Welfare Gains from Market Integration Under Competing Land Uses*

Dustin Frye† Jeffrey Hadachek‡

August 20, 2026

Preliminary Conference Draft

Abstract

The welfare gains from agricultural market integration are conventionally measured by the capitalization of market access into farmland values – an aggregate measure that omits parcels converting to non-agricultural use, which exit the Census of Agriculture's sampling frame and render farmland values a lower-bound measure of welfare gains. Using the staggered rollout of the U.S. Interstate Highway System (1940–2002), we trace where the gains accrued. Farmland values rose about 11% after a highway arrived, masking a divergence between value and cultivated acreage: the least-suitable land was retained through a shift toward livestock, middle-suitability land lost crop acreage to conversion where development pressure bound, and the best land capitalized the development option into farmland value. Combining this heterogeneity with a Fréchet productivity model and a minimum-distance GMM over development-pressure bins, we identify the parcels the aggregate omits and value them at observed developer bids at as much as $195–210 billion, concentrated on the high-development pressure fringe.

KEYWORDS: highways, agriculture, market integration, comparative advantage, farmland values

JEL CLASSIFICATION: R40, Q15, N52, F14

*Thanks to seminar participants at Kansas State University and the University of Wisconsin-Madison for helpful comments and suggestions. We would also like to thank Ting Liu for their excellent research assistance. All errors are our own.

†Department of Agricultural and Applied Economics, University of Wisconsin-Madison, 427 Lorch St., Madison, WI 53706 (e-mail: dustin.frye@wisc.edu; website: http://www.dustindfrye.com)

‡Department of Agricultural and Applied Economics, University of Wisconsin-Madison, 427 Lorch St., Madison, WI 53706 (e-mail: hadachek@wisc.edu; website: https://jhadachek.github.io)

1 Introduction

Welfare gains from agricultural market integration are conventionally measured by the capitalization of market access into farmland values (Allen and Arkolakis, 2014; Donaldson and Hornbeck, 2016). That measure is a lower bound wherever integration also raises the non-agricultural option value of land, especially at the urban periphery (Baum-Snow, 2007): a parcel that converts to developed use exits the Census of Agriculture's sampling frame, so the value it carries into its new use leaves the farmland measure entirely. This paper measures and bounds that omission, using the staggered rollout of the US Interstate Highway System as a setting where competing, non-agricultural land-use margins are active.

Our analysis centers around the capitalization of market access into county-level farmland values and changes in acreage actually under agricultural production in a county. Prior work (Donaldson and Hornbeck, 2016; Chan, 2022) on transportation and US agriculture studies railroads on an open frontier where land had fewer potential outside claimants. The IHS was introduced into a different context. The era of large-scale land settlement had passed, and output gains could no longer come from bringing new territory into farming, as they did in the railroad era. Land markets were active, non-agricultural demand was substantial, and farmland values faced competing pressures. Because farm returns and outside-option value enter the same land use margin symmetrically, no single land-value measure can separate them. The standard approach to measuring welfare gains from market access in agriculture is therefore structurally a lower bound everywhere land markets are active enough for this margin to matter.

We begin by estimating the reduced-form impacts of the Interstate Highway System. The staggered opening of Interstate Highway segments across rural counties, concentrated in 1954–1969, provides variation in when individual counties gained highway access, driven by state-level construction priorities rather than contemporaneous agricultural conditions. We use county-level outcome variables from the US Ag Census between

1940 and 2002. We exploit the staggered rollout as a difference-in-differences design; the imputation estimator of Gardner et al. (2024) recovers the average effect of IHS access on farmland values and agricultural production in counties that received a highway.

The average treatment effect of Interstate Highways on farmland values cannot decompose whether the asset-price gain comes through agricultural returns or outside-option appreciation. Development pressure is not directly observed, so we construct a pre-interstate, fine-grained measure based on proximity to urbanized areas, similar to (Burchfield et al., 2006). We first calculate a share of each county's land that is not yet developed based on a gridded, land classification dataset from 1940 (Leyk and Uhl, 2018). We then compute the share of that convertible land lying within a population-scaled reach of a nearby 1940 town, in the spirit of a market-potential function: larger towns project a wider reach than smaller towns. Because the entire measure is built from 1940 land cover and population, it is fixed before highway construction began and cannot itself be a consequence of the network it is used to help explain. This index captures a continuous measure of whether a county's agricultural land is likely to be under pressure from non-agricultural forces throughout the latter half of the twentieth century. Based on this measure, we differentiate average treatment effects by whether a county's farmland is under development pressure, uncovering heterogeneous effects across the pressure gradient.

Counties that gained IHS access saw farmland value per acre rise 11 percent on average over the post-treatment window, a figure that grows to approximately 20 percent forty years after construction. Farm counts rose but cropland declined moderately, and agricultural revenue per acre fell alongside the rising asset price. That pattern reflects the composition problem that aggregate farmland values cannot resolve: an interstate opening raises both agricultural returns and the outside-option value of land simultaneously (Floyd, 1965; Plantinga et al., 2002), and the aggregate farmland value approach (Donaldson and Hornbeck, 2016) capitalizes both without separating them. Identifying which channel dominated where requires tracing production outcomes separately across the land quality distribution.

We decompose this effect by investigating value-gains along the dimension of land suitability. The value-per-acre gains concentrate in the least-suitable land: marginal-class counties gain 16 log points and the moderately-suitable class 24, while the suitable class loses 13. These dimensions of heterogeneity demonstrate the central challenge in measuring welfare gains under extensive margin reallocation – suitable counties add marginal, lower-value parcels that pull their average down, while less-suitable counties shed their least-productive parcels and raise average land values. Farmland acreage moves oppositely, expanding in the suitable class while the less-suitable classes hold or contract.

We introduce three alternative specifications to address threats to identification and substantiate the interpretation of our primary estimate. First, IHS construction reshaped agriculture beyond the counties it directly served. We test whether market integration extends to counties within 75 kilometers of an opened interstate segment. We find that adjacent counties show effects of the same sign as directly served counties, attenuated in magnitude, confirming market integration spreads beyond the highway network. Second, staggered difference-in-differences designs recover an average treatment effect that grows with the post-treatment window as effects accumulate after adoption. To identify an effect that is invariant to the post-treatment window, we use a first-difference design to estimate the annual rate of change. Lastly, route placement and construction sequencing could be correlated with changes in the agricultural sector. To account for this, we estimate a first-difference IV specification, which instruments IHS highway access with the time-varying Pershing Map indicator of Frye (2026), whose location is fixed by a 1922 military planning document and whose timing is constructed from network betweenness centrality rather than fiscal or political sequencing. Results are robust and further illuminated across these alternative specification choices.

We complement the standard Donaldson and Hornbeck (2016) farmland-value welfare aggregate by bounding the share of total welfare that the agricultural-component decomposition omits using a Generalized Method of Moments (GMM) framework. We combine the cell-level treatment effects with a Fréchet productivity model and a pooled minimum-distance GMM that over-identifies eight structural primitives across three bins of development pressure, and quantify the omission the standard measure does not observe. The omission reaches as much as $195–210 billion across the two GAEZ surfaces (real 2025 dollars), concentrated on the high-pressure settled fringe. Hansen J-tests support the notion that the reduced form estimates are reconciled within the model. The magnitude of the bound rests on yield-based estimates of the parcel-level Fréchet tail parameter θ. We estimate that two values of these parameters are plausible for the era of U.S. agriculture, and they align with past work from Sotelo (2020) in Peru and Costinot et al. (2016) in modernized agriculture across the globe. Measured against the conventional aggregate's own gain, this implies the standard farmland-value measure captures roughly a fifth to a quarter of the total.

We complement past structural evidence regarding comparative-advantage-driven agricultural reallocation (Costinot et al., 2016; Sotelo, 2020; Pellegrina and Sotelo, 2025; Farrokhi and Pellegrina, 2023; Adamopoulos, 2025). However, the mechanism varies by context: migration (Pellegrina and Sotelo, 2025), productivity gains (Adamopoulos, 2025), and producer selection on heterogeneous land (Sotelo, 2020). Chan (2022) provides the only prior reduced-form test in the same US agricultural context, using identical FAO GAEZ suitability data and a market access design, and finds no comparative-advantage-driven specialization for 19th-century railroads. Output gains came through extensive-margin expansion on an unsettled frontier. On the settled frontier of the interstate era, by contrast, we show that market integration reorganizes existing land along comparative advantage. The Interstate Highway System introduced diverging patterns in agricultural welfare gains: a divergence in the value-per-acre response across the suitability distribution, rising on the least-suitable land and muted on the most-suitable, with the least-suitable land reorganizing toward livestock rather than exiting agriculture. That same integration raises the non-agricultural option value of land alongside agricultural returns, and uncovering the true aggregate welfare gains requires a model that rationalizes both.

We also contribute to a literature on the historical economic impacts of environmental shocks and interventions key to agricultural production, including infrastructure interventions (Hornbeck and Keskin, 2014; Hornbeck, 2010; Lewis and Severnini, 2020), climate change and weather shocks (Schlenker et al., 2005; Hornbeck, 2012; Deschênes and Greenstone, 2007), and public agricultural research investment (Kantor and Whalley, 2019). We show that the interstate highway system played an important role in reshaping where agricultural production occurred and forming the outside pressures that remain in modern agriculture. This tension between agricultural and non-agricultural land use is also addressed in the urban and spatial economics literature (Plantinga et al., 2002; Heblich et al., 2020; Fajgelbaum and Schaal, 2020; Nagy, 2022). We add to this literature by introducing a framework that better captures agricultural welfare gains in contexts where major extensive margin land reallocation is plausible.

Lastly, we contribute to a literature that measures the economic value of transit infrastructure in the US (Baum-Snow, 2007; Michaels, 2008; Chandra and Thompson, 2000; Frye, 2026). Michaels (2008) shows that highways raised the relative return to skill where skills were abundant and reduced it elsewhere, the labor market analog of the land quality heterogeneity we document; we provide the land-use analog, replacing the skill distribution with the suitability distribution. Prior work attributes the mid-century agricultural transformation, farm numbers falling by two-thirds, average farm size tripling, and agriculture's employment share dropping from 12 to 2 percent between 1950 and 2000, primarily to concurrent technological forces (Olmstead and Rhode, 2001; Hornbeck and Keskin, 2014; Gardner, 2002; Dimitri et al., 2005). Kim et al. (2025) document who left agriculture under the IHS, finding large individual gains in earnings and educational attainment as workers moved out of farming and into services; we document what happened to the land they left behind, studying county-level farm structure, production intensity, and land allocation rather than individual labor market outcomes. The suitability gradient in US agriculture widened with the IHS rollout, sorting existing farms by comparative advantage on a settled frontier rather than opening new agricultural land, establishing highway access as a consolidation channel alongside the technological forces that prior work emphasizes.

2 Background: Highways and Agricultural Production

2.1 The Interstate Highway System

Planning for a national highway network began with national-defense considerations long before the Federal-Aid Highway Act of 1956. General Pershing's 1922 military road survey, the earliest national plan of its kind, identified roughly 78,000 miles of strategic routes based on military priorities. Subsequent federal planning in the 1940s refined this into a designated system of approximately 41,000 miles that directly informed the 1956 Act, which authorized a national network of limited-access, high-speed roadways for long-distance freight and passenger movement (Swift, 2011; Rose, 1990).[1]

Construction timing, by contrast, was determined after designation: although over 90 percent of planned mileage was open to traffic by 1980 (Swift, 2011), the timing of when individual counties received segments varied substantially. The IHS was approximately 90 percent federally funded, but states retained latitude to prioritize the sequencing of work within their designated routes (Karnes, 2009; Rose, 1990). Initial post-1956 construction in cities proceeded rapidly under federal-state design standards before community opposition slowed urban segment completion in the late 1960s (Brinkman and Lin, 2024). More generally, the selection process tended to target counties with stronger prospective growth trajectories earlier among those that eventually received segments (Frye, 2026).

[1] We use the Pershing Map as a robustness check on route placement; see Section 5.

2.2 American Agriculture at Mid-Century

At mid-century, the typical rural county contained roughly 1,700 farms; approximately 30 percent of farmland was operated by tenant farmers rather than owner-operators; and the market reach of any given farm was constrained by the cost and speed of available transport, particularly for perishable products that required rapid movement to processors and consumers (Gardner, 2002; Dimitri et al., 2005).

Farm numbers fell by roughly 60 percent between 1950 and 2002: county-average counts declined from approximately 1,675 in 1950 to 650 by 2002, with the steepest decline concentrated between 1950 and 1975, the decades of peak IHS construction (Dimitri et al., 2005; Cochrane, 1979; Gardner, 2002; Olmstead and Rhode, 2001). Surviving farms grew larger: the share of operations exceeding 500 acres doubled from roughly 10 percent in 1940 to over 20 percent by 1960, and continued rising gradually through the end of the study period (Dimitri et al., 2005). Mechanization accelerated throughout: tractors per farm rose more than sixfold between 1940 and 1990 as animal power gave way entirely to machinery (Olmstead and Rhode, 2001). Land classified as farmland contracted modestly from its early-1950s peak, while cultivated cropland acreage remained stable; the reduction in farm acreage was concentrated in pasture and non-harvested land rather than in the cropped area itself (Gardner, 2002). Tenancy fell sharply over the same period, from roughly 30 percent of farmland operated by tenants in 1950 to under 10 percent by 1970, as mechanization and consolidation shifted the economics of farm organization toward owner-operation (Gardner, 2002; Dimitri et al., 2005).

Trucking displaced rail as the primary mode of agricultural freight movement over the decades when the IHS was built; by the late twentieth century, trucks carried the majority of the nation's food supply by value, with the interstate network enabling regional distribution hubs, national commodity markets, and just-in-time input delivery that extended the effective market reach of individual farms well beyond what pre-highway transport had permitted (McCoy et al., 2020).

Agricultural product market value per acre roughly tripled in real terms between 1940 and 2002; farmland values per acre rose by roughly five times over the same period. The two series co-moved through the late 1960s, then diverged: farmland values rose sharply through the 1970s commodity boom, corrected during the 1980s farm financial crisis (Gardner, 2002; Melichar, 1979), and rose again through the 1990s. By 2002, farmland values carried a wider premium over agricultural production revenue than had characterized mid-century agriculture. Farmland values rose faster than agricultural rents because farmland prices capitalize non-agricultural options alongside current farm income: development potential from expanding urban areas, speculative demand linked to infrastructure access, and rents from federal commodity programs (Plantinga et al., 2002; Floyd, 1965).

3 Conceptual Model

We study land allocation between competing agricultural and non-agricultural uses in a Ricardian model of comparative advantage on heterogeneous land, following Costinot et al. (2016) and Sotelo (2020). A highway opening does several things to farmland at once. It moves the farm-gate return through two competing forces – local producers can now reach distant buyers, which raises the price, while outside sellers can now ship in, which lowers it – and it moves per-acre costs, as cheaper shipped-in inputs pull against dearer land and labor. Neither net direction is something the model fixes; both are questions the data answer. The highway also raises the value of land in non-agricultural use where development pressure binds, concentrated in the conversion fringe of growing cities. The framework works from a single index of land quality but two farm technologies – crops, the intensive use, and livestock, the extensive one – and two non-agricultural outside options a parcel can take once farming no longer pays: development where a builder bids, and a positive non-development reversion where none does. The model sorts the cross-suitability pattern into two pieces with different drivers. The *survival margin*, which land stays in farming, is set by the soil: poor land is held in agriculture by livestock, the low-cost extensive margin, and this pattern shows up even in counties with no development pressure at all. The *capitalization margin*, the gap between what farmland sells for and what it earns from farming, is set by development pressure, which raises the outside option that shows up in the land's price but not in its output. The model delivers four testable predictions about how agricultural production, farmland values, and adjacent counties reorganize, together with a testable implication for within-agriculture composition. Complete derivations and Fréchet closed forms are in Appendix B.

3.1 Environment

Counties are indexed by c, land parcels by j, and agricultural census years by t. Following Costinot et al. (2016), each parcel's land quality is

$$\phi_{jc} = A_c \cdot \omega_{jc}, \tag{1}$$

where $A_c > 0$ is county-level suitability (soil, climate, and comparative advantage in crop production) and $\omega_{jc} > 0$ is a parcel-level idiosyncratic draw (microclimate, drainage, slope). One quality index serves any use; crops and livestock differ not in the soil they read but in how steeply they reward it. Following Eaton and Kortum (2002) as applied by Sotelo (2020), ω_{jc} is distributed Fréchet with shape parameter $\theta > 1$:

$$F(\omega) = \exp\left(-\omega^{-\theta}\right), \qquad \omega > 0, \tag{2}$$

where a lower θ means a county's parcels differ more in quality, so more of them sit near the margin and a given change in returns moves more land in or out. County suitability A_c is the scale of that distribution: it lifts every parcel proportionally, good and bad alike.

Market access enters through a transport-cost wedge $\tau_{ct} \geq 0$; an interstate opening *lowers* τ_{ct}, so $\Delta\tau < 0$ is an access improvement and "access raises X" means $\partial X/\partial\tau < 0$. Adjacent counties receive an attenuated improvement $\lambda\,\Delta\tau$ with $\lambda \in (0,1)$ on the transport channel; the local development bid is a direct-county phenomenon and does not spill (Prediction 3 below).

