

Appendix to
“Welfare Gains from Market Integration Under
Competing Land Uses”

Appendix A Supplementary Exhibits




Figure A1: Counties Excluded from the Analysis

Notes. Map displays counties excluded from the main empirical analysis. Counties are excluded if they contained a Metropolitan Statistical Area in 1960, were a Yellow Book or 1943 Planned City, had an interstate before 1941, or were in the top 10 percent of counties by population.

Appendix A.1 First Differences and First-Difference IV

The reduced-form finding is robust to two alternative estimators that relax the DiD assumptions and identify a related but distinct estimand — the average annual rate of change after IHS access rather than the cumulative level effect. The first-difference design replaces the level outcome with the inter-census log change, with the treatment and controls kept in levels because IHS access is absorbing; it is invariant to the post-treatment window. Formally,

$$\Delta \ln Y_{it} = \beta_0 \text{IHS}_{it} \times \text{NoPressure}_i + \beta_1 \text{IHS}_{it} \times \text{Pressure}_i + \mathbf{X}_i' \gamma_t + \lambda_t + \Delta \varepsilon_{it}, \quad (13)$$

where β_0 and β_1 are the average annual percentage-point changes for no-pressure and under-pressure counties relative to untreated counties, and standard errors are multi-clustered at the county and state-year level. Because β is a slope while the DiD ATT is a cumulative level effect, the two identify compatible but distinct quantities. At the per-period margin the value effects are small and imprecise, but the sign of the development-pressure gradient in farmland is preserved: farmland under development pressure contracts (Table A3).

The first-difference IV design additionally addresses the endogeneity of IHS location and timing. Following [Frye \(2026\)](#), we instrument IHS_{it} with Persh_{it} , a time-varying




Figure A2: Interstate Highway System Construction Progress

Notes. Figure displays the progress of IHS construction in 1960, 1970, and 2002, based on groundbreaking dates (1960, 1970) and opening dates (2002). By 1970, the majority of the eventual IHS route miles were under construction. Counties excluded from the estimation sample are shaded gray.




Figure A3: Distribution of IHS Construction Timing

Notes. Top panels show the distribution of groundbreaking years (left) and opening years (right) across treated counties. Bottom panels show the distribution of years elapsed since groundbreaking (left) and since opening (right) at the time of each Agricultural Census wave. The gap between groundbreaking and opening dates reflects segment construction duration, which varies with route length and terrain.

Table A1: Agricultural Outcomes: Treated and Never-Treated Counties, 1950 and 2002

	Pre-IHS (1950)			Post-IHS (2002)		
	IHS	Non-IHS	Diff.	IHS	Non-IHS	Diff.
Ln(Farms)	7.26 (0.85)	7.06 (0.90)	0.20*** [0.03]	6.39 (0.78)	6.11 (0.88)	0.28*** [0.03]
Farms	1841.07 (1183.20)	1551.94 (1074.75)	289.13*** [45.28]	766.63 (524.57)	607.11 (433.22)	159.52*** [19.49]
Farms over 500 acres	93.99 (130.35)	108.72 (151.21)	-14.73*** [5.49]	111.90 (101.05)	116.83 (104.34)	-4.92 [4.05]
Shr. Farms over 500 acres	10.82 (18.84)	14.16 (22.08)	-3.34*** [0.80]	17.80 (16.63)	22.83 (19.83)	-5.03*** [0.71]
Farmland ('000 acres)	416.49 (534.66)	361.95 (337.12)	54.55*** [18.61]	322.63 (435.02)	302.37 (336.82)	20.26 [15.88]
Cropland ('000 acres)	159.15 (129.73)	153.76 (137.79)	5.38 [5.26]	143.53 (143.63)	146.62 (150.70)	-3.08 [5.80]
Shr. Harvested Acres in Top 4 Crops	0.87 (0.13)	0.89 (0.11)	-0.02*** [0.00]	0.93 (0.09)	0.94 (0.10)	-0.00 [0.00]
Farm Value p.a.	983.81 (710.27)	796.94 (585.75)	186.87*** [26.38]	4011.33 (3427.07)	3019.93 (2685.51)	991.40*** [126.06]
Mkt. Val. Agric. Products p.a.	0.28 (0.23)	0.23 (0.19)	0.05*** [0.01]	0.57 (0.72)	0.52 (1.29)	0.05 [0.04]
Mkt. Val. Crops p.h.a.	0.40 (0.56)	0.41 (3.44)	-0.01 [0.09]	0.71 (1.56)	0.52 (1.04)	0.19*** [0.06]
Mkt. Val. Livestock p.p.a.	0.60 (1.00)	0.49 (1.00)	0.11*** [0.04]	1.95 (4.85)	2.54 (15.20)	-0.59 [0.40]
Farm Labor Expend. per farm	6.76 (10.30)	5.98 (10.75)	0.79* [0.41]	15.02 (28.14)	12.56 (18.13)	2.46** [0.98]
Shr. of Tenant Farms	0.20 (0.14)	0.20 (0.13)	0.01 [0.01]	0.08 (0.07)	0.07 (0.07)	0.00 [0.00]
Observations	1,023	1,694	2,717	1,022	1,694	2,716
						5,433

Notes. Means with standard deviations in parentheses. Diff. columns report the treated-minus-never-treated difference with standard errors in brackets; DiD column reports the change in that difference from 1950 to 2002 estimated by OLS with county and year fixed effects. Monetary values are in nominal dollars per acre. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Covariates Prior to Treatment

	IHS	Non-IHS	Diff.
<i>Panel A: Population and Urbanization</i>			
Urban share, 1950 (%)	29.17 (22.28)	18.93 (21.49)	10.24*** [0.87]
Population, 1950	26923.16 (18394.15)	17344.78 (13430.20)	9578.38*** [661.17]
Population growth, 1940–1950 (%)	0.07 (0.24)	-0.01 (0.21)	0.08*** [0.01]
Distance to nearest MSA (km)	112.85 (106.96)	143.41 (117.26)	-30.56*** [4.39]
<i>Panel B: Market Centrality</i>			
Market potential, 1950	180331.31 (64988.56)	166722.43 (55747.99)	13608.88*** [2441.81]
1947 highway betweenness centrality	4360110.34 (6011514.90)	1963691.47 (3144098.90)	2396418.87*** [202863.62]
<i>Panel C: Alternative Transit Infrastructure</i>			
Railroad track (km), 1950	131.21 (83.39)	99.28 (74.79)	31.92*** [3.18]
Distance to nearest port (km)	394.15 (291.65)	442.52 (294.94)	-48.38*** [11.60]
Distance to nearest airport (km)	51.06 (28.10)	62.70 (32.23)	-11.64*** [1.18]
Distance to inland waterway (km)	147.82 (205.56)	158.22 (204.63)	-10.40 [8.13]
<i>Panel D: Geography</i>			
Latitude (degrees)	38.00 (4.74)	38.34 (4.99)	-0.35* [0.19]
Longitude (degrees)	-91.54 (11.37)	-92.80 (10.60)	1.26*** [0.44]
County area (km ²)	2705.87 (3975.56)	2408.03 (2881.35)	297.84** [142.65]
Avg. elevation (m)	444.94 (515.38)	470.60 (525.36)	-25.66 [20.56]
Terrain ruggedness (m)	76.66 (113.32)	71.93 (108.17)	4.73 [4.41]
Observations	1,023	1,694	2,717

