

Table 7: Alternative Model Specifications

This table reports the home-price effect under alternative specifications, with the log Zillow Home Value Index as the dependent variable throughout. Estimates are expressed in percent, and the baseline specification itself is reported in the ZHVI column of the preceding table. Column (1) replaces the baseline's state-by-year fixed effects with year fixed effects. Column (2) adds the county's effective property-tax rate as a control, excluding one data center county without rate coverage. Column (3) restricts the panel to counties observed in every year. The data are from Zillow. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	(1)	(2)	(3)
Treated × Post	3.82**	3.58***	4.62***
	(1.56)	(1.35)	(1.79)
County FE	Yes	Yes	Yes
State × Year FE		Yes	Yes
Year FE	Yes		
Effective Property-Tax Rate Control		Yes	
Balanced Panel			Yes
N	41,087	32,923	40,100

Table 8: Home Prices by Distance to the Nearest Data Center

This table reports the home-price response by distance to the nearest data center. The sample is a ZIP-code panel, with treatment assigned at the county level, and the treated group comprises every data center county. Each treated ZIP code is placed in a ring by the distance from its centroid to the nearest facility. N counts the treated ZIP-code years in each ring. The dependent variable is the log Zillow Home Value Index, and estimates are expressed in percent. Because distance is measured from the ZIP centroid rather than the home, a fine local gradient would be attenuated toward zero. All specifications include ZIP and state-by-year fixed effects. The data are from Zillow and the S&P 451 census. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	0–3 mi	3–10 mi	10–20 mi
Treated × Post × Ring	4.74*	4.63***	3.01*
	(2.81)	(1.48)	(1.56)
County FE	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
N	5,050	8,128	4,633

Table 9: Home Prices and Housing-Supply Elasticity

This table reports whether the home-price response to a data center varies with housing-supply elasticity. Each column interacts entry with a different elasticity measure: realized building-permit intensity, the geographic supply elasticity of Saiz (2010), the Wharton Residential Land Use Regulatory Index (WRLURI) of Gyourko, Hartley and Krimmel (2021), and the tract-level supply elasticities of Baum-Snow and Han (2024) aggregated to counties. The dependent variable is the log Zillow Home Value Index, and estimates are expressed in percent. A positive interaction would indicate a response concentrated in inelastic markets, the pattern a demand shock produces. The Saiz and WRLURI measures cover the metropolitan subset of counties, and the Baum-Snow and Han measure covers a comparable set. All specifications include county and state-by-year fixed effects. The data are from Zillow and the sources named. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	Building Permits	Saiz (2010) Elasticity	WRLURI	Baum-Snow–Han (2024) Elasticity
Treated × Post × Elasticity	−2.09	3.95	5.80	−3.11
	(2.70)	(3.78)	(3.90)	(3.75)
County FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
N	41,052	40,063	40,083	40,042

Table 10: Effect of Tax-Base Risk on Home Prices

This table reports the effect of a data center on home prices for the non-crypto data center counties, whose facilities are hyperscale, colocation, or enterprise data centers. The risk hypothesis predicts that speculative, exit-prone tax bases capitalize less, so the treated group is restricted to counties with durable tax bases. The dependent variable is the log Zillow Home Value Index, and estimates are expressed in percent. The columns report an unweighted baseline followed by two entropy-balancing specifications: the Entropy (Pre-Trend) column matches the two pre-period growth trajectories alone, our primary weighting specification, and the Entropy (Full) column matches, in addition to the growth trajectories, the pre-period levels of home value, income, and population, pre-period building permits, the county's power-plant count, and census-division indicators. All specifications include county and state-by-year fixed effects. The data are from Zillow. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	Baseline	Entropy (Pre-Trend)	Entropy (Full)
Treated × Post	7.98***	4.55***	4.16***
	(1.75)	(1.42)	(1.49)
County FE	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes
Effective Control Counties	2,564	2,521	247
N	40,380	40,380	40,380

Table 11: Risk Resolution and Home-Price Response

This table reports when the home-price response arrives, for the non-crypto data center counties under our primary weighting specification, the Entropy (Pre-Trend) weights of the preceding table. The dependent variable is the log Zillow Home Value Index, and estimates are expressed in percent. Column (1) dates treatment from the county's first data center announcement rather than from the start of operations, with the announcement clock defined as the earlier of the two dates. Column (2) includes both event clocks, so $\text{Post}_{\text{Announcement}}$ is the response that follows the announcement and $\text{Post}_{\text{Operation}}$ is the additional increment once operations begin. All specifications include county and state-by-year fixed effects. The data are from Zillow. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	(1)	(2)
Treated $\times$ $\text{Post}_{\text{Announcement}}$	4.91***	3.69**
	(1.55)	(1.58)
Treated $\times$ $\text{Post}_{\text{Operation}}$		1.67
		(1.15)
County FE	Yes	Yes
State $\times$ Year FE	Yes	Yes
N	40,370	40,370