3.2 Uses, Returns, and the Parcel's Choice

A parcel goes to its best use. The two farm uses are crops and livestock,

$$R^C = b_C(\tau_{ct})\,\phi_{jc} - c_C(\tau_{ct}), \qquad R^L = b_L(\tau_{ct})\,\phi_{jc} - c_L(\tau_{ct}), \tag{3}$$

crops the intensive use, whose return rises steeply with quality at a high per-acre cost, and livestock the extensive use, whose return is flatter and whose cost is lower: $b_C > b_L$ and $c_C > c_L$ (Assumption A3). The non-farm options are two, both flat in the parcel draw ω_{jc}, since what a parcel is worth as a house lot or in conservation does not turn on its idiosyncratic soil. A parcel in the conversion fringe faces the developed-use bid $R^O_{ct} - \kappa_{ct}$, where $\kappa_{ct} \geq 0$ is the conversion wedge (zoning, transaction, and physical costs of building); a parcel beyond the fringe faces a non-development reversion R^N_{ct} available everywhere – low-intensity use under the conservation and idle-land regime of the period (Soil Bank, unprogrammed reversion, and from 1985 the Conservation Reserve Program) – with $R^O > R^N \geq 0$. Returns are measured net of the land's amenity and holding reservation, so the choice the parcel makes is among its farm and outside uses, with the reversion wedge κ_N suppressed in the reversion term R^N_{ct} throughout (footnote 4):[2]

$$V_{jc} = \max\{\, R^C,\ R^L,\ (R^O_{ct} - \kappa_{ct})\,\mathbf{1}\{\text{fringe}\},\ R^N_{ct}\,\}\,. \tag{4}$$

[2] An idle-land amenity floor lies beneath the named uses: a parcel too poor even to revert is held at its (normalized) holding reservation. This floor is completeness scaffolding – it is never the binding option in the data and is not tested – so it is suppressed in the value-max (4) and carried only as the implicit lower envelope of V_{jc}.

A parcel farms when a farm use beats its best non-farm option; leaving farming, it develops where a builder bids and reverts to the non-development use R^N where none does, with the reservation binding only for land too poor even to revert. Leaving farming is therefore not worth zero: non-fringe exit yields the positive reversion value R^N, fringe exit the higher development value R^O.

The two farm returns rise with quality, crops faster, while the non-farm options are flat in it, so the parcel's use is pinned by quality against two cutoffs. Crops beat livestock above

$$\phi_{CL} = \frac{c_C - c_L}{b_C - b_L}, \tag{5}$$

a comparison of the two farm technologies alone. Writing ρ_{ct} for the parcel's best non-farm option ($R^O_{ct} - \kappa_{ct}$ in the fringe, $R^N_{ct} - \kappa_N$ beyond it), livestock – the marginal farm use – clears its threshold above

$$\phi^*_{ct} = \frac{\rho_{ct} + c_L}{b_L(\tau_{ct})}, \tag{6}$$

the quality at which grazing beats the parcel's best non-farm use. When $\phi^*_{ct} < \phi_{CL}$ the parcel uses sort cleanly: crops on $\phi \geq \phi_{CL}$, livestock on $\phi^*_{ct} \leq \phi < \phi_{CL}$, exit below ϕ^*_{ct} – to development in the fringe and to the non-development use R^N beyond it. Best land crops, middle grazes, worst leaves farming: one soil-ordered response to a single rising bar $R = R^O - \kappa$, which meets the best land where it intensifies and capitalizes, the middle land where it idles and waits, and the marginal land where grazing retains it.

Two cutoffs answer two questions through two forces. The crop–livestock cutoff ϕ_{CL} depends only on the farm technologies, so it sorts the farm *activity* by soil. The farm–exit cutoff ϕ^*_{ct} and the develop-versus-revert split turn on the outside option ρ_{ct}, set by proximity to development rather than by soil, so they sort the *exit destination* by pressure. Soil decides what a parcel farms; development pressure decides what becomes of it if it stops. County farmland share keeps its single-threshold form,

$$S_{ct}^{A} = 1 - F(\omega_{ct}^{*}) = 1 - \exp\bigl(-(\omega_{ct}^{*})^{-\theta}\bigr)\,, \qquad \omega_{ct}^{*} = \phi_{ct}^{*}/A_c, \tag{7}$$

with the threshold set by the livestock technology against the local non-farm option, so low-suitability land stays in farming through grazing rather than reverting. The crop share of farmland takes the same form at ϕ_{CL}/A_c: high-suitability counties carry more parcels above ϕ_{CL} and crop-dominate, low ones graze. Appendix B.16 gives the closed forms and the coexistence condition $\phi^* < \phi_{CL}$.

Two development-pressure regimes follow, divided by a pivot $\bar{\rho} \equiv b_L\,\phi_{CL} - c_L$ in the outside option – the level at which the farm–exit cutoff reaches the crop–livestock cutoff ϕ_{CL}. Below the pivot (Regime I, low-to-moderate pressure) the crop-exit cutoff lies strictly inside the retained farm range, so crops never bind their own exit margin and the livestock band is the only shed farm margin; this nests the single-margin sort exactly. Above the pivot (Regime II, high pressure) the pure grazing band is gone and the binding cutoff climbs through the crop band from below, developing the lowest-value crop parcels near ϕ_{CL}, so displacement is no longer a grazing-only phenomenon. Which counties sit in which regime is set by where the developer bar is highest, hence by where M_c is highest; the maintained correlation $\mathrm{Corr}(A_c, M_c) > 0$ (Assumption A2) then places the pressured fringe disproportionately on good and middle soil – the crop-acreage margin concentrating on middle soil and the best soil capitalizing, developed below. Appendix B.4 derives the pivot and the cutoff ordering; Appendix B.16 gives the Fréchet shares for both regimes.

3.3 The Conversion Frontier and the Exit Destination

Development pressure falls on parcels, not on whole counties, and on all developable land near a growing city, not on farmland alone. The monocentric land-rent tradition (Alonso, 1964; Mills, 1967; Muth, 1969), extended to land conversion at the growing urban fringe by Capozza and Helsley (1989, 1990), locates that pressure in a band around the built-up edge – land close enough that the developed-use bid, net of conversion cost, approaches the agricultural bid, but not yet built. The county moderator M_c is the share of a county's developable land lying in this pre-interstate fringe; because the fringe is fixed before the interstate, M_c is predetermined with respect to it. We measure M_c directly rather than from an assumed bid schedule; Section 4.1 gives the construction.

Fringe and non-fringe parcels share the same soil and the same quality draws but face a different outside option.[3] A fraction M_c of developable parcels lie in the fringe, where a builder bids $R_{ct}^{O,\text{press}} - \kappa_{ct} > 0$; the remaining $1 - M_c$ lie beyond it, where no builder bids and land leaving farming reverts to the positive non-development use $R_{ct}^N - \kappa_N > 0$, not to idleness. Soil A_c is common across the county, and the draw ω_{jc} is independent of fringe membership conditional on A_c, since M_c is set by predetermined proximity to pre-interstate cities rather than by current quality. The two groups therefore differ in one thing only, the outside option, set by proximity and not by soil (Assumption A2, which states the cross-county suitability–pressure correlation as a maintained historical fact and this within-county independence as the invariance the sort requires; any residual fringe–quality correlation would tilt the develop margin toward better parcels, reinforcing rather than reversing the results). The crop–livestock cutoff ϕ_{CL} is common to both; the farm–exit cutoff differs, because the fringe's marginal farm use must clear a builder's bid while the non-fringe's clears only the reversion value:

$$\phi_{\text{press},ct}^* = \frac{R_{ct}^{O,\text{press}} - \kappa_{ct} + c_L}{b_L} > \phi_{\text{none},ct}^* = \frac{R_{ct}^N - \kappa_N + c_L}{b_L}, \tag{8}$$

since $R^O - \kappa > R^N - \kappa_N$. The fringe holds less land in farming and develops what it sheds; the non-fringe holds more and reverts what it sheds. Because ϕ_{CL} is common

[3]The model defines S^A and M_c over developable (still-undeveloped) parcels: already-built parcels have exited agriculture and enter neither. Empirically S^A is measured as farmland over total county land area, a broader base than the developable universe, which is not separately observed.

while the fringe's exit cutoff is higher, the fringe sheds its marginal band first, to development. Under low-to-moderate pressure (Regime I) that band is livestock, and the grazing that survives concentrates beyond the fringe, where the counterfactual is reversion; under high pressure (Regime II) the same rising bar pushes the exit cutoff through the crop band from below, so the fringe also sheds its lowest-value *crop* parcels into development. Where each margin falls on the soil ladder follows from where the bar binds. The extensive grazing band is shed across the pressured fringe, which the maintained correlation $\mathrm{Corr}(A_c, M_c) > 0$ (Assumption A2) places on good and middle soil alike. The crop-acreage margin is different: the crop-exit cutoff ϕ_C^* prices out a parcel only once the bar clears its crop return $b_C A_c \omega$, largest on the best soil, where two forces oppose – the highest bar (that same correlation puts the best soil under the most pressure) pushes toward conversion while the high crop return resists it. Over the pressure range the data span the crop-return force dominates on the best soil, so those parcels are the *last* to convert: crop-acreage displacement concentrates on *middle* soil (the surface the estimates take as primary), while on the best soil the development option capitalizes into value rather than conversion (Prediction 4). The displaced band is a modeled within-county conversion – the crop-exit cutoff truncating an occupied crop band on middle soil, the extensive band shed more broadly – selected onto the pressured soil by the maintained correlation, not an artifact of averaging across counties. This is the three-state landing: retention on the far-from-city marginal land of Prediction 1, crop displacement and parking on middle soil, and capitalization on the best soil.

The county agricultural land share is the area-weighted mixture of the two subgroups,

$$S_{ct}^A = M_c\, S^A\big(\omega_{\mathrm{press},ct}^*\big) + (1 - M_c)\, S^A\big(\omega_{\mathrm{none},ct}^*\big)\,, \qquad \omega_\sigma^* = \phi_\sigma^*/A_c, \tag{9}$$

nesting the single-group share. Since M_c is set before the highway arrives, comparing the mixture before and after access leaves a county effect that rises, to first order, in step with

M_c: the effect a county would see with no land under pressure, plus an increment for every additional slice of the county that is. Appendix B.17 works out this dose-response and its curvature; the destination split – developed value R^O in the fringe, the lower reversion value R^N on reverted land beyond it – is what the welfare accounting turns on (Section 7).

A highway also opens a farm-entry margin, the mirror of exit on the reversion side. Where access raises a farm-gate return but no builder bids ($M_c = 0$), reverted or idle land re-enters farming above the entry cutoff $\omega^*_{\text{none}} = (R^N - \kappa_N + c_L)/(b_L A_c)$, and the same reach that drives Prediction 1 pushes that cutoff down. Two gross flows therefore run under one highway: an exit-to-development flow Ω_{dev} on the pressured fringe and an entry flow Ω_{ent} on marginal, unpressured land. The county farmland share of equation (9) nets them, and the difference-in-differences measures that net object, while the welfare accounting prices the *gross* conversion Ω_{dev} alone (Section 7). On the conversion cells entry and reversion partly offset a larger gross exit, so the net acreage response *understates* gross conversion and any omission computed off net acreage is biased toward zero; the county-net also mixes a rising-farmland retention cell with a shedding conversion cell, so the welfare object is built on the cell-resolved gross flow rather than the county-net response (Appendix B.4).

3.4 Assumptions and Predictions

Assumption A1 (Access incidence on farm returns: prices and costs). A highway moves each farm-gate return $b_k(\tau_{ct})$ through reach, which raises the price local producers fetch, and competition, which lowers it; the net sign of $\partial b_k/\partial\tau$ is left to the data, separately for each activity (Donaldson and Hornbeck, 2016). The one maintained restriction is within-activity uniformity: the access-induced shift in b_k is common across a county's parcels and does not fall harder on good land than on bad, so a change in b_k moves the exit threshold cleanly rather than reshuffling which parcels farm. The model imposes no ordering on how access affects crops versus livestock; $\partial b_C/\partial\tau$ and $\partial b_L/\partial\tau$ may differ, in either direction. Costs respond too: a highway moves the per-acre cost $c_k(\tau_{ct})$, cheaper shipped-in inputs pushing it down while competition for land and labor from non-farm uses pushes it up, with the net sign again left to the data. The cost response is common across activities, so the crop–livestock cost gap $c_C - c_L$ is access-invariant even as the levels move. The composition cutoff $\phi_{CL} = (c_C - c_L)/(b_C - b_L)$ may still shift, since A1 leaves the price legs unordered, but any such shift loads on good land rather than marginal land, so the model predicts no marginal-concentrated commodity-crop expansion; Appendix B.2 gives the density argument and the sense in which this refusal to order is not vacuous. The farm-versus-exit threshold ϕ^* responds to the net of the price and cost legs.

Assumption A2 (Proximity-driven outside option, two regimes). The outside option a parcel takes on leaving farming has two regimes set by proximity to development. In the conversion fringe a parcel faces the bid $R_{ct}^{O,\text{press}} - \kappa_{ct}$, read as the capitalized development *option* – land built, or held vacant in anticipation of development, since speculation needs no building and is valued at the option – and this bid rises with access as a highway pulls the built-up edge outward (Capozza and Helsley, 1989, 1990). Beyond the fringe no builder bids, and land that leaves farming reverts to the non-development use $R_{ct}^N \geq 0$, a period-average lower bound on the reversion value.[4] The retention result (Prediction 1) requires only $R^N \geq 0$: the force holding marginal land in farming is the livestock farm-gate response $\partial b_L/\partial\tau$, which operates even as R^N approaches the floor. The idle-land amenity floor (genuinely idle land) is the implicit lower envelope underneath the named uses, kept distinct from the development option so that netting out the holding value does not absorb the pressure capitalization.

[4]The reversion enters the value-max (4) as R^N rather than $R^N - \kappa_N$: conservation enrollment is typically subsidized, so the reversion wedge $\kappa_N \approx 0$ and is carried only for symmetry, and R^N is the net program or idle payment. As a period average R^N carries a conservation premium at the program endpoints (the Soil Bank of 1956–1960 and the Conservation Reserve Program from 1985) and falls toward zero in the interprogram years, when exiting marginal land idled or shifted to within-agriculture grazing rather than drawing a payment; Prediction 1 uses only the lower bound $R^N \geq 0$. Commercial forestry is left out of R^N: timber is soil-sensitive and competes for good sites, so it is not part of this flat outside use.

Cities in the study region were historically sited on valley bottoms and good, flat soil – river crossings, port sites, and rail junctions on productive land – so county suitability A_c and the fringe share M_c are positively correlated, $\mathrm{Corr}(A_c, M_c) > 0$, an externally verifiable institutional fact visible in the data: the high-development-pressure share rises monotonically across the suitability classes, from 7.9% on the marginal class to 16.9% and 22.3% on the two suitable classes. A2 maintains this correlation as a historical fact, together with the two structural invariances the sort needs – within-county quality draws independent of fringe membership given A_c, and a bid response $\partial R^O/\partial\tau$ common across soil within a pressure stratum. The correlation is meaningful but not primitive: the displacement *mechanism* is native to the two-regime structure, while *where* it falls follows from the correlation placing the high-M_c mass on good and middle soil, so a weaker correlation would leave the displacement intact but attenuate its good-soil concentration (Appendix B.6). A2 makes no claim about whether the bid tilts on soil within a location; whether a developer's willingness to pay favors good cropland is an empirical question the pressure-conditioned evidence answers directly (Appendix B.6).

Assumption A3 (Comparative advantage as a technology ordering, the activity sort against the land-rent bar, and input-for-land substitution driven by the land price). Crops and livestock are two technologies on the same land quality, ordered by how steeply they reward it: crop returns rise more steeply with quality and carry a higher per-acre cost than livestock, $b_C > b_L$ and $c_C > c_L$. A crop parcel combines land services – acres of quality ϕ, priced at the land-rent bar $R = R^O - \kappa$ – with a non-land input bundle x through a constant-returns CES aggregator of substitution elasticity $\sigma > 0$, so the optimal input–land ratio rises in the land price, $\partial \ln(x^*/\ell^*)/\partial \ln R = \sigma > 0$. This CES elasticity σ is a single scalar technology parameter; the same symbol appears in the welfare appendix as a superscript indexing the development-pressure stratum ($\sigma \in \{\mathrm{None}, \mathrm{Low}, \mathrm{High}\}$), a distinct labeling use that is never an exponent.

The slope ordering is agronomic and the cost ordering is what lets livestock exist: were crops both steeper and cheaper they would dominate every parcel and none would graze. The two orderings place a single crossing at ϕ_{CL}, with extensive grazing holding the low-cost margin below it, so marginal-land retention is livestock holding that margin where crops cannot pay. The implied allocation is testable: the crop share of farmland rises across the suitability classes from 18% on marginal land to 38% and 60% on the two suitable classes, and is unchanged among the roughly two-thirds of counties facing no development pressure, so it is comparative advantage rather than a pressure or feasibility artifact. Appendix B.16 gives the closed form, the coexistence condition, and the gradient's provenance.

Comparative advantage fixes not just which use occupies each parcel but its *position* in the quality distribution, and that placement is what the rising land-rent bar of A2 needs: grazing occupies the marginal band, so a rising development bid climbs through that band from below and prices grazing out first. Acreage falls, extensive revenue falls, and the surviving mix tilts toward input-intensive cropping; because the bar climbs with pressure, this input rise scales with M_c and is flat-to-negative at $M_c = 0$ – the tested content of Prediction 2 – and runs through R^O, not access. The bar does not reach land-light confinement, which A3 carves out of the flat-return grazing technology, so the mechanism sheds extensive grazing, not confinement (Appendix B.8).

The three assumptions determine how a highway moves the exit threshold. Within a pressure stratum σ, differentiating $\omega^*_\sigma = (\rho^\sigma + c_L)/(b_L A_c)$ with respect to τ, holding A_c fixed, gives

$$\frac{\partial \omega^*_{\sigma,ct}}{\partial \tau} = \frac{1}{b_L(\tau)\, A_c}\, B^\sigma, \qquad B^\sigma = \frac{\partial R^{O,\sigma}_{ct}}{\partial \tau} + \frac{\partial c_L}{\partial \tau} - \phi^*_\sigma \frac{\partial b_L}{\partial \tau}, \qquad \frac{\partial B^\sigma}{\partial A_c} = 0. \tag{10}$$

The $1/A_c$ in (10) makes the threshold move more on poor land, but it cannot flip the direction of that move: the sign depends entirely on a bracket B^σ with no A_c in it. Whether that bracket is common across soil within a pressure group is what separates the two margins.