Notes. Means with standard deviations in parentheses. Diff. reports the treated-minus-never-treated difference with standard errors in brackets. Market potential is the [Donaldson and Hornbeck \(2016\)](#) county-level access measure evaluated at 1950 population weights. Highway betweenness centrality is computed from the 1947 Rand McNally national road network ([Frye, 2026](#)). Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure A4: Farm Value per Acre and Agricultural Product Value per Acre, 1940–2002

Notes. Both series are indexed to 1940 = 1. Farm value per acre capitalizes expected future returns to agricultural production. Agricultural product value per acre measures realized revenue from crops and livestock. Values are deflated to 2025 dollars. Sample restricted to rural counties in the estimation sample.



(a) Farm value per acre

(b) Farmland acres

Figure A5: Trends by Land Suitability, Indexed to 1940

Notes. Each panel plots the group-mean outcome indexed to 1940 = 1, separately for counties in the top 40 percent by GAEZ suitability (Suitable, black) and the bottom 60 percent (Unsuitable, gray), the coarse binary suitability split used for the aggregate-robustness battery. These are unconditional group means.




Figure A6: Interstate Highway Counties and Adjacent Counties within 75 km

Notes. Maps display counties that IHS segments intersect (dark blue) and counties within 75 km of an IHS county (light blue) in 1960 and 2002. Treatment timing is defined by the break-ground date. Counties excluded from the sample are shaded gray.


indicator equal to one once county i 's Pershing Map segment is predicted complete under a 25-year Newman-Girvan budget horizon; its location is fixed by the 1922 Pershing Map and its timing is set by network betweenness rather than fiscal or political sequencing. The instrument is strong, with Kleibergen-Paap F-statistics between 22 and 54 across specifications, and the IV estimates again leave the per-period value effects imprecise while farmland under development pressure contracts (Table A4, Panel B: -0.007 to -0.010 per year in the preferred columns).

Appendix A.2 Robustness of the Suitability Gradient


The same pattern of divergence holds across model specifications. Table A7 displays the coefficients from the same heterogeneous treatment effect model, where “Suitable” is defined as the top 40 percent of counties according to the FAO suitability index and “Unsuitable” is the bottom 60 percent of counties.¹¹ In Panel A, the dependent variable is logged farmland value per acre. In Panel B, it is logged farmland acres. Column 1 of both panels report the coefficients from models with only county and year fixed effects and using the full sample of counties in the US. Column 2 restricts the sample according to the qualifications described in Section 5. Column 3 returns to the full sample, but adds the full set of control variables. Column 4, which is our preferred specification, adds the control variables to the restricted sample. Finally, Column 5 weights both regressions by 1950 farmland area.¹²

¹¹We collapse to two classifications for ease of interpretability. In addition to the divergence in treatment effects between the suitable and less-suitable classes in Figure 4, this definition of suitable also corresponds to the biophysical suitability classes of Good, High, and Very High (SI > 5500).


¹²Weighting in these two cases moves the estimand closer to the primary underlying unit: the average effect across farmland acres. However, we prefer the unweighted regressions in column 4, because the estimates are retrieved from the exact same model as other outcomes that we study where weighting by



(a) Number of Farms



(b) Number of Farms >500 acres



(c) Farmland Acres

Direct (IHS first)

Adjacent (first, t75)

Table A3: First-Difference Estimates

Panel A: Log Farmland Value (\$/acre)						
	(1)	(2)	(3)	(4)	(5)	(6)
IHS $\times$ No pressure	0.001 (0.001)	0.001 (0.001)	0.002*** (0.001)	0.002*** (0.001)	0.003*** (0.001)	0.003*** (0.001)
IHS $\times$ Under pressure	-0.000 (0.001)	0.001 (0.001)	0.002 (0.001)	0.002* (0.001)	0.001 (0.001)	0.001 (0.001)
Num.Obs.	39292	35068	39292	35068	39292	35068
Panel B: Log Farmland (acres)						
	(1)	(2)	(3)	(4)	(5)	(6)
IHS $\times$ No pressure	-0.000 (0.001)	-0.000 (0.001)	-0.001 (0.001)	-0.001* (0.000)	-0.001** (0.000)	-0.001** (0.000)
IHS $\times$ Under pressure	-0.003* (0.002)	0.000 (0.002)	-0.001 (0.001)	0.000 (0.001)	-0.001 (0.001)	-0.000 (0.001)
Num.Obs.	39369	35123	39369	35123	39369	35123
Controls			✓	✓	✓	✓
Weights					✓	✓
Sample	Full	Restricted	Full	Restricted	Full	Restricted

Notes. First-difference estimates with year fixed effects. “No pressure” (“Under pressure”) indicates counties with zero (positive) development pressure at the 25,000-population cutoff. Column 1: full sample, no controls. Column 2: restricted sample, no controls. Column 3: full sample with controls. Column 4: restricted sample with controls (preferred). Column 5: Column 4 weighted by 1950 farmland area. Coefficients are average annual percentage-point changes relative to untreated counties. [Conley \(1999\)](#) standard errors are in parentheses and adjust for spatial correlation with a 100km cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: First-Difference IV Estimates

Panel A: Log Farmland Value (\$/acre)						
	(1)	(2)	(3)	(4)	(5)	(6)
IHS $\times$ No pressure	-0.002 (0.004)	-0.003 (0.004)	0.002 (0.004)	0.002 (0.004)	0.005 (0.004)	0.004 (0.004)
IHS $\times$ Under pressure	-0.004 (0.005)	-0.005 (0.005)	0.000 (0.006)	-0.003 (0.005)	-0.003 (0.006)	-0.002 (0.005)
Num.Obs.	39292	35068	39292	35068	39292	35068
KP F-stat	54.45	27.74	49.19	39.14	33.32	22.52
Panel B: Log Farmland (acres)						
	(1)	(2)	(3)	(4)	(5)	(6)
IHS $\times$ No pressure	0.001 (0.002)	0.001 (0.002)	-0.000 (0.002)	-0.000 (0.002)	-0.001 (0.001)	-0.001 (0.001)
IHS $\times$ Under pressure	-0.010*** (0.003)	-0.000 (0.003)	-0.007* (0.004)	-0.001 (0.003)	-0.002 (0.003)	0.000 (0.003)
Num.Obs.	39369	35123	39369	35123	39369	35123
KP F-stat	0.00	28.07	50.50	40.28	33.48	22.61
Controls			✓	✓	✓	✓
Weights					✓	✓
Sample	Full	Restricted	Full	Restricted	Full	Restricted