Table 12: Risk Pricing: Crypto versus Non-Crypto Data Centers

This table tests whether the crypto differential in capitalization reflects the riskiness of the tax base rather than the size of the facility. The dependent variable is the log Zillow Home Value Index, and estimates are expressed in percent. Column (1) interacts Post with an indicator for non-crypto counties, so the Post row is the entry effect for crypto counties and the Post × Non-Crypto row is the capitalization premium for durable tax bases. Column (2) adds an interaction of Post with the log of data center size, column (3) with the data center's share of the county tax base, and column (4) with both. The interacted variables are mean-centered, so the Post row is the crypto entry effect at the treated-county means. All specifications include county and state-by-year fixed effects. The data are from Zillow, the S&P 451 census, and the Census of Governments. Standard errors clustered at the county level are reported in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

	(1)	(2)	(3)	(4)
Post	−1.11	0.04	−1.14	0.66
	(1.79)	(1.83)	(1.78)	(1.89)
Post × Non-Crypto	9.08***	7.70***	9.13***	6.91***
	(2.48)	(2.42)	(2.51)	(2.57)
Post × Log(Data Center Size)		1.90**		2.66***
		(0.78)		(0.99)
Post × Tax-Base Share			0.02	−0.10
			(0.09)	(0.10)
County FE	Yes	Yes	Yes	Yes
State × Year FE	Yes	Yes	Yes	Yes
N	41,081	41,081	41,081	41,081

Table 13: Effects of County Governance

This table reports how county governance shapes the division and pricing of the windfall. Panel A reports subsample difference-in-differences estimates of Data Center Size, the county's cumulative installed capacity in gigawatts, with the effective property-tax rate as the dependent variable, within the two weakly governed subsamples: commission counties, which operate without a professional appointed manager, and counties verified not to hold the Government Finance Officers Association (GFOA) reporting certificate. Governance is observed for 51 data center counties, and GFOA status is verified for 33 of them. Panel B reports, with one observation per treated county (36 counties, and 34 in column (3)), the change in log home value from the year before entry to three years after, regressed on the annual tax saving per home implied by the county's adopted nominal rate cut (Paid Out) and on the annualized change in the county government's cash and security holdings per capita (Retained), with heteroskedasticity-robust standard errors in parentheses. Column (1) includes the paid-out flow alone, column (2) adds the retained flow, and column (3) additionally interacts Retained with thin local news coverage (Opacity) and excludes the two Nevada counties, whose nonstandard posted-rate concept and assessment ratio make their dollar conversions the least reliable in the sample. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

Panel A: Rate Response by County Governance

	Data Center Size	p
Commission County	−0.085	0.453
No GFOA Certificate	−0.105	0.282

Panel B: The Retained-Dollar Test

	Δ Log Home Value $\times$ 100, Entry−1 to Entry+3		
	(1)	(2)	(3)
Paid Out ($/Year per Home)	0.026***	0.034***	0.038***
	(0.004)	(0.007)	(0.009)
Retained ($/Year per Capita)		0.009**	0.024***
		(0.004)	(0.008)
Retained $\times$ Opacity			−0.015**
			(0.008)

Table 14: Effects on Other Claimants

This table reports the effect of a data center on the county's other claimants: renters, bondholders, and schools. Each row reports a difference-in-differences estimate of Treated × Post for the variable listed: the median gross rent of renter-occupied units, in percent change; the spread on newly issued municipal bonds, in basis points; student test scores, in standard deviations, measured three years after entry; teacher full-time equivalents per one thousand students; and teacher salary per student, both in percent changes. The rent row is estimated on the all-data-center county sample, the bond spread at the deal level across all data center counties, and the school rows on the non-crypto county sample. The p-value is reported alongside each estimate. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively. All variables are defined in Appendix B.

Claimant	Variable	β	p
Renters	Median Gross Rent (% change)	−0.53	0.54
Bondholders	New-Issue Spread (bps)	4.98	0.11
Schools	Student Test Scores (SD, 3-Year Lag)	0.03	0.73
	Teacher FTE per 1,000 Students (% change)	−2.91	0.15
	Teacher Salary per Student (% change)	0.58	0.60

Figure 1: Home Values around Data-Center Entry

This figure plots event-study estimates of home values, in log points times 100 relative to the year before the event, with 95% confidence intervals. The Data Center series comprises the 279 counties with a verified operating facility, on an entry clock. The Announced but Never Built Data Center series comprises the 18 counties where a publicly announced project was never built, aligned on the announcement year. Estimates use county and state-by-year fixed effects, with standard errors clustered by county. All variables are defined in Appendix B.