On the capitalization margin the loading comes from the outside option, which differs across pressure strata but shows no detectable soil tilt within a stratum once exposure M_c is held fixed (Appendix B.6); on the survival margin the retained activity is itself soil-ordered, so the cross-suitability survival gradient is comparative advantage, not pressure (Appendix B.4).

Proposition 1 (Capitalization incidence) *Fix a development-pressure stratum $\sigma \in \{\text{none}, \text{press}\}$ and an interior solution ($\omega^*_{ct} \in (0, \infty)$, $f(\omega^*_{ct}) > 0$). The value–revenue capitalization wedge $\Delta \ln(\textit{value/acre}) - \Delta \ln(\textit{revenue/acre})$ is governed by development pressure. Farmland value per acre loads on the development option $R^{O,\sigma}$ through composition – which parcels survive and their average quality – and through direct capitalization – near-margin survivors in the fringe are valued near $R^{O,\sigma}$; agricultural revenue per acre loads on $R^{O,\sigma}$ through composition alone, since a parcel that stays in farming sells the same output regardless of its option value. The wedge therefore isolates the capitalization of the option value, and the cutoff response factors as in* (10) *with $\partial B^\sigma / \partial A_c = 0$. Conditioning on M_c orders the wedge – larger on the pressured subgroup, where the development bid binds – with suitability entering only as a secondary amplifier through the revenue base $b_L(\tau)A_c$.*

Proposition 1 ties the capitalization gap to pressure, not to soil: the wedge is present where a builder bids and small beyond the fringe, where the outside option is the lower reversion value $R^N \ll R^O$ and value and revenue move nearly together. Soil scales the magnitude through the revenue base but does not set the ranking, so the wedge attenuates once pressure is controlled – the falsifiable content of test N1 (Appendix B.5, Appendix B.8).

Lemma 1 (Survival by comparative advantage) *Under Assumption A3, which land stays in farming is set by the marginal-use (livestock) breakeven. Beyond the fringe this is the farm-versus-reversion margin, $\omega^*_{\text{none},ct} = (R^N_{ct} - \kappa_N + c_L)/(b_L A_c)$, present at $M_c = 0$ and independent of development pressure. Because comparative advantage makes low-suitability counties livestock-*

dominated, retention of marginal land is retention of livestock; the survival gradient is ordered by soil through composition, not by pressure, and is not collapsed by conditioning on M_c.

Lemma 1 keeps the survival margin apart from the capitalization margin. Survival is ordered by soil and runs against a reversion counterfactual; capitalization is ordered by pressure and runs against a development counterfactual. Survival is held by livestock's low-cost extensive margin: where reach raises the livestock farm-gate b_L (A1), ω^*_{none} falls and marginal land that would otherwise revert stays in grazing (Appendix B.8).

Under A1–A3 an access improvement reorganizes the agricultural base along two axes: the survival margin runs on soil-linked comparative advantage (Lemma 1), and the capitalization margin runs on development pressure (Proposition 1). Appendix Table A23 summarizes the predicted directions by development-pressure stratum. Table 1 states the four predictions the model delivers and what would falsify each.

A testable implication (corollary of Prediction 1 and A3): the marginal extensive-grazing margin has two destinations, sorted by pressure. Comparative advantage makes marginal counties livestock-dominated, so the marginal land the model speaks to is extensive-grazing land. Where there is no development pressure ($M_c = 0$), access *retains* this land in grazing: ω^* falls, the marginal band is held rather than reverting, and pastureland and beef-cattle inventory rise on the retained base. Where pressure binds, the *same* marginal grazing land is the first use the rising land-rent bar prices out. The confinement carve-out separates the two poles: the bar sheds extensive grazing while land-light confinement is untouched to expanding, so the A3 carve-out is shown rather than posited. At the prime end the conversion pull is strongest, but holding continuous exposure M_c fixed collapses the within-location good-land tilt (Appendix B.9, Appendix B.6).

The split rests on three empirical conditions: (C1) development pressure falls on good land – coarsely, a marginal-versus-suitable contrast that is flat within the suitable classes; (C2) the capitalization response is common across suitability within a pressure stratum, so M_c orders the wedge; and (C3) a net farm-gate slope sign on the no-pressure subgroup,

Table 1: The four testable predictions and what would falsify each

	Prediction and mechanism	Falsified if
P1	**Marginal land is retained in farming, against a reversion counterfactual.** Livestock holds marginal land where crops cannot pay; a highway that raises the livestock farm-gate ($\partial b_L/\partial\tau < 0$) retains land that would otherwise revert. The pattern is soil-ordered and present at $M_c = 0$, because the counterfactual is reversion, not development.	Marginal retention vanishes at $M_c = 0$, or collapses under M_c-conditioning.
P2	**Where pressure binds, acreage and revenue per acre fall as the land-rent bar prices out extensive grazing; value per acre is flat.** The rising bar R^O enters the farm–exit cutoff ϕ^*, not the technology-only ϕ_{CL}, so it climbs through the grazing band from below: acreage contracts and livestock revenue per acre collapses, while surviving crop parcels substitute non-land inputs for dearer land, so input use rises with M_c. Value per acre is unsigned – the sum of this revenue decline and the P4 wedge – so its flatness is a joint over-identifying restriction.	Acreage does not fall under pressure; revenue per acre *rises* with pressure; input use rises under access alone rather than scaling with pressure; or the displaced animal margin is confinement rather than extensive grazing.
P3	**Adjacent counties show attenuated, same-direction effects, by channel.** Neighbors receive $\lambda\,\Delta\tau$ with $\lambda < 1$ on the transport channel, so revenue and input margins propagate. The capitalization channel is local – R^O is a direct-county phenomenon – so it does not spill.	A single common λ across the revenue and value effects, rather than channel-specific pass-through.
P4	**Farmland value and agricultural revenue diverge where the outside option capitalizes.** Value per acre records the parcel's best use including the development option; revenue per acre records current output only. The gap $\Delta\ln(\text{value/acre}) - \Delta\ln(\text{revenue/acre})$ therefore isolates the capitalized non-farm option (Floyd, 1965; Plantinga et al., 2002). It is pressure-first: large in the fringe, small beyond it where $R^N \ll R^O$, with soil a secondary amplifier only.	The wedge is not ordered by M_c, or fails to attenuate once pressure is controlled (falsification test N1).

Notes. P2 is identified by its scaling with pressure, which separates it from a selection story (revenue per acre *up*, no input rise) and from an access story (input use rising at $M_c = 0$). P4 is distinct from P2, not a relabeling: they differ in comparison object – acreage and revenue per acre versus the value−revenue gap – and bind in the same place. On a no-pressure marginal county neither pressure statement operates: P1 holds while the P4 wedge is near zero and the P2 contraction is absent. Appendix B.8 and Appendix B.9 derive P2, the rejected alternatives, and the discriminating tests; Appendix B.16 treats the held, not-yet-converted band on middle land, whose acreage-versus-revenue wedge resolves P2 one layer finer; Appendix B.5 states N1–N4 formally.

which the data measure. The falsification tests N1–N4 and the survival corollary are stated formally in Appendix B.5. Section 6 maps every assumption, the proposition, the lemma, the four predictions, the livestock implication, and the falsification tests to their evidence and verdict.

Appendix B provides Fréchet closed forms for the agricultural land share and welfare decomposition and maps model objects to Agricultural Census outcomes; Appendix B.5 states the falsification tests N1, N3, and N4 as corollaries of Proposition 1, with the acreage-margin test N2 developed in Appendix B.9, and the survival corollary of Lemma 1.

4 Data

Agricultural Data We assembled county-level data on US agriculture from 1940–2002 available from Haines et al. (2018). These data come from the US agricultural census, which is collected roughly (but not exactly) every five years. They provide a comprehensive record of agricultural outcomes at 14 points in time before, during, and after the era of the rollout of the Interstate Highway System. The variables of interest span three categories of economic outcomes.

First, agricultural land values measure the potential value of production on a per-acre basis. Agricultural land values capitalize expected future returns to agricultural production (Mendelsohn et al., 1994). Similar to Donaldson and Hornbeck (2016), this measure serves as our summary measure of how highway access affected the overall expected productivity of local agriculture.

Second, we extract six variables describing the local agricultural economy's structure, including the number of farms, the number of farms over 500 acres, farmland acres, cropland acres, pastureland acres, and the cows and calves inventory. These variables allow us to assess whether interstate highways facilitated farm consolidation, altered the extensive margin of land use, and/or shifted the composition of agricultural activity between cropping and livestock.

Third, input choice and crop production outcomes are captured by farm labor ($ per farm), fertilizer use ($ per acre), share of cropland irrigated, and the county's four-crop concentration ratio (i.e., share of agricultural land devoted to the county's top four crops). These outcomes speak to whether improved market access induced changes in input intensity, technology adoption, or the degree of crop specialization.

Crop Suitability To classify a county's land into suitability for crop production, we extract gridded suitability indices from the Food and Agriculture Organization of the United Nations' Global Agro-Ecological Zones (GAEZ v.4).[5] These data provide an assessment of a location's climatic, natural resource, and soil endowments, paired with agronomic crop growth models for different crop types. They comprise a 9 km by 9 km grid across the globe, measuring each cell's attainable yield for a crop. These data are commonly used as the basis for determining a location's comparative advantage or natural endowment for growing specific crops (Costinot and Donaldson, 2012; Bustos et al., 2016; Costinot et al., 2016; Adamopoulos and Restuccia, 2022).

In this study, we focus less on how Interstate Highways impacted specific choices between crops within the US, and instead use these data to characterize US counties by their suitability for the composite of leading US crops. We aggregate the rainfed low-input surface to the multi-crop level by taking the pixel-wise mean across the SUIT-9 commodity bundle.[6] Each pixel's mean suitability score is classified into the four stan-

[5]In particular, we select the rainfed, low-input suitability indices over the period of 1961–1990 for the baseline specifications. The alternative options – high-input and irrigated suitability indices – incorporate returns to mechanization, fertilizer application, improved varieties, and irrigation infrastructure – factors that are themselves influenced by market access and thus potentially endogenous to highway treatment. We treat rainfed, low-input as our baseline to achieve a closer measure of a place's natural endowment, similar to Adamopoulos and Restuccia (2022). We test the sensitivity of our results to this choice in Appendix B.14.

[6]The SUIT-9 bundle is corn, wheat, soybean, sorghum, barley, oats, cotton, Irish potato, and rice. Mean aggregation, rather than the pixel-wise maximum across crops, is the specification: the cross-crop mean is structurally closer to the relevant scale parameter A_c in the model than the maximum, which over-weights the best-suited crop and discards the diversification margin.

dard FAO classes on the 0–$10{,}000$ scale – N (not suitable, score < 2500), S3 (marginally suitable, 2500–5000), S2 (moderately suitable, 5000–7500), and S1 (highly suitable, ≥ 7500) – and a county's suitability is the vector of its area shares in the four classes, which sum to one. The four-share parametrization preserves the cross-county composition heterogeneity that a single dominant-class indicator would discard. The rainfed low-input surface is the primary specification, and the surface used for the θ estimation; the rainfed high-input and irrigated surfaces are alternative input-intensity parameterizations carried as robustness, with the irrigated surface collapsing the marginal N class to a thin residual not used to read the marginal-land margin.

4.1 Measuring Development Pressure

The development-pressure moderator M_c is a predetermined index of each county's exposure to urban development pressure: the share of its convertible 1940 land lying within reach of a nearby town. It is built in three steps – locate the 1940 towns, extend a population-scaled reach around each, and take the share of a county's convertible land those reaches cover. Appendix B.18 gives the construction in full; the essentials follow.

Towns and their reach. A town is a 1940 incorporated census place of at least $25{,}000$ residents, taken from the complete full-count census-place layer (Appendix Figure A12). Each town is assigned a circular reach whose radius grows with its population, $r_j = \min\{k\, P_j^{1/3},\, \bar{r}\}$: the cube-root scaling is urban allometry – built-up area scales with population to the two-thirds power, so the radius scales to the one-third – and k is set so a city of one million residents reaches sixty kilometers, with the cap $\bar{r}$ binding for almost no 1940 city. The reach widens with city size, since a larger city's development frontier extends farther, but sub-linearly, so the largest cities do not dominate the surface. It is read from a town's census population, never from any post-1940 information.

Convertible land. The denominator nets developed land and water out of each county. A $250\,\mathrm{m}$ cell is convertible when its 1940 built-up intensity is below $1{,}000\,\mathrm{m}^2$ of gross indoor floor area and it is not open water. The intensity threshold de-dusts the built-up record: HISDAC-US built-up intensity counts every structure type, including agricultural buildings (Leyk and Uhl, 2018; McShane et al., 2022), so a presence rule reads dispersed farmsteads as developed, most severely in densely settled, well-recorded farm regions; the threshold is calibrated non-circularly, as the level that best separates incorporated census places of at least $2{,}500$ residents from rural background. Because the built-up record under-records towns in some regions, a coverage-gap town's own developed footprint is *imputed* from the population-and-region relationship fit on the towns the record does cover, so its built core is netted out of the convertible denominator like any other; the match uses population and region only, so it carries no market-access information (Appendix B.18).

The measure. The county index is the share of a county's convertible land that falls within at least one town's reach,

$$M_c = \frac{\displaystyle\int_c \mathbf{1}\{x \in \textstyle\bigcup_j \mathrm{reach}_j\}\,\mathbf{1}\{x \text{ convertible}\}\,dx}{\displaystyle\int_c \mathbf{1}\{x \text{ convertible}\}\,dx} \in [0,1]. \tag{11}$$

Reaches from neighboring towns are unioned, so overlap is counted once, and developed and water cells carry no weight in either the numerator or the denominator. Netting built and water land out makes M_c the share of a county's *remaining convertible* land under pressure, not the share of its total area. Both the towns and their reach are read from 1940-and-earlier geography, so M_c cannot respond to the interstate. We report M_c at four town-population cutoffs – $2{,}500$, $10{,}000$, $25{,}000$, and $50{,}000$ – with the $25{,}000$-resident threshold as the baseline and the other three as a robustness family. Figure 1 shows the national surface: pressure concentrates in the Northeast megalopolis and the Great Lakes belt, with the agricultural interior near zero – a concentrated pattern rather than a smooth continental gradient, which is what separates M_c from market access.

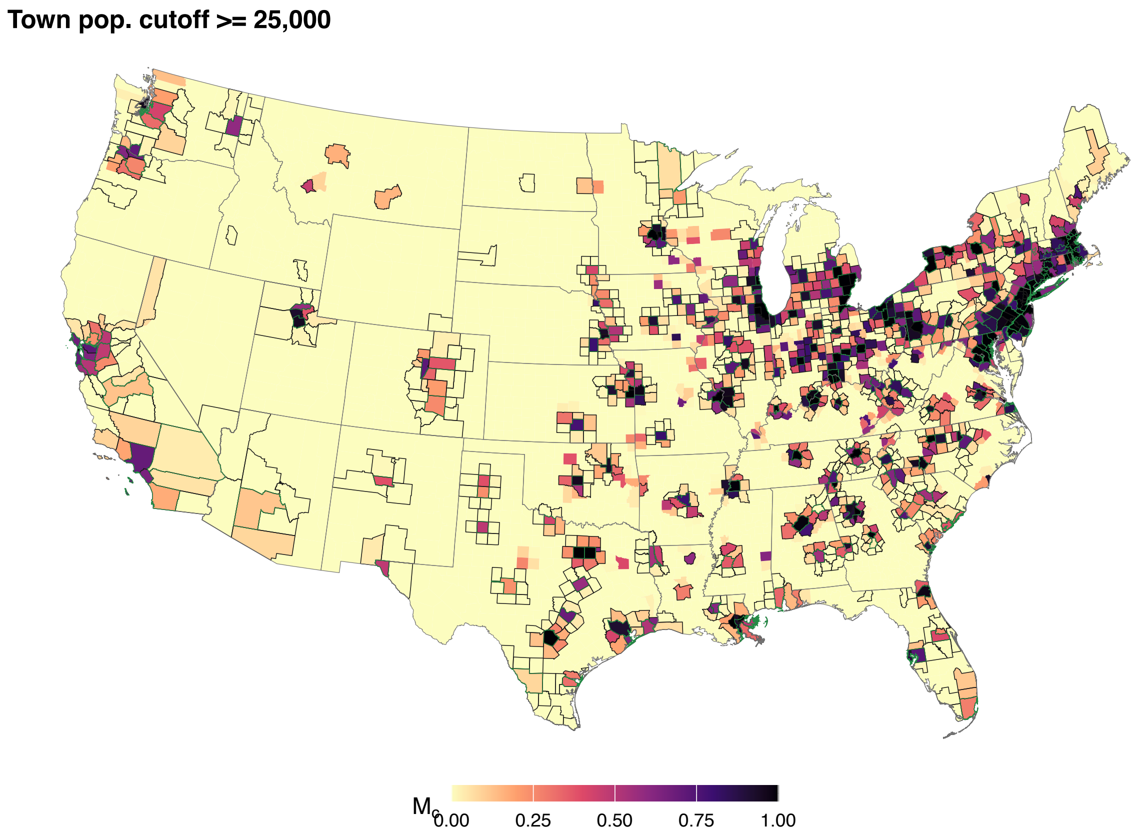

Figure 1: County development pressure M_c, towns $\geq 25{,}000$. The share of each county's convertible 1940 land within a city-size-scaled reach of a town. Pressure concentrates in the Northeast megalopolis and the Great Lakes industrial belt, with scattered Southern, Texan, and coastal-California metros; the agricultural interior, Great Plains, and Mountain West are near zero.