Notes. First-difference IV estimates. IHS treatment instrumented by $Persh25_{it}$, a time-varying indicator equal to one once the county's Pershing Map segment is predicted complete under a 25-year Newman-Girvan budget horizon (Frye, 2026), interacted with the development-pressure strata. "No pressure" ("Under pressure") indicates counties with zero (positive) development pressure at the 25,000-population cut-off. Column definitions as in Table A3. KP F-stat is the Kleibergen-Paap weak-instrument F-statistic. Standard errors multi-clustered at county and state-year in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Impact of IHS on Farmland Values: Heterogeneity

	log(Farmland Value)			
	(1)	(2)	(3)	(4)
IHS × N (not suitable)	0.111** (0.053)	0.047 (0.084)	−0.076 (0.072)	0.065 (0.056)
IHS × S3	−0.024 (0.040)	0.062 (0.058)	0.067 (0.057)	0.086 (0.089)
IHS × S2S1 (suitable)	0.114*** (0.043)	0.094** (0.047)	0.094** (0.041)	0.013 (0.037)
IHS × N (not suitable) × Suburb	−0.040 (0.080)			
IHS × S3 × Suburb	0.135** (0.060)			
IHS × S2S1 (suitable) × Suburb	0.061* (0.032)			
IHS × N (not suitable) × High Other Roads		0.095* (0.055)		
IHS × S3 × High Other Roads		0.047 (0.039)		
IHS × S2S1 (suitable) × High Other Roads		0.077** (0.031)		
IHS × N (not suitable) × High Rail			0.166*** (0.051)	
IHS × S3 × High Rail			0.040 (0.038)	
IHS × S2S1 (suitable) × High Rail			0.074** (0.032)	
IHS × N (not suitable) × Waterway				0.113 (0.095)
IHS × S3 × Waterway				0.034 (0.037)
IHS × S2S1 (suitable) × Waterway				0.139*** (0.040)
Num.Obs.	37 852	37 852	37 852	37 852

Notes. Two-stage DiD estimates ([Gardner et al., 2024](#)) interacting the IHS treatment with the county's FAO suitability-class shares (marginal N, S3, and suitable S2+S1) and, in turn, with four proximity and pre-existing-transit dimensions, one per column. Metropolitan proximity is defined as a county centroid within 75 km of a 1960 Metropolitan Statistical Area centroid. Other Roads is an indicator for above-median 1947 highway betweenness. Rail is an indicator for above-median railroad density. Waterway is an indicator for a county centroid within 75 km of a navigable waterway. All specifications include county and year fixed effects and the full control vector interacted with year indicators. [Conley \(1999\)](#) standard errors in parentheses, 100 km cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Impact of IHS on Farmland Acres: Heterogeneity by Suitability

	log(Farmland)			
	(1)	(2)	(3)	(4)
IHS × N (not suitable)	0.170*** (0.051)	0.010 (0.098)	0.177** (0.083)	0.251*** (0.048)
IHS × S3	0.017 (0.030)	−0.064 (0.044)	−0.059 (0.049)	−0.064 (0.076)
IHS × S2S1 (suitable)	−0.023 (0.025)	0.050* (0.025)	−0.001 (0.027)	0.031 (0.030)
IHS × N (not suitable) × Suburb	0.214** (0.091)			
IHS × S3 × Suburb	−0.105* (0.059)			
IHS × S2S1 (suitable) × Suburb	0.066** (0.028)			
IHS × N (not suitable) × High Other Roads		0.250*** (0.044)		
IHS × S3 × High Other Roads		−0.027 (0.041)		
IHS × S2S1 (suitable) × High Other Roads		0.026 (0.026)		
IHS × N (not suitable) × High Rail			0.183*** (0.049)	
IHS × S3 × High Rail			−0.021 (0.034)	
IHS × S2S1 (suitable) × High Rail			0.046* (0.026)	
IHS × N (not suitable) × Waterway				0.005 (0.097)
IHS × S3 × Waterway				−0.027 (0.035)
IHS × S2S1 (suitable) × Waterway				0.034 (0.030)
Num.Obs.	37 888	37 888	37 888	37 888

Notes. As in Table A5, for log farmland acres: the IHS treatment interacted with the FAO suitability-class shares (marginal N, S3, and suitable S2+S1) and with each of the four proximity and pre-existing-transit dimensions, one per column. Heterogeneity indicators defined as in Table A5. Conley (1999) standard errors in parentheses, 100 km cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Impact of IHS on Farmland Values and Farmland by Land Suitability

Panel A					
	ln(Farmland Value (\$/acre))				
	(1)	(2)	(3)	(4)	(5)
IHS × Unsuitable	0.115*** (0.030)	0.153*** (0.033)	0.201*** (0.029)	0.193*** (0.033)	0.242*** (0.041)
IHS × Suitable	−0.174*** (0.028)	−0.160*** (0.032)	−0.025 (0.027)	−0.065** (0.030)	−0.051* (0.030)
Num.Obs.	42 440	37 852	42 440	37 852	37 852
Panel B					
	ln(Farmland (acres))				
	(1)	(2)	(3)	(4)	(5)
IHS × Unsuitable	−0.172*** (0.045)	−0.095** (0.040)	−0.153*** (0.030)	−0.109*** (0.031)	−0.093*** (0.021)
IHS × Suitable	0.040 (0.029)	0.117*** (0.018)	0.097*** (0.022)	0.155*** (0.020)	0.104*** (0.016)
Num.Obs.	42 492	37 888	42 492	37 888	37 888
Controls			✓	✓	✓
Weights					✓
Sample	Full	Restricted	Full	Restricted	Restricted

Notes. Two-stage DiD estimates (Gardner et al., 2024). “Suitable” is the top 40 percent of counties on the FAO rainfed suitability index and “Unsuitable” the bottom 60 percent. Column 1: full sample, no controls. Column 2: restricted sample, no controls. Column 3: full sample with controls. Column 4: restricted sample with controls (preferred). Column 5: Column 4 weighted by 1950 farmland area. Conley (1999) standard errors are in parentheses and adjust for spatial correlation with a 100km cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Impact of IHS on Farmland Values and Farmland: Continuous Development Pressure