The reach allometry. A town's reach scales with its population through the urban size law: built-up area grows roughly in proportion to $P^{2/3}$, so the radius of a town's developed core grows as $P^{1/3}$. We set each town's reach to $r_j = \min\{k\, P_j^{1/3},\ 75\ \text{km}\}$, fixing the exponent at the allometric $1/3$ and pinning the scale k by anchoring a city of one million residents to a sixty-kilometer reach. The implied reaches are eight kilometers at the $2{,}500$ cutoff, eighteen at $25{,}000$, and twenty-eight at $100{,}000$; the seventy-five-kilometer cap binds for only a handful of the largest metropolitan cores. Because the reach uses

1940 population alone and no post-1940 information, it is predetermined to the interstate. Reaches from neighboring towns are unioned, so overlapping frontiers are not double counted, and water and land already developed by 1940 carry no reach.

Distinctness from market access. M_c measures *local* conversion exposure, not position on the continental market-access gradient, and the two are empirically separable. Across estimation-sample counties M_c correlates 0.33 with log 1950 market potential, the first-stage control, so about 89 percent of its variation is orthogonal to that continental gradient; adding the four suitability shares leaves 88 percent unexplained – suitability alone accounts for only two percent – and that residual variation is what identifies the heterogeneity, so the suitability$\times$pressure interaction is not mechanically confounded. Independence from market access increases with the city-inclusion cutoff, and county rankings are stable across the alternative cutoffs. The pressure coefficient is read holding the market-access gradient fixed: log 1950 market potential enters the difference-in-differences first stage year-interacted, so M_c operates on the residual conversion exposure rather than on position along that gradient. Predetermination to 1940 geography does not by itself establish exogeneity of a moderator, and it does not fully rule out that M_c proxies a post-1940 suburbanization trend that would reshape fringe agriculture regardless of the Interstate. We defend the pressure margin as a maintained identifying assumption on two grounds – the low correlation with market potential and the year-interacted market-potential control – while conceding that the sharpest separation of pressure from soil rests on the marginal-versus-suitable suitability contrast, which the analysis of Section 6.3 treats as coarse.

Interstate Highways IHS route locations come from the Federal Highway Administration's National Highway System, distributed through the Bureau of Transportation Statistics National Transportation Atlas Database (Federal Highway Administration, 2024). Construction timing comes from the Brooks and Liscow (2023) replication of the federal PR-511 highway database, which underlies the approach in Baum-Snow (2007). The PR-

511 reports progress at selected mile markers across the network rather than at every mile; Brooks and Liscow (2023) match these sparse records to the true construction history for each segment. The raw PR-511 includes several construction-phase dates, and we follow Baum-Snow (2007) and Brooks and Liscow (2023) in classifying two: the groundbreaking date, when construction began on a segment's first mile, and the opening date, when that mile was first open to traffic. Appendix Figure A3 shows the distribution of each. Groundbreaking dates are concentrated in a narrow window beginning in 1954, while opening dates are spread across two decades as construction proceeded at different paces across states and terrain. This divergence motivates the treatment-date distinction across estimators: the groundbreaking date serves as the primary treatment timing measure in the difference-in-differences specifications, and the opening date serves as the endogenous regressor in the first-difference IV specifications. Where PR-511 records are incomplete, we supplement with Rand McNally historical highway maps and state highway department construction archives assembled in Frye (2026).

The sample is limited to IHS construction completed by 2002, the final year of the agricultural census panel. Counties whose first segment opened after 2002 are coded as never-treated. Appendix Figure A2 maps the network at 1960, 1970, and 2002. The scattered pattern of completed segments in 1960 and 1970, distributed across nearly every state rather than advancing on a single geographic front, reflects that states controlled their own construction timelines under the federal cost-sharing formula. Most of the eventual system was open by 1970, so the bulk of the identifying variation falls in the mid-1950s through the 1960s and later-treated counties serve as valid controls throughout most of the panel. We further restrict the sample to rural counties outside Metropolitan Statistical Areas and exclude counties at IHS network junctions, where multiple routes intersect and highway access reflects the network's nodal structure rather than the market-integration margin the model identifies. Appendix Figure A1 maps the excluded counties.

Control Variables The historical priorities that shaped IHS route placement (Section 2.1) imply systematic differences between counties that received the IHS and counties that did not. Michaels (2008) provides a foundational covariate set for a similarly rural-focused IHS sample, which we expand following Frye (2026). We group these baseline county characteristics into three dimensions. The first is population and urbanization at midcentury (Schroeder et al., 2025). The second is market centrality, captured by two complementary measures: market potential and a betweenness centrality measure constructed from the 1947 Rand McNally national route plan. Market potential weights surrounding population by distance, capturing the population a county can reach; betweenness centrality counts the number of shortest-path trips between intersection nodes that pass through the county, capturing its position in the underlying road network (Appendix Figure A11). The third is alternative transit infrastructure, comprising the rail network (Atack, 2016) and proximity to ports, airports, and inland waterways (Frye, 2026). Appendix A describes variable construction.

Appendix Table A2 reports baseline characteristics for IHS and non-IHS counties, organized along these three dimensions. IHS counties were larger, more urban, and growing faster at midcentury, with closer proximity to existing metropolitan centers. They were more central in the pre-IHS road network and had greater access to surrounding population. They were also better connected to alternative transit modes, including rail, ports, and airports. These differences are consistent with the planning narrative in Section 2.1.

Descriptive Evidence The aggregate trajectories of US farm value per acre and agricultural product value per acre tracked closely from 1940 through the late 1960s before diverging into the 1970s and beyond (Appendix Figure A4). By 2002, farm value per acre had risen to roughly 4.9 times its 1940 level, while agricultural product value per acre had risen to about 3.1 times.

Appendix Table A1 reports county-level agricultural outcomes for IHS and non-IHS counties in 1950 and 2002. Both groups experienced the consolidation that characterized US agriculture during this era: farm counts fell by roughly sixty percent, farmland contracted, and labor expenditure per farm roughly doubled. The largest IHS-vs-non-IHS divergence is in farmland values per acre: IHS counties moved from $983 in 1950 to $4,006 by 2002, non-IHS counties from $797 to $3,025, and the gap widened from $187 per acre to $982. Crop revenue per harvested acre and the share of cropland in the top four crops rose more in IHS counties; livestock revenue rose less.

Stratifying counties by GAEZ crop suitability uncovers substantial heterogeneity behind these aggregate trends (Appendix Figure A5). Unsuitable counties grew farm value per acre roughly twice as fast as Suitable counties in proportional terms, ending the period at 6.6× their 1940 level versus 3.3× for Suitable counties. They also followed a more pronounced boom-and-bust trajectory in farmland, expanding to a peak of ∼1.14× the 1940 base by 1955 before contracting faster than Suitable counties through the late twentieth century. These descriptive patterns motivate the analysis below.

5 Empirical Strategy

Our staggered difference-in-differences strategy measures how the local agricultural economy in a county evolves after IHS access, comparing IHS counties to non-IHS counties. We code IHS access at the break-ground date, when construction begins on a county's first segment. The first-difference IV specification of Appendix A.1 instead codes IHS access at the opening date, when the first segment is opened to traffic, because the Pershing instrument's timing is constructed from network-completion betweenness centrality and aligns with the sequence of actual segment openings rather than the construction-start sequence that depended on state-level fiscal conditions. The two date conventions index the same absorbing treatment event under different labels, so the choice between break-ground and opening dates concerns first-stage alignment, not what is identified. The DiD and the first-difference IV *designs*, by contrast, target distinct but compatible estimands – a cumulative post-treatment level effect versus an average per-period rate of change – which Appendix A.1 reconciles.

$$Y_{it} = \beta_1 \, \text{IHS}_{it} + \mathbf{X}_i' \boldsymbol{\gamma}_t + \alpha_i + \lambda_t + \varepsilon_{it}, \tag{12}$$

where Y_{it} is the outcome (measured in logs or shares) in county i in year t. IHS_{it} is a binary indicator of whether a county has an interstate highway constructed by year t, coded at the break-ground date rather than highway opening, as the length of the construction period is potentially endogenous to local economic conditions. $\mathbf{X}_i' \boldsymbol{\gamma}_t$ comprises the vector of baseline control variables described in Section 4 interacted with a year t indicator. The variables α_i and λ_t capture county- and year-fixed effects, respectively. Standard errors follow Conley (1999), allowing for spatial correlation across counties within 100 kilometers.[7] The first-difference and first-difference IV specifications, which difference out the county fixed effect, instead cluster two-way by county and state-year, matching the residual structure a differenced design leaves; the level difference-in-differences, where post-fixed-effect spatial correlation is the binding concern, uses the Conley spatial HAC. Each inference scheme is chosen to match the residual correlation its own design leaves, not to mimic a single convention. In our full specifications, we include regression weights based on 1950 county farmland acres. This weighting moves from a regression where each county receives equal weight – since the variables from the Census are aggregated to the county – to one where regression units are aligned with the underlying farmland distribution in the U.S.

Because construction onsets are staggered across the panel, concentrated between 1954 and 1969 with a thin tail into the 1990s (see Appendix Figure A3), this research design is

[7] Results are robust to alternative clustering schemes; Appendix Figure A13 compares the preferred Conley standard errors to county-level and county-by-state-year multi-clustering for the farmland value event study.

vulnerable to bias when estimated with standard OLS (Goodman-Bacon, 2021). We estimate equation 12 using the two-stage difference-in-differences (DID2s) estimator of Gardner et al. (2024). The DID2s estimator falls into a class of imputation estimators, including Borusyak et al. (2024).[8] This estimator removes covariates, group, and period effects in the first stage from a regression using *only* never-treated and not-yet-treated observations; the second stage then uses estimated residuals for *all* observations to identify the IHS effect. Identification requires that this first-stage factor structure imputes the correct counterfactual agricultural trajectory for IHS counties, conditional on the year-interacted control vector.

The heterogeneity specification is the triple interaction of IHS exposure with land suitability and development pressure. Suitability enters as the four FAO rainfed-low classes (N, $S3$, $S2$, $S1$), parametrized as county area shares s_c^k that sum to one, so each cell coefficient is the IHS effect for a county lying entirely in the indicated class. Development pressure enters as the binary M_c^b, which splits counties at the median of the development-pressure distribution of Section 4.1 into "no-pressure" ($M_c^b = 0$) and "under-pressure" ($M_c^b = 1$) groups; at the $25{,}000$-resident city-inclusion cutoff the estimation-sample median is zero, so a county is under development pressure when any qualifying 1940 town lies within its reach, and the continuous dose-response companion uses M_c directly. The suitability partition is carried in *both* estimation stages – the first-stage counterfactual is imputed within suitability class, since pooling common controls across classes is a homogeneity restriction the data reject – while the development-pressure moderator enters only the second stage, so the first stage is not contaminated by the post-access conversion margin the moderator is meant to capture. The eight suitability×pressure cells are the primary heterogeneity object; the continuous dose-response in M_c is the companion that traces the same gradient across the full range of the pressure share. Suitability enters

[8] We choose to estimate with Gardner et al. (2024) because the software package `did2s` by Butts and Gardner (2022) allows for the flexible inclusion of treatment-effect interactions and continuous treatment variables.

the analysis at three granularities, each matched to its task: a binary suitable-versus-unsuitable split for the aggregate-robustness battery (the coarsest and best-powered object), the four FAO classes for the cell-level heterogeneity (which resolve the three-state landing across N, $S3$, and $S2{+}S1$ – retention on the marginal class, crop-acreage displacement and parking on the middle class, and capitalization on the suitable class), and a three-class collapse that merges the thin prime class for the welfare GMM.

Identification comes from the staggered rollout of IHS construction across counties, concentrated in 1954–1969. Route placement was settled before construction began: General Pershing's 1922 military road survey identified strategic routes, subsequent federal planning in the 1940s refined the designated system, and the 1956 Federal-Aid Highway Act funded its construction (Swift, 2011; Rose, 1990). Construction timing varied across counties after designation, with states retaining latitude to prioritize sequencing within their designated routes (Karnes, 2009); among counties that eventually received segments, the selection process tended to target counties with stronger prospective growth trajectories earlier (Frye, 2026). The sample restricts to rural counties outside Metropolitan Statistical Areas and excludes counties at IHS network junctions, where highway access reflects network geometry rather than the market-integration margin the model identifies.

The control vector groups baseline county characteristics into three categories, each measured before the IHS construction era and interacted with year fixed effects so that counties differing on these dimensions can follow distinct agricultural trajectories absent the IHS. Population and urbanization measures address the most direct threat: counties closer to cities and with larger populations were both more likely to receive early IHS access and were already on faster urbanization trajectories that independently reshaped agricultural activity.

Market centrality measures comprise 1950 market potential and the betweenness centrality of each county in the 1947 highway network. Without these controls, the IHS estimate would partly recover the long-run effects of pre-existing market access; the 1947 betweenness specifically addresses the re-centering concern that Borusyak and Hull (2023) raise for instruments constructed from network exposure such as the Pershing Map.

Alternative transit infrastructure measures comprise rail access (Atack, 2016) along with proximity to ports, airports, and inland waterways. Rail in particular predicts both where IHS routes ran and pre-IHS agricultural structure; without this control, the IHS coefficient would partly recover the long-run effects of 19th-century railroad placement.

A Gardner DID2s event study using Persh25_{it} as the absorbing treatment variable supports the instrument's conditional independence and is consistent with the exclusion restriction. Pre-IHS coefficients are jointly indistinguishable from zero across all four outcomes ($p \in \{0.22, 0.79, 0.16, 0.76\}$ for $\ln$ farmland value per acre, $\ln$ farmland, $\ln$ number of farms, and $\ln$ crops market value per harvested acre), and post-period coefficients carry the same signs as the FD-IV second-stage estimates. The event-study construction follows the appendix of Frye (2026).

6 Results

Farmland Values and Agricultural Structure. Figure 2 reports estimates from an event study of Interstate Highway construction on county agricultural land values per acre. Land values significantly increased in the decades following the construction of Interstate Highways relative to land values in non-Interstate counties. Forty years after construction, agricultural land values were about 20% higher in Interstate counties, and about 11% higher on average over the entire post-treatment window. This figure also supports the parallel trends assumption. Conditional on fixed effects and control variables, agricultural land values did not differ significantly between counties that eventually had interstates constructed and those that did not.

Figure 3 reports event study coefficients across other structural measures of the local

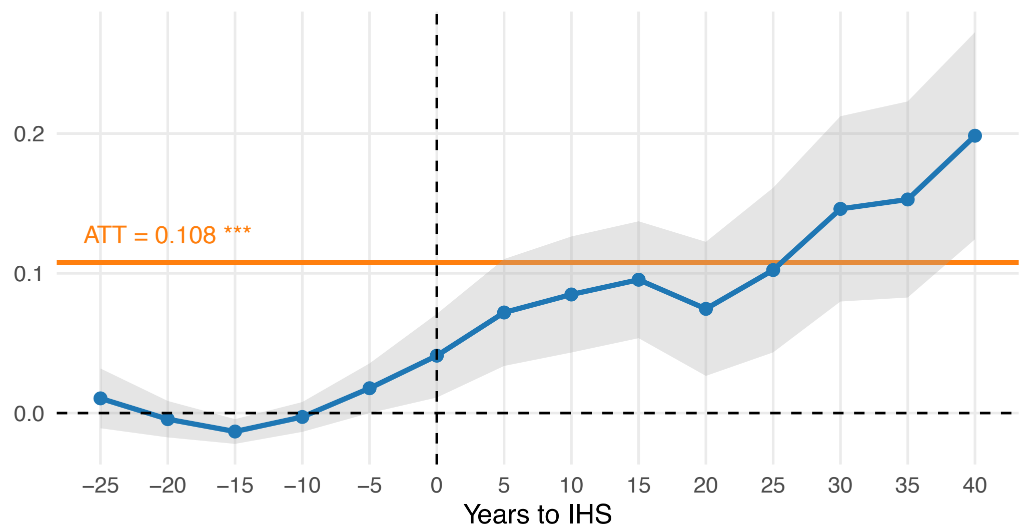

Figure 2: Interstate Highways and Farmland Values

Note: Figure displays coefficients from an event study regression estimated with the two-stage difference-in-differences estimator of Gardner et al. (2024). The dependent variable is logged farmland value ($ per acre). The regression controls for county and year fixed effects and controls for initial cross-sectional measures of urban proximity, market accessibility, and other transit infrastructure, multiplied by agricultural census year indicators. Standard errors are adjusted for spatial correlation within 100km. The gray area displays 95% confidence intervals. The orange line indicates the average treatment effect over all years after IHS construction, and *** indicates statistical significance at the 0.01 level.

agricultural industry. Across all outcomes, there is again support for the parallel trends assumption, as no pre-treatment coefficients are significantly different from zero.[9] After Interstate highways, the total number of farms increased (Panel (a)), but the number of larger farms significantly decreased (Panel (b)). Panel (c) shows that farmland acres in a county were stable before and after highway construction on average, but cropland experienced a moderate decline (Panel (d)). Panels (e) and (f) show that total agricultural revenue and revenue per farmland acre both declined in the decades following Interstate construction. These results run counter to those on farmland values per acre, given that farmland values should reflect expectations about future revenue from that county's land. Treatment-effect heterogeneity by land suitability resolves this contradiction (next subsection).

The Role of Land Suitability Market integration can generate diverging effects across agriculture according to comparative advantage (Section 3). We interact IHS_{it} in Equation 12 with the county's area shares in the FAO suitability classes – marginal (N), S3, and suitable (S2+S1) – so the treatment effect varies with a county's suitability composition.

Figure 4 reports the effects on farmland value per acre (Panel a) and farmland acres (Panel b) by suitability class, estimated without and with suitability-class year trends. Without the trends, the effects diverge sharply along suitability. The farmland-value gains concentrate in the less-suitable classes – $+16$ log points in the marginal class and $+24$ in S3 – while the suitable class loses 13 log points. Farmland acres mirror this: the suitable class expands its base by 21 log points, while the less-suitable classes are flat to contracting. Suitable counties add farmland as their value per acre falls, and less-suitable counties do the reverse – the extensive-margin reallocation and specialization the comparative-advantage model predicts. The average effect across all interstate counties (Figure 3)

[9] Appendix Figure A8 performs a Rambachan and Roth (2023) sensitivity analysis of parallel trend violation for this event-study. The results remain statistically distinguishable from zero at the 95% confidence level when hypothetical post-treatment parallel-trend violations are bounded by the magnitude of the largest observed pre-treatment deviation ($\bar{M} = 1$).

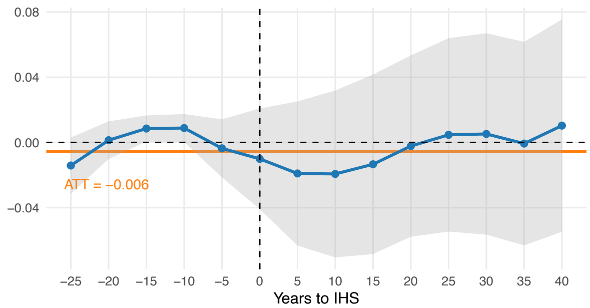

(a) Number of Farms

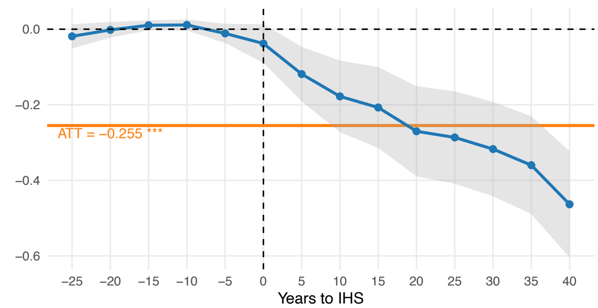

(b) Number of Farms >500 acres

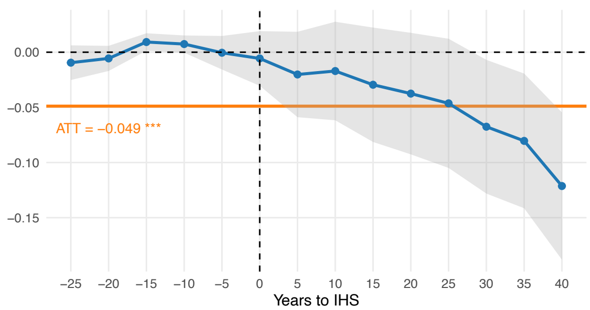

(c) Farmland Acres

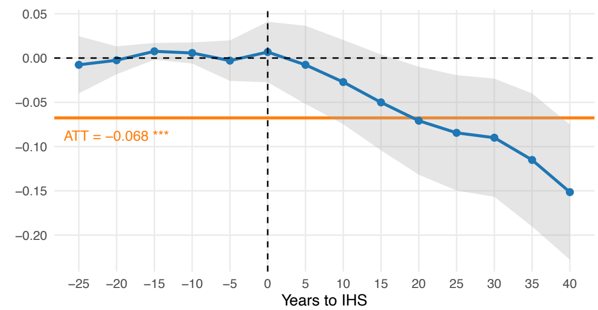

(d) Cropland Acres

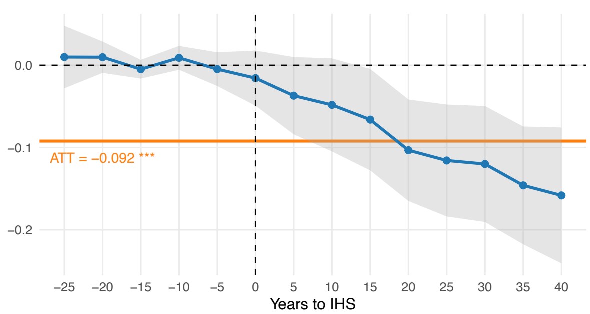

(e) Ag Revenue (Total $)

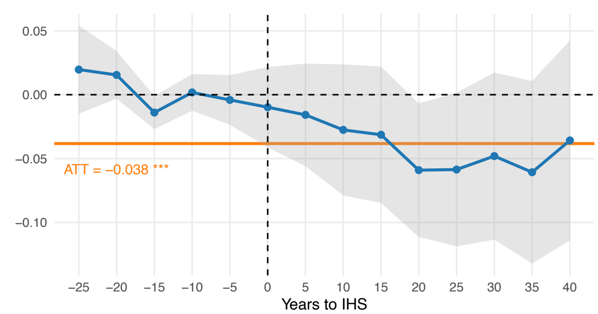

(f) Ag Revenue ($ Per Acre)

Figure 3: IHS Break Ground Date Event Studies

Note: Figure displays coefficients from event study regressions estimated with the two-stage difference-in-differences estimator of Gardner et al. (2024). The regressions control for county and year fixed effects and control for initial cross-sectional measures of urban proximity, market accessibility, and other transit infrastructure, multiplied by agricultural census year indicators. Standard errors are adjusted for spatial correlation within 100km. The gray area displays 95% confidence intervals. The orange line indicates the average treatment effect over all years after IHS construction, and * indicates statistical significance at the 0.1 level.

masks this heterogeneity.

The divergence in value per acre is the accounting shadow of the acreage reallocation: adding marginal, lower-value land in suitable counties pulls their average value down, while shedding the least-productive land in less-suitable counties raises theirs. Adding suitability-class year trends attenuates the raw contrast – the marginal class instead retains farmland ($+19$ log points) and the value differences compress – so the formal analysis below tracks these trends throughout, and the cell-level specification recovers the pattern that survives them.

The divergence is robust across specifications – full versus restricted sample, with and without the control vector, and area-weighted – and holds under a continuous development-pressure measure (Appendix Tables A7 and A8).

6.1 Cell-level evidence: retention, displacement, and capitalization by suitability

The suitability gradient above sharpens into three identification classes built from the four FAO grades – marginal (N), S3, and a combined suitable class (S2+S1) – each interacted with development pressure, and the disaggregation locates the action at the absolute bottom of the distribution. The prime grade S1 is not broken out on its own: it has thin identifying support – the least of the four grades on the baseline rainfed-low surface – and its cells fall below the reliability bar under any finer pressure stratification, so it is merged into the suitable class (Appendix A.4). Counties in FAO class N – the worst grain-suitability classification – retain agricultural land and hold agricultural value rather than contract with the broad low-suitability average. Table 3 reports the highway effect across the three classes interacted with development pressure for the primary outcomes; the appendix battery (Tables A9, A10, and A11) reports the full outcome set.

On the rainfed-low surface, the marginal class at low development pressure retains 17 log points of farmland (N×none, Table 3, column 1) and gains 25 log points of to-

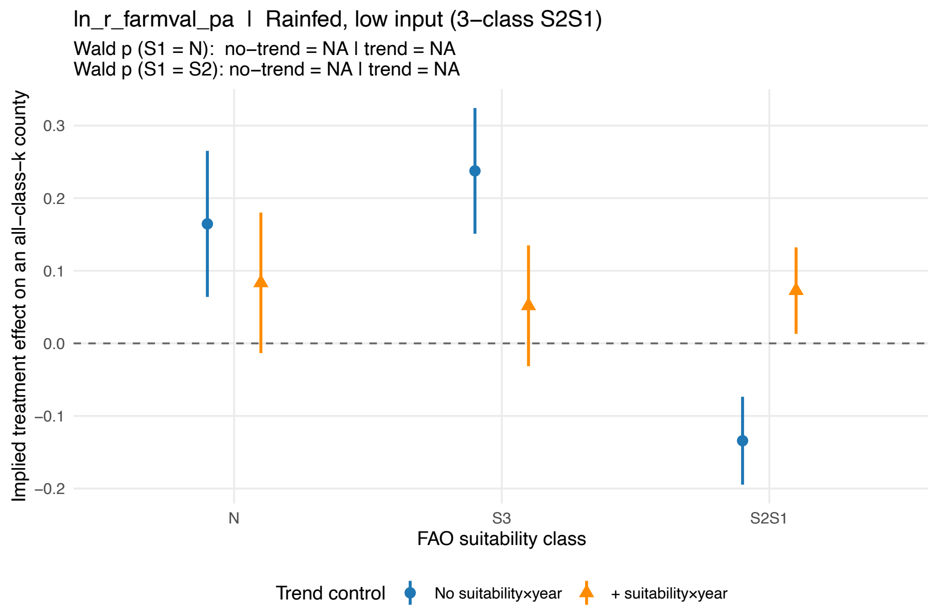

(a) Farmland Value ($ per acre)

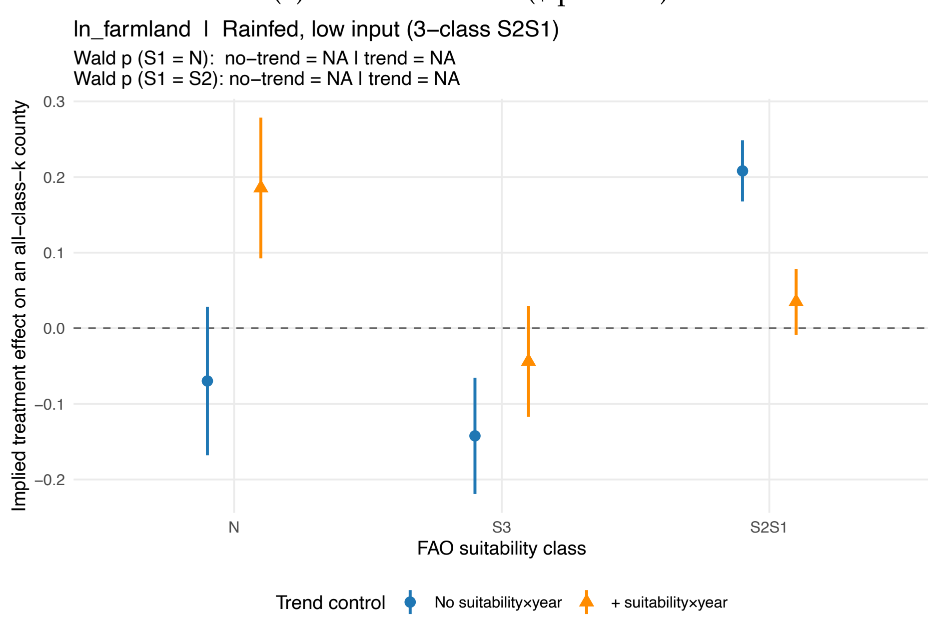

(b) Farmland Acres

Figure 4: IHS Treatment Effect by FAO Suitability Class

Note: Figure displays heterogeneous treatment effects of Interstate Highways by FAO suitability class, from a Gardner (2024) two-stage difference-in-differences that interacts the IHS treatment with county area shares in the marginal (N), S3, and suitable (S2+S1) classes. Each point is the implied effect for a county fully in that class; the two estimates per class are estimated without and with suitability-class×year trend controls. The dependent variables are logged farmland value ($ per acre) and logged farmland acres in Panels (a) and (b). All models include county and year fixed effects and year-interacted controls for urban proximity, market access, and prior transit infrastructure. Conley (1999) standard errors with a 100 km cut-off; bars are 95% confidence intervals.

Table 3: Highway effects on agriculture by suitability class and development pressure.

	Survival		Inten.	Land value		Composition		
Dep. var. (ln):	Farmland (1)	Harv. crop (2)	Tractors/farm (3)	Total value (4)	Value/acre (5)	Cattle (6)	Lvstk. rev. (7)	Perishable (8)
Panel A: Rainfed-low suitability								
N (marginal) × none	0.17***	0.11	−0.02	0.25***	0.09	0.32***	0.28***	0.69***
	(0.05)	(0.09)	(0.03)	(0.06)	(0.06)	(0.07)	(0.09)	(0.22)
N (marginal) × pressure	−0.05	−0.03	0.05	−0.10*	−0.06	−0.10**	−0.06	−1.31**
	(0.09)	(0.10)	(0.03)	(0.08)	(0.07)	(0.11)	(0.14)	(0.26)
S3 × none	−0.01	−0.04	−0.03	0.01	0.02	0.04	−0.04	−0.09
	(0.03)	(0.06)	(0.02)	(0.04)	(0.04)	(0.05)	(0.09)	(0.08)
S3 × pressure	−0.13***	−0.10*	0.08**	−0.09*	0.03	−0.23***	−0.59***	−0.06
	(0.05)	(0.08)	(0.03)	(0.04)	(0.05)	(0.06)	(0.13)	(0.10)
S2+S1 (suitable) × none	0.00	−0.04	−0.05**	0.07	0.07*	−0.07	−0.05	−0.02
	(0.03)	(0.05)	(0.02)	(0.04)	(0.04)	(0.06)	(0.05)	(0.12)
S2+S1 (suitable) × pressure	0.01	0.09	0.08**	0.05***	0.04***	−0.22***	−0.30***	0.04
	(0.03)	(0.04)	(0.02)	(0.03)	(0.03)	(0.06)	(0.06)	(0.09)
Panel B: Rainfed-high suitability								
N (marginal) × none	0.10***	0.02	−0.03	0.16***	0.06	0.14**	0.14*	0.37**
	(0.03)	(0.06)	(0.02)	(0.04)	(0.05)	(0.06)	(0.08)	(0.18)
N (marginal) × pressure	−0.06	−0.09	0.05	0.13***	0.18***	−0.05	0.17*	−0.82**
	(0.06)	(0.08)	(0.03)	(0.07)	(0.07)	(0.12)	(0.16)	(0.21)
S3 × none	0.02	0.05	−0.02	0.10	0.09	0.10	0.07	0.10
	(0.05)	(0.06)	(0.03)	(0.07)	(0.06)	(0.08)	(0.09)	(0.14)
S3 × pressure	0.03	0.03	0.04	0.05*	0.03	0.04*	−0.05	−0.21
	(0.07)	(0.09)	(0.03)	(0.08)	(0.08)	(0.08)	(0.12)	(0.21)
S2+S1 (suitable) × none	−0.06	−0.12*	−0.05***	0.03	0.09**	0.02	−0.03	−0.08
	(0.04)	(0.06)	(0.02)	(0.04)	(0.04)	(0.05)	(0.08)	(0.09)
S2+S1 (suitable) × pressure	0.01	0.10	0.11***	0.03**	0.02***	−0.25***	−0.37***	0.01
	(0.05)	(0.07)	(0.02)	(0.03)	(0.04)	(0.05)	(0.07)	(0.08)

Notes. Each cell reports the IHS-opening treatment effect from a Gardner (2024) two-stage difference-in-differences on the KeepSet county panel; Conley (100 km) standard errors in parentheses. The three identification classes (N marginal, S3, and S2+S1 suitable) enter as continuous county area-shares summing to one; development pressure is the binary $M_c^b = \mathbf{1}\{M_c > \mathrm{med}(M_c)\}$ splitting the estimation sample at the median of M_c, "none" (below median) versus "pressure" (above). The "none" row is the effect for a county fully in the class at low pressure; the "pressure" row is the additional effect under development pressure. Suitability shares interact with year in the first stage; the moderator enters only the second stage. Panels vary the GAEZ input parameterization (rainfed-low and rainfed-high). Perishable-acreage outcomes use the pre-1978 consistent-definition window (1940–1974). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

tal agricultural value (column 4), both significant at the one percent level, alongside a positive but statistically insignificant 9 log points of farmland value per acre (column 5). The retention survives the alternative rainfed surface: under rainfed-high suitability the same cell holds 10 log points of farmland and 16 log points of total value, both significant at the one percent level, so the pattern is not an artifact of one GAEZ parameterization. The mid-distribution class S3 and the suitable class S2+S1 are null or weakly contracting in their dominant low-pressure cells, so the low-pressure retention concentrates at the marginal extreme rather than spreading across the low-suitability range. Under development pressure the three classes trace the three states the model leaves on the soil ladder: the marginal class N retains its base, the middle class S3 is where the crop margin is displaced, and the suitable class S2+S1 is where the development option capitalizes into value. The middle-soil displacement is visible directly in Table 3: S3×pressure sheds farmland (-0.13) and harvested cropland (-0.10) as its extensive use collapses (cattle -0.23, livestock revenue -0.59) and its survivors intensify (tractors $+0.08$). The N-class gains are composition-driven: the agricultural base is held while the activity mix shifts, not expanded in commodity grain, as the next subsection documents.

Production and Structural Margins. The structural margins behind the value gradient sort along the same suitability axis, and away from development pressure they run toward livestock rather than crops. The marginal class holds its agricultural base rather than shedding it: farmland ($+0.17$), pastureland ($+0.18$), and cattle inventory ($+0.32$) all rise, livestock revenue rises 28 log points, and vegetable and orchard acreage expand sharply ($+0.83$ and $+0.52$) – the perishable-and-livestock reorganization for which marginal land retains a comparative advantage (Table 3; Appendix Tables A9–A11). Crop revenue, by contrast, is positive but imprecise on marginal land ($+0.12$, not significant) and machinery intensity is flat, so the retention operates through extensive livestock and perishable use, not crop intensification. In the most-suitable class the value response is muted ($+0.07$ total value) and crop revenue does not differentially rise (-0.11, not significant): the crop-intensification response is not a land-quality effect. Where it appears is under development pressure – suitable counties under pressure raise crop revenue ($+0.26$) as livestock revenue falls (-0.30) and value per acre rises ($+0.04$) – the best-soil landing, where the development option capitalizes into land value while the crop base is retained and intensified rather than displaced. This is what separates the suitable class from the middle: development pressure sheds the extensive grazing use on both, but only on the middle soil does the crop acreage itself convert, whereas on the best soil it is held and the option prices into value.