Panel A						
	ln(Farmland Value (\$/acre))					
	(1)	(2)	(3)	(4)	(5)	(6)
IHS	0.031 (0.028)	0.027 (0.028)	0.091*** (0.025)	0.086*** (0.025)	0.116*** (0.027)	0.114*** (0.027)
IHS × Dev. Pressure	−0.060 (0.041)	0.063 (0.071)	0.079** (0.036)	0.033 (0.069)	0.113** (0.050)	−0.070 (0.078)
Num.Obs.	42 441	37 853	42 441	37 853	42 440	37 852
Panel B						
	ln(Farmland (acres))					
	(1)	(2)	(3)	(4)	(5)	(6)
IHS	0.078*** (0.026)	0.050* (0.026)	0.021 (0.022)	0.006 (0.021)	−0.019 (0.015)	−0.022 (0.015)
IHS × Dev. Pressure	−0.392*** (0.090)	−0.128 (0.085)	−0.262*** (0.059)	0.004 (0.087)	−0.118*** (0.041)	0.042 (0.067)
Num.Obs.	42 494	37 890	42 494	37 890	42 492	37 888
Controls			✓	✓	✓	✓
Weights					✓	✓
Sample	Full	Restricted	Full	Restricted	Full	Restricted
Panel C						
	ln(Total Farmland Value)					
	(1)	(2)	(3)	(4)	(5)	(6)
IHS	0.105*** (0.025)	0.076*** (0.024)	0.109*** (0.022)	0.091*** (0.023)	0.097*** (0.024)	0.092*** (0.025)
IHS × Dev. Pressure	−0.440*** (0.075)	−0.067 (0.058)	−0.177*** (0.062)	0.032 (0.056)	−0.005 (0.056)	−0.030 (0.056)
Num.Obs.	42 441	37 853	42 441	37 853	42 440	37 852
Controls			✓	✓	✓	✓
Weights					✓	✓
Sample	Full	Restricted	Full	Restricted	Full	Restricted

Notes. Two-stage DiD estimates ([Gardner et al., 2024](#)). Development pressure is a continuous measure of urban development in 1950. Column 1: full sample, no controls. Column 2: restricted sample, no controls. Column 3: full sample with controls. Column 4: restricted sample with controls (preferred). Column 5: Column 4 weighted by 1950 farmland area. [Conley \(1999\)](#) standard errors are in parentheses and adjust for spatial correlation with a 100km cutoff. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.




Figure A8: [Rambachan and Roth \(2023\)](#) Sensitivity Test on Farmland Values

Note: This figure reports sensitivity analysis following [Rambachan and Roth \(2023\)](#). The orange interval is the baseline 95% confidence interval for the farmland value ATT at 40 years post-treatment, estimated under the standard parallel trends assumption. The blue intervals (C-LF) show how the confidence set expands as we relax parallel trends by allowing post-treatment violations of increasing magnitude. The horizontal axis $\bar{M}$ parameterizes the maximum post-period deviation from linearity, expressed as a multiple of the largest pre-period absolute deviation observed in the event study.


Appendix A.3 Suitability surfaces and the heterogeneity battery

This section reports the full cell-stratified battery behind the headline heterogeneity of Section 6.1. Each table breaks the IHS effect out by the three identification classes (marginal N, S3, and the suitable class S2+S1) interacted with the binary development-pressure split, and reports every cell on both GAEZ parameterizations — the rainfed-low and rainfed-high suitability surfaces — so that the sensitivity of each margin to the surface choice is visible in one place. The three tables cover farm survival and the extensive margin, land value and output, and within-agriculture composition in turn.


Appendix A.4 Identifying Support Across Suitability Cells

Every cell-stratified estimate in the paper is identified off the treated counties that carry the relevant suitability share. The estimator is a share-weighted Gardner DID2s, in which a suitability $\times$ pressure cell coefficient is recovered from those treated counties weighted by their (aggregated) FAO class share; the effective number of identifying counties is the Kish effective sample size $(\sum w)^2 / \sum w^2$ of these share weights. This diagnostic governs the reduced-form heterogeneity battery and the continuous dose-response, and it governs the welfare GMM, whose twelve moments are the same cell ATTs. Cluster-robust inference is outside its reliable range below roughly forty


farmland acres would be inappropriate or in different units (e.g., # of farms).



(a) Crop Revenue (\$ per cropland acre)



(b) Livestock Revenue (\$ per pastureland acre)



(c) Farm Labor Expenditure (\$ per farm)

Direct (IHS first)

Adjacent (first, t75)




Figure A10: County Area Shares in the Four FAO Suitability Classes

Notes. Each panel maps the share of a county's area in one FAO GAEZ suitability class — N (marginal), S3, S2, and S1 (prime) — on the rainfed-low input surface, the primary specification. Shares across the four classes sum to one. Analysis sample (KeepSet); excluded counties in grey.




Figure A11: 1947 Highway Network Betweenness Centrality

Notes. Map displays the quintile of total betweenness centrality for each county in the 1947 Rand McNally national highway network. Counties excluded from the analysis sample are shaded gray.

Town pop. cutoff $\geq 2,500$ – CPP towns




Figure A12: The 1940 Census-Place Layer Used to Seed and Validate the City Blocks

Notes. The georeferenced incorporated places of the full-count population census, sized by 1940 population; presence is read from the census rather than from the built-up raster, so the layer is complete even where the raster under-records small towns. The map shows all places, not the KeepSet-filtered analysis sample.




Figure A13: Interstate Highways and Farmland Values: Alternative Clustering

Notes. Figure displays coefficients from an event study regression estimated with `did2s`. The gray shaded area reflects 95 percent confidence intervals adjusted for spatial correlation within 100 km, identical to the preferred primary specification in Figure 2. The solid line displays 95 percent confidence intervals with standard errors clustered at the county level. The dashed line displays 95 percent confidence intervals multi-clustered at the county and state-year levels.

Table A9: Full heterogeneity battery: farm survival and the extensive margin.

Dependent variable:	Farmland (1)	Cropland (2)	Harv. crop (3)	Pasture (4)	Farms 500+ (5)
<i>Panel A: Rainfed-low suitability</i>					
N (marginal) $\times$ none	0.17*** (0.05)	0.12* (0.06)	0.11 (0.09)	0.18*** (0.06)	0.07 (0.10)
N (marginal) $\times$ pressure	-0.05 (0.09)	-0.10 (0.08)	-0.03 (0.10)	-0.14 (0.12)	-0.31*** (0.09)
S3 $\times$ none	-0.01 (0.03)	-0.01 (0.04)	-0.04 (0.06)	0.04 (0.05)	0.00 (0.05)
S3 $\times$ pressure	-0.13*** (0.05)	-0.16** (0.07)	-0.10* (0.08)	0.01 (0.06)	-0.50*** (0.08)
S2+S1 (suitable) $\times$ none	0.00 (0.03)	-0.02 (0.03)	-0.04 (0.05)	-0.01 (0.05)	-0.03 (0.08)
S2+S1 (suitable) $\times$ pressure	0.01 (0.03)	0.01 (0.03)	0.09 (0.04)	-0.18*** (0.05)	-0.10* (0.07)
<i>Panel B: Rainfed-high suitability</i>					
N (marginal) $\times$ none	0.10*** (0.03)	0.07* (0.04)	0.02 (0.06)	0.14*** (0.03)	0.02 (0.06)
N (marginal) $\times$ pressure	-0.06 (0.06)	-0.13 (0.06)	-0.09 (0.08)	-0.15 (0.12)	-0.33*** (0.10)
S3 $\times$ none	0.02 (0.05)	0.08 (0.06)	0.05 (0.06)	0.03 (0.07)	-0.04 (0.07)
S3 $\times$ pressure	0.03 (0.07)	0.01 (0.07)	0.03 (0.09)	0.08 (0.09)	-0.27*** (0.10)
S2+S1 (suitable) $\times$ none	-0.06 (0.04)	-0.08* (0.04)	-0.12* (0.06)	0.01 (0.05)	-0.11 (0.07)
S2+S1 (suitable) $\times$ pressure	0.01 (0.05)	0.00 (0.05)	0.10 (0.07)	-0.07 (0.05)	-0.16*** (0.09)

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 $M_c^b = \mathbf{1}\{M_c > \text{med}(M_c)\}$ 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). Orchard, vegetable, and perishable acreage use the pre-1978 window (1940–1974). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Full heterogeneity battery: land value and output.

Dependent variable:	Total value (1)	Value/acre (2)	Crop rev. (3)	Lvstk. rev. (4)
<i>Panel A: Rainfed-low suitability</i>				
N (marginal) $\times$ none	0.25*** (0.06)	0.09 (0.06)	0.12 (0.12)	0.28*** (0.09)
N (marginal) $\times$ pressure	-0.10* (0.08)	-0.06 (0.07)	0.16** (0.12)	-0.06 (0.14)
S3 $\times$ none	0.01 (0.04)	0.02 (0.04)	-0.08 (0.09)	-0.04 (0.09)
S3 $\times$ pressure	-0.09* (0.04)	0.03 (0.05)	0.16 (0.08)	-0.59*** (0.13)
S2+S1 (suitable) $\times$ none	0.07 (0.04)	0.07* (0.04)	-0.11 (0.09)	-0.05 (0.05)
S2+S1 (suitable) $\times$ pressure	0.05*** (0.03)	0.04*** (0.03)	0.26** (0.07)	-0.30*** (0.06)
<i>Panel B: Rainfed-high suitability</i>				
N (marginal) $\times$ none	0.16*** (0.04)	0.06 (0.05)	-0.03 (0.09)	0.14* (0.08)
N (marginal) $\times$ pressure	0.13*** (0.07)	0.18*** (0.07)	0.12 (0.10)	0.17* (0.16)
S3 $\times$ none	0.10 (0.07)	0.09 (0.06)	-0.01 (0.09)	0.07 (0.09)
S3 $\times$ pressure	0.05* (0.08)	0.03 (0.08)	0.14 (0.11)	-0.05 (0.12)
S2+S1 (suitable) $\times$ none	0.03 (0.04)	0.09** (0.04)	-0.17* (0.10)	-0.03 (0.08)
S2+S1 (suitable) $\times$ pressure	0.03** (0.03)	0.02*** (0.04)	0.27 (0.07)	-0.37*** (0.07)

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 $M_c^b = 1\{M_c > \text{med}(M_c)\}$ 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). Orchard, vegetable, and perishable acreage use the pre-1978 window (1940–1974). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A11: Full heterogeneity battery: within-agriculture composition.

Dependent variable:	Cattle (1)	Top4 harv (2)	Top4 tot (3)	Orchard (4)	Veg. (5)	Perish. (6)	Tenant sh (7)
<i>Panel A: Rainfed-low suitability</i>							
N (marginal) $\times$ none	0.32*** (0.07)	−0.01 (0.02)	0.00 (0.01)	0.52** (0.26)	0.83*** (0.24)	0.69*** (0.22)	0.01 (0.01)
N (marginal) $\times$ pressure	−0.10** (0.11)	−0.02 (0.02)	−0.02 (0.01)	−0.85 (0.32)	−0.65 (0.29)	−1.31** (0.26)	0.01 (0.02)
S3 $\times$ none	0.04 (0.05)	0.01 (0.01)	0.00 (0.01)	−0.07 (0.09)	−0.12 (0.11)	−0.09 (0.08)	0.00 (0.01)
S3 $\times$ pressure	−0.23*** (0.06)	0.00 (0.01)	0.00 (0.01)	0.06 (0.11)	0.11 (0.12)	−0.06 (0.10)	0.06*** (0.02)
S2+S1 (suitable) $\times$ none	−0.07 (0.06)	0.00 (0.01)	0.01 (0.01)	0.03 (0.11)	−0.11 (0.14)	−0.02 (0.12)	−0.02* (0.01)
S2+S1 (suitable) $\times$ pressure	−0.22*** (0.06)	0.03*** (0.01)	−0.01 (0.01)	−0.07 (0.09)	0.34* (0.12)	0.04 (0.09)	0.05*** (0.01)
<i>Panel B: Rainfed-high suitability</i>							
N (marginal) $\times$ none	0.14** (0.06)	−0.02 (0.01)	0.00 (0.01)	0.33 (0.22)	0.50** (0.20)	0.37** (0.18)	0.02** (0.01)
N (marginal) $\times$ pressure	−0.05 (0.12)	0.00 (0.02)	−0.01 (0.01)	−0.73 (0.28)	−0.26 (0.24)	−0.82** (0.21)	0.00 (0.01)
S3 $\times$ none	0.10 (0.08)	0.03 (0.02)	0.00 (0.01)	0.16 (0.17)	0.00 (0.16)	0.10 (0.14)	0.01 (0.01)
S3 $\times$ pressure	0.04* (0.08)	0.00 (0.02)	−0.01 (0.02)	−0.14 (0.19)	−0.14 (0.22)	−0.21 (0.21)	0.03*** (0.01)
S2+S1 (suitable) $\times$ none	0.02 (0.05)	0.00 (0.01)	0.00 (0.01)	−0.10 (0.08)	−0.08 (0.10)	−0.08 (0.09)	−0.01 (0.01)
S2+S1 (suitable) $\times$ pressure	−0.25*** (0.05)	0.01** (0.01)	0.00 (0.01)	0.07 (0.08)	0.22 (0.09)	0.01 (0.08)	0.04*** (0.01)

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 $M_c^b = 1\{M_c > \text{med}(M_c)\}$ 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). Orchard, vegetable, and perishable acreage use the pre-1978 window (1940–1974). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

effective counties; the counts below are the Kish effective count of treated counties behind each share-weighted cell, computed under three pressure structures on both GAEZ surfaces (the suitability-only estimator, the binary median split, and a finer three-way split).

Under the suitability-only pooling and the binary median split, every suitability class clears the reliability bar on both surfaces: the smallest cell is the prime (S1) class away from pressure on the rainfed-low surface, at 61 effective counties. Identifying support fails only for the prime class under a finer three-way pressure split, where the S1 none-pressure cell falls to 8 effective counties on the rainfed-low surface and 20 on rainfed-high. The prime class also carries the least total share mass — 181 effective counties under the suitability-only pooling on rainfed-low, the smallest of the four classes — so a standalone S1 share-coefficient extrapolates toward the fully-prime county off the thinnest support.

We therefore carry the three-class partition — marginal (N), S3, and the suitable class (S2+S1) — interacted with the binary median pressure split, a structure in which every cell clears the bar on both surfaces: the merged suitable cell holds at least 97 effective counties under all three pressure structures. Merging the prime grade into the suitable class is what removes the exposure; breaking S1 out on its own would rest the finer contrasts on its thin support. This structure supplies the cell ATTs the welfare GMM uses as moments. The outside-option channel is identified, as throughout, off the cross-suitability differential and the observed development response rather than off a continuous within-class pressure gradient.