6.2 Spatial Spillovers

The model predicts that reductions in τ are county-specific: the threshold ω^* shifts for the county receiving direct IHS access, not for its neighbors. Yet counties within reach of an IHS ramp may also face lower effective transport costs, generating spillovers that operate through the same market integration mechanism at an attenuated level. Appendix Figure A6 illustrates that this concern is empirically non-trivial: by 2002, a substantial and growing share of rural counties fell within 75 kilometers of an IHS county, raising the possibility that adjacent counties in the control group are themselves partially treated. We address this by augmenting the DID2s second stage to include both the direct treatment indicator and a separate indicator for counties within 75 kilometers of a treated county, estimating the two effects jointly so that only counties in neither category serve as controls in the first stage.

Direct county effects mirror the main specification throughout the pre- and post-treatment window for both farmland value per acre and aggregate agricultural revenue per acre (Figure 5). The two outcomes diverge in adjacent counties, but in a window-dependent way. The static-fit average effect on farmland value per acre is 1.4 percent in adjacent counties versus 9.4 percent in direct counties, statistically indistinguishable from zero on average across the post-treatment window; the adjacent event-study profile rises late in the window, suggesting that capitalization in adjacent counties accumulates more slowly than in direct counties and may not have completed within the 1940–2002 Census panel. Agricultural revenue per acre falls 8.9 percent in adjacent counties versus 13.2 percent in direct counties, an adjacent effect approximately two-thirds the direct magnitude that emerges earlier in the post-treatment window. Appendix Figures A7 and A9 report the full set of adjacency event studies: farm counts, farm labor expenditure, livestock revenue per acre, and the farm-size and acreage composition margins exhibit attenuated adjacent effects of the same sign as the direct effects. The asymmetry between the static averages and the late-window capitalization profile follows from the model's county-specific threshold prediction: only direct IHS access shifts ω^*, so capitalization concentrates in direct counties, while spillovers in transport costs propagate to revenue and input margins in adjacent counties more rapidly than they capitalize into land values.

6.3 The Development-Pressure Gradient

The continuous dose-response. A continuous companion traces the pressure gradient across the full range of the coverage share (Table 4), sorting the marginal grazing base into the two destinations the model predicts (Section 3.3). Where no builder bids, the base is retained: at zero pressure the marginal class holds farmland ($+0.19$), cattle inventory ($+0.36$), and livestock revenue ($+0.35$), each significant at the one percent level. Where pressure rises, that same band is displaced, and the displacement concentrates on the extensive-grazing margin the comparative-advantage mechanism singles out: the coverage slopes on cattle (-1.01) and livestock revenue (-1.38) are large and significant at the one percent level, and total agricultural value falls with pressure (-0.57, significant at five percent), while the farmland and value-per-acre slopes are same-signed but imprecise (-0.49 and -0.06). Prime land is squeezed more deeply the more of the county is exposed (S1 farmland slope -0.30, significant at one percent). Under development pressure the

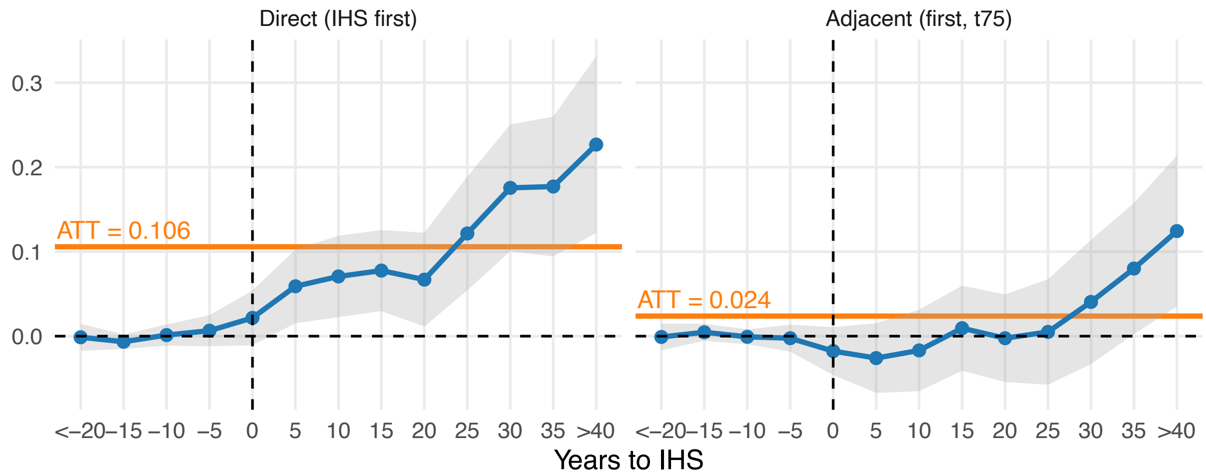

(a) Farmland Value ($ per acre)

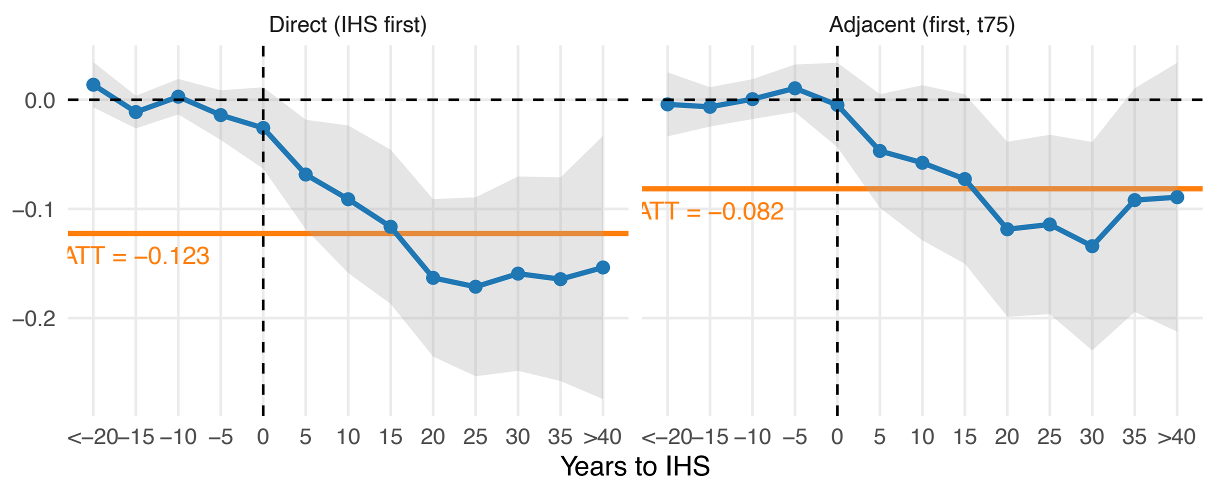

(b) Agricultural Revenue ($ per acre)

Figure 5: Spatial Spillovers: Direct and Adjacent County Effects

Notes. Each panel displays event study coefficients for directly treated counties (left) and counties within 75 km of a treated county (right), estimated jointly in the DID2s second stage. Only counties in neither category serve as controls in the first stage. Regressions control for county and year fixed effects and for 1950 measures of urban proximity, market accessibility, and prior transit infrastructure, each interacted with census year indicators. Standard errors adjusted for spatial correlation within 100 km. Shaded area is 95 percent confidence interval. Orange horizontal line marks the static-fit average ATT for each series.

marginal grazing band thus cedes to the residential bid – the displaced destination of the two-destinations corollary – rather than being held as it is where pressure is absent. The continuous dose therefore locates the pressure response in the grazing-band displacement rather than isolating a value-per-acre capitalization from an acreage response; the discrete pressure contrast and the value levels carry that reading. The broken-out S1×pressure farmland contraction (-0.15) is the cleanest single-cell sign of conversion, though the model organizes rather than signs it: at the prime end the net farm-gate incidence (A1) and the A2 conversion pull oppose each other, and prime land away from pressure is flat rather than expanding.

Table 4: Continuous dose-response: the IHS effect and the share of county land under pressure.

	Marginal (N)		Prime (S1)	
Outcome (ln)	level	slope/share	level	slope/share
Farmland	0.19***	−0.49	0.00	−0.30***
Harvested crop	0.11	−0.22	−0.09	0.11
Tractors per farm	−0.02	0.17	−0.06	0.26***
Total ag. value	0.26***	−0.57**	0.14	−0.36***
Value per acre	0.07	−0.06	0.14	−0.05
Cattle	0.36***	−1.01***	0.08	−0.97***
Livestock revenue	0.35***	−1.38***	0.07	−1.18***
Perishable acres	0.52***	−3.30***	0.37	−0.89**

Notes. The moderator M_c enters continuously as the share of county convertible 1940 land within reach of a 1940 town, interacted with the suitability class. *level* is the IHS effect at zero coverage; *slope/share* is the change per unit of coverage. Marginal-class (N) retention is positive at zero pressure and declines as more of the county comes under pressure (negative slopes), the attenuation Assumption A2 predicts; prime-class (S1) farmland is squeezed more deeply the more of the county is exposed. Because M_c is concentrated near zero, the slopes read most reliably as sign and significance; the binary six-cell grid (Table 3) is the primary contrast. Gardner DiD2s, rainfed-low surface, KeepSet; Conley (100 km) standard errors. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The composition mechanism: livestock. The marginal-class retention is composition-driven, operating through Assumption A3 (livestock). On the rainfed-low surface, cattle inventory in the marginal class away from pressure rises by 32 log points and livestock revenue by 28 log points (Table 3, columns 6–7; both significant at one percent), and the pattern survives the rainfed-high surface (cattle $+14$, livestock revenue $+14$). The cattle and livestock-revenue aggregates corroborate the extensive-grazing retention the model describes but cannot separate it from access-driven dairy and feedlot expansion the single-technology model does not own; both load with the same sign. Under the two-destinations corollary the model predicts this marginal grazing band is retained where no builder bids and displaced under development pressure. The binary under-pressure interaction on livestock revenue is small and imprecise – its minimum detectable effect, roughly 0.42 log points, exceeds the point estimate – while the continuous dose-response registers the displacement directly (livestock-revenue slope -1.38, Section 6.3), so the coarse binary split understates a displacement the continuous measure makes clear. Perishable acreage in the marginal class also rises, by 69 log points at low pressure – driven by vegetables ($+83$ log points, significant at one percent; orchards $+52$, significant at five percent; Appendix Table A11) – and by 37 log points on the rainfed-high surface; the model assigns this margin no structural role.

Robustness of the headline heterogeneity. Appendix Tables A5–A6 show the suitability gradient is robust to interacting IHS access, by suitability class, with four proximity and pre-existing-transit dimensions – metropolitan proximity, 1947 highway betweenness, rail density, and inland-waterway proximity.

What the pressure contrast identifies. The pressure claim runs through the value–revenue wedge (Proposition 1, falsification test N1) rather than through the agricultural land share: conditioning on M_c orders the wedge but does not collapse the cross-suitability survival gradient, which is present in the no-pressure stratum and soil-ordered (Lemma 1). Because M_c is a coarse correlate of suitability – a marginal-versus-suitable contrast that is flat within the suitable classes (condition C1) – the wedge mediation has limited within-suitable pressure variation to exploit, and the decisive separation of pressure from soil rests on the marginal-versus-suitable margin.

7 Welfare Decomposition

This section quantifies the welfare component that the Census farmland-value aggregate omits when converted parcels leave the sampling frame. When Interstate Highway construction makes a rural county more accessible to nearby cities, some of that county's farmland is eventually sold to housing developers, retailers, or industrial users. The land does not become more agriculturally productive; it becomes more valuable in a different use. Once it converts, it leaves the Census of Agriculture's sampling frame – the Census records the value of land *that remains in farms*. A welfare measure built from the Census farmland-value aggregate, the standard empirical input to the Donaldson and Hornbeck (2016) framework, therefore captures only part of the land-value gain that a transportation improvement generates.

In the Donaldson and Hornbeck (2016) framework, the welfare gain from access improvement equals the change in total county land value, which decomposes into an agricultural component W^A_{ct} (acres remaining in farms, valued in their agricultural use) and an outside-use component W^O_{ct} (acres converted to non-agricultural use, valued in their new use). The Census of Agriculture measures W^A_{ct}; W^O_{ct} is what the Census misses. Appendix B.10 derives the model objects formally and Appendix Table A20 maps each to its Agricultural Census counterpart. In the open-frontier expansion setting that Donaldson and Hornbeck (2016) study, W^O_{ct} is small relative to W^A_{ct} and the Census aggregate captures essentially the full welfare gain. In the settled-frontier conversion setting we study, W^O_{ct} is non-trivial in counties where the conversion margin is active, and the Census aggregate is a lower bound on the welfare gain rather than the full measure.

Section 7.1 documents the cross-suitability heterogeneity in total farmland-value response that distinguishes revenue-channel from composition-channel counties. Section 7.2

reports a direct US estimate of the one parameter the welfare decomposition borrows from outside the pooled treatment effects – the parcel-level productivity tail index θ – estimated separately on each of the two GAEZ suitability surfaces the analysis uses. Section 7.3 maps the pressure-stratified treatment effects to the model's channel primitives through a minimum-distance GMM and reports the omission as a set-identified quantity, bounded above by the external development value of the converted land.

7.1 Total Farmland Value by Suitability

Total farmland value – value per acre times farmland acres – rises in IHS-treated counties, with the response concentrated in the least-suitable land. Away from development pressure, total value in the marginal (N) class rises by 25 log points (Table 3, column 4), against a modest and insignificant 7 log points in the most-suitable (S2–S1) class; under development pressure the marginal-class response reverses to -10 log points, the displacement margin.

The model decomposes this response into three channels. A *revenue channel* operates on parcels that remain in agriculture: as access improves, output prices rise and the value of every farming acre increases. A *cost channel* operates through input markets: as access improves, local labor markets may tighten as workers commute to non-farm jobs while transport-cost reductions lower per-unit prices of fertilizer, seed, and other tradeable inputs. A *composition channel* operates at the participation margin: as access improves, the return to non-agricultural use rises, and marginal parcels exit farming. The acres that remain in farms are systematically the most agriculturally valuable ones in the county, so the per-acre value of the surviving land rises further. The composition channel adds value to the farmland-value aggregate even when within-agriculture productivity has not changed – it is the capitalization of conversion-option value into the prices of the remaining agricultural base.

The input-market evidence in Section 6 shows reorganization within the cost channel: a mechanization-and-intensification substitution where farms remain – consistent with workers leaving agriculture for non-farm employment as remaining farms scale up – against a joint decline in inputs where agriculture deactivates at the participation margin. The net per-acre cost elasticity recovered by the structural GMM in Section 7.3 is modest because these within-cost-channel components substantially offset in suitable counties; the cost channel is active in the data but small in net effect relative to the outside-use channel that drives the welfare bound. Appendix B.4's capitalization-incidence proposition (Proposition 1) admits this cost-channel activity; the value–revenue capitalization wedge is robust to it because the common cost-channel component cancels in the cross-county subtraction.

In high-suitability counties, the typical parcel lies far from the participation margin (its agricultural use dominates outside use by a comfortable margin), so the revenue channel does most of the work. In low-suitability counties, the typical parcel sits closer to the margin, so the composition channel is more important. The empirical asymmetry in Table 3 – a total-value response of 25 log points in the marginal class away from pressure against 7 in the most-suitable class – is consistent with the composition channel adding substantial value in the marginal cells beyond what the revenue channel alone would generate. This decomposition into revenue, cost, and composition channels is the input to the welfare bound in Section 7.3.

7.2 Calibrating the Tail Parameter θ

The welfare decomposition requires one parameter that does not appear in the pressure-stratified treatment effects of Section 6: θ, the tail index of the parcel-level productivity distribution. Conceptually, θ governs how much agricultural productivity varies across parcels within a county. A high θ corresponds to relatively homogeneous parcels (all land in a county has roughly similar productivity); a low θ corresponds to substantial within-county variation (some parcels are very productive, others are marginal).

We estimate θ directly from US data using the Sotelo (2020) eq (4) ratio identification applied to the US Census of Agriculture panel. Stacking county-crop-year cells across 1940–1974 for the nine Ricardian crops in the paper's suitability index (corn, wheat, soybean, sorghum, barley, oats, cotton, white potato, wetland rice; alfalfa excluded on residual-structure and GAEZ-collinearity grounds documented in Appendix B.15) and conditioning on county fixed effects and crop-by-year fixed effects, with corn as the numeraire, the slope of relative log acreage on relative log suitability identifies θ. Because the GAEZ surface enters on the right-hand side, the estimate is internally tied to the suitability surface used to construct it, and we estimate θ once on each of the two surfaces the welfare analysis uses: $\hat{\theta} = 2.13$ on the rainfed-low surface and $\hat{\theta} = 1.61$ on the rainfed-high surface. The higher-input surface compresses the cross-crop productivity differences and so implies a lower tail index. We pair each surface with its own θ in the welfare GMM rather than imposing a single common value, so that the structural dispersion parameter and the suitability data enter consistently on the same surface.