Appendix B Theoretical Framework: Full Derivations

This appendix presents the complete theoretical framework underlying the predictions in Section 3. The appendix is self-contained: a reader can follow the argument without referring back to the main text, though we note where main-text equations appear in reduced form.

Appendix B.1 Environment

Counties are indexed by c , land parcels by j , and agricultural census years by t . We summarize county market access through an *access wedge* $\tau_{ct} \in \mathbb{R}_+$, where higher values indicate worse access.

An interstate highway opening reduces τ_{ct} . We write:

$$\tau_{ct} = \bar{\tau}_{ct} - \Delta\tau \cdot \text{Open}_{ct} - \lambda \Delta\tau \cdot \text{Adj}_{ct}, \quad 0 < \lambda < 1, \quad (14)$$

where $\bar{\tau}_{ct}$ is the baseline access wedge absent interstate treatment; $\text{Open}_{ct} \in \{0, 1\}$ indicates direct interstate access by year t ; $\text{Adj}_{ct} \in \{0, 1\}$ indicates adjacency to a directly treated county; $\Delta\tau > 0$ is the magnitude of the access improvement from direct treatment; and $\lambda \in (0, 1)$ captures the attenuated shock received by adjacent counties.

Equation (14) is a reduced-form representation of how direct and nearby exposure shift the local access wedge. It is sufficient to organize the comparative statics; it is not a full network market-access model.

Appendix B.2 The Bundled Access Shock: Where τ Enters

An interstate opening does not change a single price. It simultaneously shifts a bundle of local prices and costs, and it raises the value of land in non-agricultural use. We identify the channels through which τ_{ct} enters the framework, now with two farm technologies and two non-farm outside options.

Revenue channel: $b_C(\tau_{ct})$ and $b_L(\tau_{ct})$

Better access connects farms to more distant buyers, raising farm-gate prices net of marketing and delivery costs, and lets rival sellers into the markets the county serves, lowering them. The effective return per unit of output is defined separately for each farm activity $k \in \{C, L\}$ (crops, livestock):

$$b_k(\tau_{ct}) = P_{k,t} + \beta_k(\tau_{ct}) - g_k(\tau_{ct}), \quad (15)$$

where $P_{k,t}$ is an aggregate activity price common across counties, $\beta_k(\tau_{ct})$ captures local price premia from buyer access, and $g_k(\tau_{ct})$ is the per-unit delivery and marketing cost. Reach raises β_k and competition lowers it (through $P_{k,t}$ and g_k); under Assumption A1 the net sign of $\partial b_k / \partial \tau$ is left to the data, separately for each activity, with no ordering imposed between crops and livestock. The revenue channel scales with productivity: the

gain from any access improvement is proportional to output, so within an activity it shifts the farm-gate return common across a county's parcels.

Cost channel: $c_C(\tau_{ct})$ and $c_L(\tau_{ct})$

Market access changes factor costs in two opposing directions:

- *Input prices.* Delivered prices for fertilizer, seed, fuel, and equipment fall as better access connects farmers to national suppliers. Define $p^M(\tau_{ct})$ as the delivered materials price, with $\partial p^M / \partial \tau > 0$. Lower input prices reduce per-acre costs.
- *Wages and land.* Local wages and land prices rise as highway access expands non-farm employment and residential demand and strengthens workers' outside options. Define $w(\tau_{ct})$ as the local factor price, with $\partial w / \partial \tau < 0$. Higher factor prices raise per-acre costs.

We collect these effects into a net per-acre cost for each activity:

$$c_k(\tau_{ct}) = w(\tau_{ct}) \ell_k^* + p^M(\tau_{ct}) m_k^* + r_K(\tau_{ct}) k_k^*, \quad (16)$$

where ℓ_k^* , m_k^* , and k_k^* are activity-specific optimal per-acre factor demands and $r_K(\tau_{ct})$ is the local cost of capital services. The sign of $\partial c_k / \partial \tau$ is ambiguous and, under Assumption A1, left to the data. The cost *response* is taken common across activities, so the crop–livestock cost gap $c_C - c_L$ is access-invariant even as the levels move; the *level* ordering $c_C > c_L$ is part of Assumption A3. The composition cutoff $\phi_{CL} = (c_C - c_L) / (b_C - b_L)$ of the main text can still move with access, because A1 declines to order the price legs and a wedge between $\partial b_C / \partial \tau$ and $\partial b_L / \partial \tau$ shifts the denominator; its movement is left to the data and signs no marginal-band crop prediction (the density argument below confines any ϕ_{CL} shift to good land). Neither cost leg scales with parcel productivity ω_{je} : the cost baselines are common across a county's parcels within an activity at a given time.

Outside-option channel: development $R_{ct}^O(\tau_{ct})$ and reversion $R_{ct}^N(\tau_{ct})$

Better market access raises the value of land in non-agricultural use, and a parcel that leaves farming faces one of two options set by its location. In the conversion fringe a builder bids the developed-use return $R_{ct}^O(\tau_{ct})$, with $\partial R_{ct}^O / \partial \tau < 0$: as commuting costs fall, residential and commercial demand raises the bid, read as a capitalized development *option* (land built, or held vacant in anticipation of development). Beyond the fringe no builder bids, and land that leaves farming reverts to a positive non-development use $R_{ct}^N(\tau_{ct})$ – low-intensity conservation or idle reversion (Soil Bank, unprogrammed reversion, and from 1985 the Conservation Reserve Program) – with $R_{ct}^O > R_{ct}^N \geq 0$.¹³ A2 fixes this two-regime, proximity-driven structure but takes no stand on whether the bid's response tilts on soil within a location. The benchmark against which the data are read is

¹³The reversion wedge κ_N is approximately zero, since conservation enrollment is typically subsidized; R^N is the net program or idle payment. Commercial forestry is left out of R^N : timber is soil-sensitive and competes for good sites.

the no-tilt case, in which what a developer or a conservation program will pay responds to access by location alone:

$$\frac{\partial^2 R_{ct}^O}{\partial \omega_{jc} \partial \tau} = \frac{\partial^2 R_{ct}^N}{\partial \omega_{jc} \partial \tau} = 0. \quad (17)$$

Proximity orders outside-use responsiveness; the soil tilt is measured. Under the no-tilt benchmark of (17) the outside option's response to access does not depend on a parcel's farm productivity ω_{jc} : what a developer will pay for the land, and how sharply that bid rises once a highway shortens commutes and brings housing and commercial demand, is set by where the parcel sits relative to a growing city. A parcel inside the conversion fringe sees its outside option rise sharply with access; a parcel well beyond the fringe barely stirs. This proximity structure is the content of Assumption A2 in the main text. A2 does not impose the no-tilt benchmark; whether the bid's response also tilts on soil within a location is an empirical question, and a within-location good-land tilt would arise if development historically lands on what was good farmland. The direct test conditions on proximity and asks whether non-farm development responds more strongly to IHS access for proximate (fringe) parcels than for distant ones, and whether any residual soil tilt survives once continuous exposure is held fixed. [Appendix B.6](#) reports this test.

What A1 does and does not restrict. $\phi_{CL} = (c_C - c_L)/(b_C - b_L)$ may still shift, because A1 leaves the price legs $\partial b_C/\partial \tau$ and $\partial b_L/\partial \tau$ unordered and a difference between them moves the denominator; whether the crop band extends or retreats with access is therefore left to the data. Any such shift, however, loads on good land, not marginal: the parcel density at the cutoff is $g(\phi_{CL})$, where g is the induced density of quality $\phi = A_c \omega$ (the scale A_c acting on the Fréchet draw ω), and $g(\phi_{CL})$ carries a factor A_c^θ and vanishes as the county quality scale $A_c \rightarrow 0$, so a ϕ_{CL} shift moves negligible crop acreage in marginal (low- A_c) counties and bites only where mass sits near the cutoff, on good land. The model therefore predicts no marginal-concentrated commodity-crop expansion – commodity cropland on the marginal class is flat to good-land-concentrated, not crops climbing down onto poor soil. The farm-versus-exit

A1 refuses to order the cross-activity response rather than claiming the input bundles are identical, and the refusal is not vacuous: it forbids marginal-concentrated commodity-crop expansion – the A_c^θ density argument confines any ϕ_{CL} shift to good land – and imposes within-activity uniformity, so a b_k shift moves the exit threshold rather than reshuffling which parcels farm. The model's composition content runs through the extensive margin, which land is retained or shed, with livestock's breakeven responding to access through $\partial b_L/\partial \tau$, the source of the retention result ([Appendix B.2](#)).

Appendix B.3 Productivity, the Fréchet Specification, and Agricultural Returns

Productivity decomposition. Each parcel's agricultural productivity is

$$\phi_{jc} = A_c \cdot \omega_{jc}, \quad (18)$$

as in equation (1) of the main text. Following Sotelo (2020), the parcel-level idiosyncratic draw ω_{jc} is distributed independently across parcels with the Fréchet CDF:

$$F(\omega) = \exp(-\omega^{-\theta}), \quad \omega > 0, \quad (19)$$

with density $f(\omega) = \theta \omega^{-(\theta+1)} \exp(-\omega^{-\theta})$ and shape parameter $\theta > 1$. This is equation (2) of the main text. The shape parameter θ governs within-county productivity dispersion: lower θ implies a heavier right tail, more parcels near the selection margin, and a larger extensive-margin response to changes in returns.

Under this specification, $\phi_{jc} \sim \text{Fréchet}(A_c^\theta, \theta)$: county suitability A_c enters as the scale parameter of the composite productivity distribution, so that every quantile of the within-county productivity distribution is proportionally higher in high-suitability counties. Two moments of the Fréchet distribution that appear in the aggregation below:

$$\mathbb{E}[\omega^k] = \Gamma\left(1 - \frac{k}{\theta}\right), \quad k < \theta, \quad (20)$$

$$\int_x^\infty \omega f(\omega) d\omega = \gamma\left(1 - \frac{1}{\theta}, x^{-\theta}\right), \quad (21)$$

where $\Gamma(\cdot)$ is the gamma function and $\gamma(a, x) = \int_0^x t^{a-1} e^{-t} dt$ is the lower incomplete gamma function.¹⁴

Production technology. A farm on parcel j running activity $k \in \{C, L\}$ produces:

$$Q_{jct} = s_{jc} \cdot \phi_{jc} \cdot f_k(k_{jct}, \ell_{jct}, m_{jct}), \quad (22)$$

where s_{jc} is parcel acreage, $f_k(\cdot)$ is an activity-specific constant-returns-to-scale production function in per-acre capital k , labor ℓ , and materials m , and ϕ_{jc} is Hicks-neutral productivity. One quality index ϕ_{jc} serves either activity; crops and livestock differ in how steeply they reward it (Assumption A3).

Per-acre agricultural returns. A farm running activity k chooses per-acre inputs to maximize profit per acre:

$$R^k(\phi_{jc}, \tau_{ct}) = \max_{\kappa, \ell, m \geq 0} \left\{ b_k(\tau_{ct}) \phi_{jc} f_k(\kappa, \ell, m) - r_K(\tau_{ct}) \kappa - w(\tau_{ct}) \ell - p^M(\tau_{ct}) m \right\}. \quad (23)$$

Under constant returns to scale, optimal input ratios are independent of the scale of ϕ_{jc} . Substituting yields the two-activity returns of the main text (equation (3)):

$$R^C(\phi_{jc}, \tau_{ct}) = b_C(\tau_{ct}) \phi_{jc} - c_C(\tau_{ct}), \quad R^L(\phi_{jc}, \tau_{ct}) = b_L(\tau_{ct}) \phi_{jc} - c_L(\tau_{ct}), \quad (24)$$

¹⁴Derivation of the tail (first incomplete) moment: with $f(\omega) = \theta \omega^{-(\theta+1)} e^{-\omega^{-\theta}}$, substitute $u = \omega^{-\theta}$ so $\omega = u^{-1/\theta}$ and $d\omega = -(1/\theta) u^{-1/\theta-1} du$. As ω runs from x to ∞ , u runs from $x^{-\theta}$ to 0. Then $\int_x^\infty \omega f(\omega) d\omega = \int_0^{x^{-\theta}} u^{(1-1/\theta)-1} e^{-u} du = \gamma(1 - 1/\theta, x^{-\theta})$. Note this is the *first incomplete moment* of the Fréchet distribution, not the conditional (truncated) mean; the conditional mean is $\gamma(1 - 1/\theta, (x)^{-\theta})/[1 - F(x)]$, i.e., divided by the survival probability.

where $b_k(\tau)$ absorbs the technology constant and encompasses both the output-price effects of equation (15) and the optimal input-mix adjustment as τ changes. The key property is linearity in ϕ_{jc} : per-acre returns are the product of an access-dependent effective price $b_k(\tau)$ and parcel quality, less an activity-specific cost $c_k(\tau)$. Crops are the intensive use ($b_C > b_L$, $c_C > c_L$) and livestock the extensive one (Assumption A3).

Reach-versus-competition incidence (Assumption A1). A highway pulls each farm-gate return $b_k(\tau_{ct})$ two ways at once, the reach-versus-competition force of [Donaldson and Hornbeck \(2016\)](#): it lets local producers reach distant buyers, which raises the return, and it lets rival sellers into the markets the county serves, which lowers it. The net sign of $\partial b_k / \partial \tau$ is left to the data, separately for each activity, with no ordering imposed between crops and livestock. Costs respond too, $\partial c_k / \partial \tau$ ambiguous in sign, with the response common across activities so the cost gap $c_C - c_L$ is access-invariant. And because each R^k is linear in ϕ_{jc} , whichever way a return moves it moves the same way on every parcel in a county: a highway does not shift the price more on a county's better land than on its worse land, so the sign is an empirical question, not a model primitive.