The literature provides comparators that bracket the band. The closest analog in the structural Ricardian agriculture literature is Sotelo (2020, p. 2715), who estimates $\hat{\theta} = 1.658$ on plot-level land qualities in Peruvian agricultural data at the district level (his finest spatial unit, on the subsample of 4 of Peru's 24 departments with the strongest GAEZ spatial match). Costinot et al. (2016, Tables 2, 5) estimate $\hat{\theta}^{\text{CDS}} = 2.46$ on within-field productivity from cross-country crop gravity flows on a 50-country sample, with their 29-country developed-country subsample (including the US) yielding $\hat{\theta}^{\text{CDS}}_{\text{Developed}} = 2.29$ and the 21-country developing-country subsample yielding $\hat{\theta}^{\text{CDS}}_{\text{Developing}} = 2.71$. Our dual-surface band $[1.61, 2.13]$ lies between the Sotelo Peruvian-district estimate and the CDS developed-country subsample, with the rainfed-high estimate essentially on top of the Sotelo value and the rainfed-low estimate just below the CDS developed-country figure.[10]

[10] We do not adopt trade-elasticity estimates from Eaton and Kortum (2002) or US settings (Herzog, 2021; Jaworski et al., 2023) as comparators: the trade-elasticity parameter in those papers operates at a different

The qualitative cross-sectional prediction – that the omitted share is concentrated in low-suitability counties – does not require numerical estimation of θ. As shown formally in Appendix B.10, W_{ct}^{A}/W_{ct} is monotone increasing in agricultural suitability A_c for any $\theta > 0$: better-suited counties have fewer marginal parcels sitting near the participation boundary, so a smaller share of county welfare derives from the outside-use component, regardless of how steep the productivity distribution is. The cross-sectional monotonicity argument additionally requires that the outside-use return is approximately invariant in suitability at fixed access, because soil quality and growing-degree days are irrelevant to a parcel's value as residential or commercial land. Section 5's controls for population, urbanization, market centrality, and alternative transit infrastructure absorb the residual variation in outside-use returns that does correlate with suitability.

7.3 Quantifying the Omitted Share

The pooled two-activity welfare GMM is the bridge from the pressure-stratified treatment effects above to a structurally interpretable welfare aggregate. It over-identifies a vector of eight structural Jacobian primitives, stratified by development pressure M_c – the pre-determined city-reach exposure of Section 4.1 at the 25,000-population cutoff, resolved into $K = 3$ strata by a median split of positive exposure:

$$\boldsymbol{\pi} = \left(\pi_{b_C},\ \pi_{b_L},\ \pi_c^{\text{None}},\ \pi_c^{\text{Low}},\ \pi_c^{\text{High}},\ \pi_R^{\text{None}},\ \pi_R^{\text{Low}},\ \pi_R^{\text{High}}\right)',$$

where π_{b_C} and π_{b_L} are the crop and livestock output-price slopes, π_c^{σ} is the cost-channel slope, and π_R^{σ} is the outside-option slope, the latter two indexed by the development-pressure stratum $\sigma \in \{\text{None}, \text{Low}, \text{High}\}$. The two farm-gate slopes clear in national output markets and are held common across strata; the cost and outside-option slopes

aggregation level than the parcel-level Fréchet shape required here. The within-county farm size distribution would identify θ only under a monotone size-productivity mapping that the activity-conditional dependence in Appendix B.7 rules out.

clear in local factor and land markets and split by pressure stratum. The moments are formed from stratum-level ATTs on four base outcomes,

$$\ln V^A_{ct} = \ln\big(\text{farmland}_{ct} \cdot \text{value per acre}_{ct}\big)\,, \quad \ln S^A_{ct}, \quad \ln \overline{\text{Rev}}^C_{ct}, \quad \ln \overline{\text{Rev}}^L_{ct},$$

farmland value per acre, agricultural land share, crop revenue, and livestock revenue. The two-activity split of the revenue outcome lets the crop and livestock farm-gate slopes enter separately, which the earlier value/acreage pair collapsed. The moment system stacks four outcomes across three pressure strata into twelve moments and estimates eight primitives, leaving four Hansen degrees of freedom (Appendix B.11). Appendix B.11 formalizes the moment system and the bootstrap procedure.

Where the stratum ATTs come from. All four base outcomes are read off one common estimation sample: every county-year enters $\ln V^A$, $\ln S^A$, $\ln \overline{\text{Rev}}^C$, and $\ln \overline{\text{Rev}}^L$ jointly non-missing, so the four stratum ATTs share a single sample and a disclosure-suppressed crop-revenue subset cannot leak leverage into one moment alone. The ATTs come from the pooled Gardner DID2s estimated over the $K = 3$ pressure strata, with suitability entering as a first-stage year-interacted trend control rather than a second-stage interaction. This pooling – pressure resolved into strata, suitability held in the first stage – is the canonical structural specification. The finer per-cell suitability$\times$pressure imputation grid is retained only as a demoted robustness because it lacks identifying power in the thin prime cells: breaking the prime (S1) class out under a finer split drops its support to as few as eight to twenty effective treated counties, too few to resolve a cell-specific slope, so it is merged into the suitable class (Appendix A.4). The pooled system is chosen on that power ground, not on a formal test rejecting the grid. The aggregation that maps stratum-level welfare derivatives to the national bound nonetheless uses each county's full four-share suitability composition, so no county-level information is discarded in forming the aggregate. We estimate the system on two internally consistent surface–θ runs: the rainfed-low surface pairs the rainfed-low baselines with $\theta = 2.13$, and the rainfed-high surface pairs the rainfed-high baselines with $\theta = 1.61$, so the surface enters consistently on both the data and the structural side.

The pressure strata. Pressure enters the canonical system through $K = 3$ strata, None, Low, and High, where None collects counties with no city-reach exposure and Low and High split the counties with positive exposure at the median of M_c. Suitability is not a second-stage cut in the canonical system; it enters the first stage as a year-interacted trend control, so the strata pool pressure over suitability along the dimension the DID2s second stage varies. The earlier per-cell grid, which crossed the four FAO suitability classes with the pressure split, is retained only as demoted robustness because the thin prime (S1) cells lack the support to identify cell-specific slopes (Appendix A.4). The four base outcomes are stacked across the three strata to form the moment vector, and the system over-identifies the eight primitives.

The omission and its external anchor. The farmland aggregate omits the development value of converted land. The acreage that enters this omission is the *gross* farm-to-development flow, while the difference-in-differences measures the *net* acreage change, which nets entry and reversion against that exit (Section 3.3); on the conversion cells the net change therefore understates the gross conversion the omission prices, so building the omission off the net response is conservative. Quantifying that value fully requires the developer's bid R^O, which the GMM cannot identify: the estimated system identifies the exit *sum* at the marginal parcel, roughly the land's agricultural use value ($\approx 1\times$), while the developer bid runs an order of magnitude higher ($\approx 8.6\times$ on the fringe). The economically-complete omission therefore needs an external anchor, and we report the development-value figure as an anchored upper bound rather than an identified point. Table 5 states the external figure across the two surfaces. Valued at observed developer bids, the omission reaches \$200.7–210.1 billion on the rainfed-low surface and \$195.0–199.0 billion on the rainfed-high surface, the internal range spanning the POST- and PRE-highway property anchor. Across the $1{,}338$ IHS-treated counties, this amounts to roughly $150 million per treated county.

The development-value figure re-prices the converted inframarginal stock as a level term from its agricultural use value up to the developer bid R^O, drawn from observed fringe property values (Haines et al., 2018; Appendix B.11). It is an external calibration, disclosed as such and not identified by the design: it is the economically-complete "at most," the development value the farmland books miss when converted land is valued at observed developer bids. The standard farmland-value aggregate is therefore a lower bound on the welfare gain, and pinning the omission's precise magnitude is the subject of ongoing work. For scale, that aggregate itself captures a ΔV^A of $54.3 billion on the rainfed-low surface and $65.9 billion on the rainfed-high surface; if the omission stands at its disclosed upper bound, the conventional measure would capture roughly a fifth to a quarter of the total gain.

The property comparator establishes why the development-value figure is an upper bound rather than a point. Pricing converted land at R^O requires a clean off-network counterfactual for the developer bid, and no such comparator exists: interstate placement is collinear with the urban proximity that generates any development premium, so a spillover-clean never-treated county resembling a treated fringe county has essentially no development premium and its matched anchor is degenerate. R^O is therefore not point-identifiable off a comparator, and the development-value figure inherits a bound rather than a point. The internal PRE$\rightarrow$POST sub-range bounds the double-count: on the high-pressure fringe the observed development premium is flat between the pre-highway and post-highway periods, so valuing induced conversion at observed prices does not substantially recount a premium the highway itself created, and the maximum double-count is bounded by the small post-minus-pre gap. The development-value figure is the disclosed, externally anchored "at most" that this comparator licenses as an upper bound,

Table 5: The welfare omission valued at observed developer bids

Surface	ΔW^O ($B)	Identification	Provenance / assumption
Low-suitability surface ($\theta = 2.13$)			
Developer-bid value (PRE anchor)	210.1	external calibration	Haines property R^O, PRE-highway anchor
Developer-bid value (POST anchor)	200.7	external calibration	Haines property R^O, observed fringe
High-suitability surface ($\theta = 1.61$)			
Developer-bid value (PRE anchor)	199.0	external calibration	Haines property R^O, PRE-highway anchor
Developer-bid value (POST anchor)	195.0	external calibration	Haines property R^O, observed fringe

Notes. The omitted outside-use component ΔW^O in USD 2025 billions, stock at $r = 0.05$, valued at the developer bid R^O from observed fringe property values (Haines et al., 2018). This is an external calibration, disclosed and not identified by the design, and has an internal PRE→POST sub-range from the property comparator; it is the economically-complete "at most" on the omission. Pressure enters as the $K = 3$ median split $\{\text{None}, \text{Low}, \text{High}\}$ of the 25,000-population city-reach exposure M_c; suitability enters as a first-stage trend control. The figure is reported on both GAEZ surfaces (rainfed-low, $\theta = 2.13$; rainfed-high, $\theta = 1.61$). Ninety-five percent percentile confidence intervals are from the $1{,}000$-iteration stable-county panel bootstrap.

not as an estimate.

The over-identifying restriction and the bound. The pooled two-activity system over-identifies the eight primitives, and the Hansen J tests whether the cross-stratum restrictions the model imposes on the ATT vector are consistent with the homogeneity assumptions. The Hansen J is 14.7 on the rainfed-low surface (four degrees of freedom, $p = 0.0055$) and 9.2 on the rainfed-high surface (four degrees of freedom, $p = 0.056$); the cross-stratum restrictions are rejected on the low surface and borderline on the high surface. On the estimation sample the residual misfit is carried by the high-pressure stratum: dropping either the farmland-acreage or the livestock-revenue moment in that stratum collapses the statistic below conventional rejection thresholds, and the rejection is sharper, not softer, under a bootstrap reference distribution (Appendix B.11, Tables A16 and A15). A rejected over-identification test, where it arises, impugns point identification of the eight primitives, not the lower-bound reading the decomposition commits to: the welfare aggregate is a lower bound on the total land-value gain, and a rejected over-identification leaves that bound intact while denying the eight primitives a precise point. The omitted component is carried by the high-pressure fringe. There the cost and outside-option responses do not separately point-identify – near cities, farmland value is pinned by the development option, so the agricultural signal that would separate operating cost from the outside option is faint. The fringe enters the aggregate through an interval, not a point, and the width of that interval is conditional rather than fitted. The GMM identifies only the acreage sum $(\pi_c + \pi_R)^{\text{High}}$ on the fringe; the c/R split is not identified, so the outside-option response π_R^{High} is pinned only once the sign restrictions of Assumption A5 (Appendix B.11) are imposed. Under A5 the fringe contribution lies in a non-negative interval whose width is conditional on those restrictions rather than fitted, as the sign grid in Appendix B.11, Table A18, records; the identification-robust confidence set for the omission is correspondingly wide (Table A17). The over-identifying restrictions are rejected on the low surface and borderline on the high surface, so the bound the decomposition commits to is assumption-conditional and disclosed as such, not a fitted magnitude the rejected system independently supports.

Summary The farmland value aggregate omits a large, positive welfare component concentrated on settled urban-fringe land, and that concentration – not a sharp dollar figure – is what the decomposition establishes. The standard farmland-value aggregate is a lower bound on the welfare gain in the sense of Donaldson and Hornbeck (2016); valued at observed developer bids, the omission it misses reaches as much as \$195–210 billion across surfaces, and pinning its precise magnitude is the subject of ongoing work. The set-identification structure and the sampling uncertainty are separate statements and should be read as such. The sign restrictions of Assumption A5 (Appendix B.11) guarantee that the fringe contribution to the omission lies in a non-negative interval, so the sign structure of the omission is set-identified regardless of sampling noise. The bootstrap confidence interval is a distinct object: it is wide, and on the conservative surface it spans zero. The magnitude is therefore reported as an anchored upper bound carrying that caveat, not as a precise point estimate: the development-value figure on both GAEZ surfaces is stated in Table 5, with the conservative rainfed-low surface the tightest and most defensible reading and the rainfed-high surface the upper surface rather than a co-equal headline, since pairing the lower tail index with the same system implies more cross-stratum dispersion and a larger, more sharply signed omission. Read across the two surfaces, the omission is large and fringe-concentrated in direction; the standard aggregate bounds the welfare gain from below, and the omission is bounded above, at most, by the externally anchored development value of the converted land.

The omission concentrates in the pressure fringe. The eight primitives hold a single sign across the surface runs. Within a pressure stratum, however, the cost and outside-option responses do not separately identify: they enter the farm-exit cutoff symmetrically, so only their sum $(\pi_c + \pi_R)^\sigma$ is pinned and their bootstrap draws are near-perfectly anti-correlated. We therefore do not read the stratum-by-stratum split of π_R as evidence; the retention-attenuation pattern is carried by the reduced-form acreage and value gradients (Section 6), not by an identified π_R ordering. The reported omission concentrates in the high-pressure fringe, where the pressured strata carry the identified magnitude. The aggregate channel-decomposition magnitudes are reported in Appendix Table A19.

What the GMM does not identify. The two-activity revenue split lets the crop and livestock farm-gate slopes (π_{b_C}, π_{b_L}) enter separately, but the pooled GMM does not decompose the captured value change into activity-specific components. It does not estimate how much of the gain accrues via crop intensification versus livestock retention in the value aggregate; separating those would require activity-specific value moments – for example $\ln V^A_{\text{crops}}$ and $\ln V^A_{\text{livestock}}$ estimated separately – which we do not build here. The four base moments are sufficient for the welfare claim, which is a statement about total agricultural welfare aggregating over composition. The welfare aggregate is computed from the identified revenue sum, not the split into activity-specific returns: the two revenue legs enter the moment system as an accounting identity, but the crop and livestock slopes b_C and b_L are not separately identified in levels, because each activity's per-acre revenue is observed only over its own quality band (Appendix B.16). Nor does a rejected over-identification deliver point-identified primitives; the eight Jacobian slopes are sign-coherent and bounded, but the test denies them a precise point, so we read the aggregate as a decomposition into a captured and an omitted component rather than as a structural point estimate.

Comparison with frontier-expansion settings. Chan (2022) attributes a null specialization response in nineteenth-century US railroads to the frontier expansion mechanism, where new agricultural land entered at the boundary rather than reorganizing within an established margin. The settled-frontier analog documented here operates at the opposite margin: conversion rather than creation, with outside-use returns driving the welfare-versus-output divergence on already-settled land where land markets are deep and the exit margin is real. The two settings can produce similar reduced-form patterns through different mechanisms; their welfare implications differ because in the open-frontier setting, rising farmland values reflect genuine agricultural expansion at the extensive margin, while in a settled-frontier setting where marginal parcels exit agriculture, they reflect outside-use gains that the Agricultural Census does not record.

The natural empirical extensions of this welfare decomposition are calibrated counterfactual exercises that vary the cross-sectional configuration of A_c – for example projected climate-driven shifts in suitability, or comparison with the suitability distributions of other agricultural economies – under the calibrated parameters and the IHS access shock. Such exercises sharpen the decomposition from a level statement to a comparative-statics result and quantify how the agricultural-component fraction of welfare depends on the spatial pattern of suitability. They are reserved for future work.

8 Conclusion

Interstate Highway System access raised average farmland values by about 11 percent, with effects growing to roughly 20 percent at forty years post-construction. Farmland values fell in the high-suitability GAEZ classes and rose in the low-suitability classes, while agricultural acreage moved in the reverse direction: expanding where values fell, contracting where values rose. This pattern is consistent with a strong outside-use channel, as market access raises returns to non-agricultural land use alongside agricultural returns and the outside option competes most readily where agricultural productivity is low. Where outside uses prevailed at the margin, acreage contracted and the parcels that remained in agriculture were, on average, higher quality, which pushed per-acre values up; where outside uses did not prevail, acreage expanded onto lower-quality land, diluting average parcel quality.

The broad-average contraction in low-suitability counties masks a reversal at the very bottom of the distribution. Resolving the four FAO suitability classes as continuous county shares, the marginal class – land not suitable for rainfed crop production – retains its agricultural base rather than abandoning it, through a within-agriculture composition shift toward livestock rather than a contraction. That retention attenuates along a predetermined measure of local development pressure, M_c, the share of a county's convertible 1940 land within reach of its 1940 towns: one access shock leaves three states across the soil ladder. Where pressure is absent, the livestock retention is strongest and the marginal base is held; where it binds, the rising land-rent bar sheds the lowest-value uses, displacing crop acreage into conversion or parking on middle soil while the best soil capitalizes the development option into land value and retains its intensified crop base. The composition shift on marginal land, and its gradient in M_c, reconcile the cell-level evidence with the aggregate. Because the same highway sheds land to development on the pressured fringe while drawing reverted land back into farming on the unpressured margin, the difference-in-differences measures the net acreage change, whereas the welfare object is the gross farm-to-development flow – so the net-acreage read is a conservative lower bound on gross conversion, reinforcing the omission as a bound. In the US context, this conversion margin operated against a backdrop of functioning private land markets, secure parcel-level property rights, and credit access that made conversion financially executable.

We complement the standard Donaldson and Hornbeck (2016) farmland-value welfare aggregate by bounding the welfare it omits. The qualitative result is the headline: the farmland-value aggregate omits a large, positive welfare component concentrated on the high-pressure settled urban fringe. We combine the pressure-stratified treatment effects with a Fréchet productivity model and a pooled minimum-distance GMM that over-identifies eight structural primitives across three development-pressure strata, with yield-based estimates of the land-heterogeneity parameter θ on two GAEZ surfaces. The standard farmland-value aggregate is a lower bound on the welfare gain in the sense of Donaldson and Hornbeck (2016); the omission it misses is a set-identified quantity, bounded above rather than pinned to a point, because the cost and outside-option responses enter the farm-exit cutoff symmetrically and the developer bid that would value converted land in its new use requires an external anchor. Valued at observed developer bids – an external property anchor we disclose as calibration, not identification – the omission reaches as much as \$195–210 billion across surfaces (real 2025 dollars, capitalized at $r = 0.05$), concentrated on the high-pressure settled fringe. A clean off-network comparator for the developer bid does not exist, since interstate placement is collinear with the urban proximity that generates any development premium; the external figure is therefore an anchored upper bound, and the observed fringe premium is flat before and after the highway, so valuing induced conversion at observed prices does not substantially double-count. The set-identification structure and the sampling uncertainty are separate statements. The sign restrictions guarantee that the fringe contribution lies in a non-negative interval, so the direction is set-identified regardless of sampling noise; the bootstrap interval is a distinct object, and it is wide and on the conservative surface spans zero. The over-identification test rejects on the low surface (Hansen $J = 14.7$) and is borderline on the high ($J = 9.2$); the bound is corroborated as a set-identified object rather than resting on the point fit, and pinning its precise magnitude is the subject of ongoing work. The reason it remains a bound is weak identification of the fringe primitives – near cities, farmland value is pinned by the development option, so the agricultural signal that would separate operating cost from the outside option is faint. We read the omission as present, fringe-concentrated, and bounded above, at most, by the externally anchored development value of the converted land.

Two findings travel separately. The suitability gradient itself – high-suitability counties expanding, low-suitability counties contracting on average, with opposite-signed land-value responses, and the marginal-class retention nested beneath it – is the direct empirical signature of comparative-advantage reallocation, and is portable to any context where market access raises agricultural returns differentially across the productivity distribution. The magnitude of the welfare omission depends on the conversion margin being active, which requires functioning land markets, secure parcel-level property rights, and credit access. The direction of the comparative-advantage mechanism generalizes; the magnitude of the welfare correction is conditional on institutional preconditions absent in many development settings.

Across the developing world, governments and multilateral institutions have made rural transport investment a primary instrument of development policy, with the stated goals of expanding market participation and raising agricultural productivity (Gorgulu et al., 2023). The World Bank alone maintains $33.3 billion in active transport commitments across 165 projects, approximately 10 percent of its total lending (Gonzalez-Navarro et al., 2024), and the empirical evidence spans national highway systems and rural road programs across China (Faber, 2014; Banerjee et al., 2020), India (Ghani et al., 2016; Asher and Novosad, 2020; Alder, 2025; Shamdasani, 2021), Indonesia (Gertler et al., 2024), Brazil (Morten and Oliveira, 2024), Ethiopia (Adamopoulos, 2025), and 39 countries in Sub-Saharan Africa (Jedwab et al., 2022). Where the outside-use channel is operative, agricultural outcomes capture only part of the welfare effect, and the part they miss is concentrated where the conversion margin is active – the settled fringe within reach of expanding towns, not the acreage far from development pressure.

Whether the magnitude of the welfare correction this paper estimates for the United States carries over to a developing setting depends on which of those preconditions hold. Property rights at the parcel level are a documented binding constraint: Adamopoulos et al. (2024) find that land insecurity in China was at least as large a barrier to reallocation out of agriculture as all other mobility frictions combined. Credit access and land market depth vary substantially across and within developing economies and are not guaranteed by road construction alone. The empirical exercise we propose for these settings is the suitability-stratified analog of the one this paper implements: stratify outcomes by GAEZ suitability before reaching for a welfare aggregate, and supplement farmland measures with direct non-agricultural-margin observables such as land conversion rates and non-farm employment. The presence of the suitability gradient diagnoses whether the comparative-advantage mechanism is operative; the steepness of the conversion-margin response diagnoses whether the welfare-correction magnitude carries over.

References

Adamopoulos, Tasso, "Spatial Integration and Agricultural Productivity: Quantifying the Impact of New Roads," *American Economic Journal: Macroeconomics*, 2025, *17* (1), 343–78.

— **and Diego Restuccia**, "Geography and Agricultural Productivity: Cross-Country Evidence from Micro Plot-Level Data," *The Review of Economic Studies*, 7 2022, *89* (4), 1629–1653.

— **, Loren Brandt, Chaoran Chen, Diego Restuccia, and Xiaoyun Wei**, "Land Security and Mobility Frictions," *The Quarterly Journal of Economics*, 7 2024, *139* (3), 1941–1987.

Alder, Simon, "Chinese roads in India: The effect of transport infrastructure on economic development," *Journal of International Economics*, 9 2025, *157*, 104140.

Allen, Treb and Costas Arkolakis, "Trade and the Topography of the Spatial Economy," *The Quarterly Journal of Economics*, 8 2014, *129* (3), 1085–1140.

Alonso, William, *Location and Land Use: Toward a General Theory of Land Rent*, Cambridge, MA: Harvard University Press, 1964.

Asher, Sam and Paul Novosad, "Rural Roads and Local Economic Development," *American Economic Review*, 2020, *110* (3), 797–823.

Atack, Jeremy, "Historical Geographic Information Systems (GIS) Database of U.S. Railroads for 1911," 2016.

Banerjee, Abhijit, Esther Duflo, and Nancy Qian, "On the road: Access to transportation infrastructure and economic growth in China," *Journal of Development Economics*, 6 2020, *145*, 102442.

Baum-Snow, Nathaniel, "Did Highways Cause Suburbanization?," *The Quarterly Journal of Economics*, 5 2007, *122* (2), 775–805.

Borusyak, Kirill and Peter Hull, "Nonrandom Exposure to Exogenous Shocks," *Econometrica*, 11 2023, *91* (6), 2155–2185.

— **, Xavier Jaravel, and Jann Spiess**, "Revisiting Event Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Brinkman, Jeffrey and Jeffrey Lin, "Freeway Revolts! The Quality of Life Effects of Highways," *The Review of Economics and Statistics*, 9 2024, *106* (5), 1268–1284.

Brooks, Leah and Zachary Liscow, "Infrastructure Costs," *American Economic Journal: Applied Economics*, 2023, *15* (2), 1–30.

Burchfield, Marcy, Henry G. Overman, Diego Puga, and Matthew A. Turner, "Causes of Sprawl: A Portrait from Space," *The Quarterly Journal of Economics*, 5 2006, *121* (2), 587–633.

Bustos, Paula, Bruno Caprettini, and Jacopo Ponticelli, "Agricultural Productivity and Structural Transformation: Evidence from Brazil," *American Economic Review*, 6 2016, *106* (6), 1320–65.

Butts, Kyle and John Gardner, "did2s: Two-Stage Difference-in-Differences," *R Journal*, 12 2022, *14* (3), 162–173.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Robust Inference With Multiway Clustering," *Journal of Business & Economic Statistics*, 2011, *29* (2), 238–249.

Capozza, Dennis R. and Robert W. Helsley, "The Fundamentals of Land Prices and Urban Growth," *Journal of Urban Economics*, 1989, *26* (3), 295–306.

_ and _ , "The Stochastic City," *Journal of Urban Economics*, 1990, *28* (2), 187–203.

Chan, Jeff, "Farming output, concentration, and market access: Evidence from the 19th-century American railroad expansion," *Journal of Development Economics*, 6 2022, *157*, 102878.

Chandra, Amitabh and Eric Thompson, "Does public infrastructure affect economic activity? Evidence from the rural interstate highway system," *Regional Science and Urban Economics*, 7 2000, *30* (4), 457–490.

Cochrane, Willard W, *The Development of American Agriculture: A Historical Analysis*, Minneapolis, MN: University of Minnesota Press, 1979.

Conley, T. G., "GMM estimation with cross sectional dependence," *Journal of Econometrics*, 9 1999, *92* (1), 1–45.

Costinot, Arnaud and Dave Donaldson, "Ricardo's Theory of Comparative Advantage: Old Idea, New Evidence," *American Economic Review*, 2012, *102* (3), 453–58.

_ , _ , and Cory Smith, "Evolving Comparative Advantage and the Impact of Climate Change in Agricultural Markets: Evidence from 1.7 Million Fields around the World," *Journal of Political Economy*, 2 2016, *124* (1), 205–248.

Deschênes, Olivier and Michael Greenstone, "The Economic Impacts of Climate Change: Evidence from Agricultural Output and Random Fluctuations in Weather," *American Economic Review*, 2007, *97* (1), 354–385.

Dimitri, Carolyn, Anne Effland, and Neilson Conklin, "The 20th Century Transformation of U.S. Agriculture and Farm Policy," *USDA Economic Research Service*, 2005, (EIB-3).

Donaldson, Dave and Richard Hornbeck, "Railroads and American Economic Growth: A "Market Access" Approach," *The Quarterly Journal of Economics*, 5 2016, *131* (2), 799–858.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

Faber, Benjamin, "Trade Integration, Market Size, and Industrialization: Evidence from China's National Trunk Highway System," *Review of Economic Studies*, 2014, *81* (3), 1046–1070.

Fajgelbaum, Pablo D. and Edouard Schaal, "Optimal Transport Networks in Spatial Equilibrium," *Econometrica*, 7 2020, *88* (4), 1411–1452.

Farrokhi, Farid and Heitor S. Pellegrina, "Trade, Technology, and Agricultural Productivity," *https://doi.org/10.1086/724319*, 9 2023, *131* (9), 2509–2555.

Federal Highway Administration, "National Highway System (NHS)," National Transportation Atlas Database, Bureau of Transportation Statistics, U.S. Department of Transportation 2024.

Floyd, John E., "The Effects of Farm Price Supports on the Returns to Land and Labor in Agriculture," *https://doi.org/10.1086/259003*, 4 1965, *73* (2), 148–158.

Frye, Dustin, "Transportation Networks and the Geographic Concentration of Employment," *Review of Economics and Statistics*, 1 2026, pp. 1–11.

Gardner, Bruce, *American agriculture in the twentieth century: How it flourished and what it cost*, Harvard University Press, 2002.

Gardner, John, Neil Thakral, Linh T Tô, and Luther Yap, "Two-Stage Differences in Differences," Technical Report 2024.

Gertler, Paul J., Marco Gonzalez-Navarro, Tadeja Gračner, and Alexander D. Rothenberg, "Road maintenance and local economic development: Evidence from Indonesia's highways," *Journal of Urban Economics*, 9 2024, *143*, 103687.

Ghani, Ejaz, Arti Grover Goswami, and William R Kerr, "Highway to Success: The Impact of the Golden Quadrilateral Project for the Location and Performance of Indian Manufacturing," *Economic Journal*, 2016, *126* (591), 317–357.

Gonzalez-Navarro, Marco, Navarro Roman, David Zarate, Remi Jedwab, and Nick Tsivanidis, "Land Transport Infrastructure," *VoxDev*, 10 2024, *9* (1).

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 12 2021, *225* (2), 254–277.

Gorgulu, Nisan, Vivien Foster, Dhruv Jain, Stéphane Straub, and Maria Vagliasindi, "The Impact of Infrastructure on Development Outcomes: A Meta-Analysis," *World Bank: Policy Research Working Papers*, 3 2023.

Haines, Michael, Price Fishback, and Paul Rhode, "United States Agriculture Data, 1840 - 2012," 2018.

Heblich, Stephan, Stephen J. Redding, and Daniel M. Sturm, "The Making of the Modern Metropolis: Evidence from London," *The Quarterly Journal of Economics*, 11 2020, *135* (4), 2059–2133.

Herzog, Ian, "National transportation networks, market access, and regional economic growth," *Journal of Urban Economics*, 3 2021, *122*, 103316.

Hornbeck, Richard, "Barbed Wire: Property Rights and Agricultural Development," *The Quarterly Journal of Economics*, 5 2010, *125* (2), 767–810.

_ , "The Enduring Impact of the American Dust Bowl: Short- and Long-Run Adjustments to Environmental Catastrophe," *American Economic Review*, 6 2012, *102* (4), 1477–1507.

_ **and Pinar Keskin**, "The Historically Evolving Impact of the Ogallala Aquifer: Agricultural Adaptation to Groundwater and Drought," *American Economic Journal: Applied Economics*, 2014, *6* (1), 190–219.

Jaworski, Taylor, Carl Kitchens, and Sergey Nigai, "Highways and Globalization," *International Economic Review*, 11 2023, *64* (4), 1615–1648.

Jedwab, Rémi, Adam Storeygard, Simon Alder, Treb Allen, Pol Antras, Nathaniel Baum-Snow, Kristian Behrens, Gharad Bryan, Kerem Cosar, Victor Couture, Dave Donaldson, Esther Duflo, Gilles Duranton, Benjamin Faber, Edward Glaeser, Vernon Henderson, Eoin Mcguirk, Melanie Morten, Paul Novosad, Elias Papaioannou, Harris Selod, Matthew Turner, Anthony Venables, Leonard Wantchekon, Yasmin Abisourour, Karen Chen, Rose Choi, Taher Elsheikh, Yury Higuchi, Erin Mcdevitt, and Emily Ryon, "The Average and Heterogeneous Effects of Transportation Investments: Evidence from Sub-Saharan Africa 1960–2010," *Journal of the European Economic Association*, 2 2022, *20* (1), 1–38.

Kantor, Shawn and Alexander Whalley, "Research Proximity and Productivity: Long-Term Evidence from Agriculture," *Journal of Political Economy*, 2019, *127* (2), 819–854.

Karnes, Thomas L, *Asphalt and Politics: A History of the American Highway System*, Jefferson, NC: McFarland, 2009.

Kim, Heejin, John Anders, Craig Carpenter, Marcus Bernard, and Trey Malone, "The Long-Run Effects of Interstate Highways on Rural America," 2025.

Lewis, Joshua and Edson Severnini, "Short- and Long-Run Impacts of Rural Electrification: Evidence from the Historical Rollout of the U.S. Power Grid," *Journal of Development Economics*, 2020, *143*, 102412.

Leyk, Stefan and Johannes H. Uhl, "Data descriptor: HISDAC-US, historical settlement data compilation for the conterminous United States over 200 years," *Scientific Data*, 9 2018, *5* (1), 180175–.

McCoy, Shawn, Xin Zhao, Chris Milner, and Christopher Heiberger, "The Role of Trucking in the U.S. Food Supply Chain," *Journal of Food Distribution Research*, 2020, *51* (2), 17–36.

McShane, Caitlin M., Johannes H. Uhl, and Stefan Leyk, "Gridded Land Use Data for the Conterminous United States 1940–2015," *Scientific Data*, 2022, *9*, 512.

Melichar, Emanuel, "Capital Gains versus Current Income in the Farming Sector," *American Journal of Agricultural Economics*, 1979, *61* (5), 1085–1092.

Mendelsohn, Robert, William D Nordhaus, and Daigee Shaw, "The Impact of Global Warming on Agriculture: A Ricardian Analysis," *The American Economic Review*, 1994, *84* (4), 753–771.

Michaels, Guy, "The Effect of Trade on the Demand for Skill: Evidence from the Interstate Highway System," *The Review of Economics and Statistics*, 11 2008, *90* (4), 683–701.

Mills, Edwin S., "An Aggregative Model of Resource Allocation in a Metropolitan Area," *American Economic Review*, 1967, *57* (2), 197–210.

Morten, Melanie and Jaqueline Oliveira, "The Effects of Roads on Trade and Migration: Evidence from a Planned Capital City," *American Economic Journal: Applied Economics*, 2024, *16* (2), 389–421.

Muth, Richard F., *Cities and Housing: The Spatial Pattern of Urban Residential Land Use*, Chicago: University of Chicago Press, 1969.

Nagy, Dávid Krisztián, "Trade and Urbanization: Evidence from Hungary," *American Economic Journal: Microeconomics*, 8 2022, *14* (3), 733–790.

Olmstead, Alan L. and Paul W. Rhode, "Reshaping the Landscape: The Impact and Diffusion of the Tractor in American Agriculture, 1910–1960," *The Journal of Economic History*, 9 2001, *61* (3), 663–698.

Pellegrina, Heitor S. and Sebastian Sotelo, "Migration, Specialization, and Trade: Evidence from Brazil's March to the West," *https://doi.org/10.1086/738147*, 12 2025, *133* (12), 3993–4049.

Plantinga, Andrew J., Ruben N. Lubowski, and Robert N. Stavins, "The effects of potential land development on agricultural land prices," *Journal of Urban Economics*, 11 2002, *52* (3), 561–581.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *The Review of Economic Studies*, 9 2023, *90* (5), 2555–2591.

Rose, Mark H, *Interstate: Express Highway Politics, 1939–1989*, revised ed., Knoxville, TN: University of Tennessee Press, 1990.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher, "Will U.S. Agriculture Really Benefit from Global Warming? Accounting for Irrigation in the Hedonic Approach," *American Economic Review*, 2005, *95* (1), 395–406.

Schroeder, Jonathan, David Van Riper, Steven Manson, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles, "IPUMS National Historical Geographic Information System: Version 20.0," 2025.

Shamdasani, Yogita, "Rural road infrastructure & agricultural production: Evidence from India," *Journal of Development Economics*, 9 2021, *152* (3), 102686.

Sotelo, Sebastian, "Domestic Trade Frictions and Agriculture," *Journal of Political Economy*, 7 2020, *128* (7), 2690–2738.

Swift, Earl, *The Big Roads: The Untold Story of the Engineers, Visionaries, and Trailblazers Who Created the American Superhighways*, Boston, MA: Houghton Mifflin Harcourt, 2011.

von Thünen, Johann Heinrich, *Der isolierte Staat in Beziehung auf Landwirtschaft und Nationalökonomie*, Hamburg: Perthes, 1826.