Crop and grazing input bundles do differ – cropping is chemical- and fuel-intensive, extensive grazing land- and labor-intensive – so a shipped-in-input cost cut plausibly falls harder on crops. Imposing a common cost response declines a price-driven composition prediction the data do not support rather than asserting the bundles are identical: A1 makes no claim about whether access nudges a county's crop–livestock mix, leaving that to the data, and the model's composition content instead runs through the extensive margin.

Appendix B.4 The Participation Condition and Threshold

The pivot and the two-regime cutoff ordering. Two development-pressure regimes follow, divided by a pivot $\bar{\rho}$ in the outside option. Let $\bar{\rho} \equiv b_L \phi_{CL} - c_L$ be the level of the outside option at which the farm–exit cutoff ϕ^* reaches the crop–livestock cutoff ϕ_{CL} ; by the definition of ϕ_{CL} this is also $b_C \phi_{CL} - c_C$, so at $\rho = \bar{\rho}$ both farm returns equal the outside option at ϕ_{CL} at once and the two exit cutoffs meet ϕ_{CL} at a single point. Below the pivot ($\rho < \bar{\rho}$, low-to-moderate pressure) the sort is the familiar one of equation (6): a distinct crop-exit cutoff $\phi_C^* = (\rho + c_C)/b_C$ lies strictly between the livestock-exit cutoff $\phi^* = (\rho + c_L)/b_L$ and ϕ_{CL} – the ordering $\phi^* < \phi_C^* < \phi_{CL}$ holds because the crop-exit cutoff falls more slowly in ρ than the livestock-exit cutoff ($\partial \phi_C^* / \partial \rho = 1/b_C < 1/b_L = \partial \phi^* / \partial \rho$, since $b_C > b_L$) – so a crop parcel ($\phi \geq \phi_{CL}$) clears both exit cutoffs, crops never bind their own exit margin, and the livestock band is the only shed farm margin. This is Regime I, and it nests the single-margin sort exactly, with ϕ_C^* slack inside the retained farm range. Above the pivot ($\rho \geq \bar{\rho}$, high pressure) the livestock-exit cutoff has risen past ϕ_{CL} , the pure grazing band is gone, and the binding farm–exit cutoff on the crop range becomes $\phi_C^* = (R^O - \kappa + c_C)/b_C \geq \phi_{CL}$, which climbs through the crop band from below and develops the lowest-value crop parcels near ϕ_{CL} . This is Regime II – the co-equal crop-exit margin (ϕ_C^* here is the $\phi_C^\dagger$ of equation (77)) – and displacement is no longer a grazing-only phenomenon. Which counties reach Regime II is set by where the developer bar is

highest, hence by where M_c is highest; the maintained correlation $\text{Corr}(A_c, M_c) > 0$ (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.16](#) carries the Fréchet shares for both regimes and the pivot.

Land allocation. Parcel j in county c takes its best use at time t ,

$$V_{jc} = \max \left\{ R^C, R^L, (R_{ct}^O - \kappa_{ct}) \mathbf{1}\{\text{fringe}\}, R_{ct}^N - \kappa_N, 0 \right\}, \quad (25)$$

where $\kappa_{ct} \geq 0$ is the conversion wedge (zoning, transaction, and physical costs of building) and the 0 is the amenity/holding reservation, all returns measured net of it. This is equation (4) of the main text. The conversion wedge creates a band of inaction: land remains in agriculture even when an outside option marginally exceeds current returns, capturing inertia from zoning, transaction costs, and the irreversibility of development.

Two cutoffs. The two farm returns rise with quality, crops faster, while the non-farm options are flat in ω_{jc} , so a parcel's use is pinned against two cutoffs. Crops beat livestock when $(b_C - b_L)\phi \geq c_C - c_L$, i.e. above the crop–livestock cutoff

$$\phi_{CL} = \frac{c_C - c_L}{b_C - b_L}, \quad (26)$$

a comparison of the two farm technologies alone, free of soil and of the outside option. Livestock – the marginal farm use – beats the best non-farm option ρ_{ct} ($R_{ct}^O - \kappa_{ct}$ in the fringe, $R_{ct}^N - \kappa_N$ beyond it) when $b_L\phi - c_L \geq \rho_{ct}$, i.e. above the farm–exit cutoff

$$\omega_{ct}^* = \frac{\rho_{ct} + c_L(\tau_{ct})}{b_L(\tau_{ct}) A_c}, \quad \phi_{ct}^* \equiv A_c \omega_{ct}^* = \frac{\rho_{ct} + c_L}{b_L}. \quad (27)$$

This is equation (6) of the main text. When $\phi_{ct}^* < \phi_{CL}$ (coexistence) the 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. Because crops are the steeper technology, $\phi \geq \phi_{CL} > \phi^*$ implies crops also beat the outside option, so no separate crop–exit margin binds in the interior (coexistence) regime; the squeeze-out boundary, where intense development pressure lifts ϕ^* past ϕ_{CL} and develops the marginal crop parcels, is in [Appendix B.16](#).

The farm–exit cutoff ω_{ct}^* governs the agricultural land share and carries the comparative statics below. Several properties are immediate:

- $\partial\omega^*/\partial A_c < 0$: higher county suitability lowers the threshold; more parcels remain in agriculture.
- $\partial\omega^*/\partial\rho_{ct} > 0$: a higher best non-farm option (a builder's bid in the fringe, the reversion value beyond it) raises the threshold; more parcels exit.

- $\partial\omega^*/\partial b_L < 0$: a higher livestock farm-gate return lowers the threshold; more parcels stay.
- $\partial\omega^*/\partial c_L > 0$: higher livestock cost raises the threshold; more parcels exit.
- $\partial\omega^*/\partial\kappa < 0$: a higher conversion wedge lowers the effective outside option; more parcels stay.

Comparative statics with respect to τ . Fix a development-pressure stratum $\sigma \in \{\text{none, press}\}$ with best non-farm option ρ^σ ($\rho^{\text{press}} = R_{ct}^O - \kappa_{ct}$ in the fringe, $\rho^{\text{none}} = R_{ct}^N - \kappa_N$ beyond it).¹⁵ The farm-exit cutoff is $\omega_\sigma^* = (\rho^\sigma + c_L)/(b_L A_c)$. Differentiating with respect to τ_{ct} :

$$\frac{\partial\omega_{\sigma,ct}^*}{\partial\tau_{ct}} = \frac{1}{A_c} \cdot \frac{\partial}{\partial\tau} \left[\frac{\rho^\sigma + c_L(\tau)}{b_L(\tau)} \right]. \quad (28)$$

Applying the quotient rule to the bracket, with $N \equiv \rho^\sigma + c_L(\tau)$ and $D \equiv b_L(\tau)$ (the $1/A_c$ factor is already outside):

$$\frac{\partial}{\partial\tau} \left(\frac{N}{D} \right) = \frac{N'D - ND'}{D^2} = \frac{(\partial\rho^\sigma/\partial\tau + \partial c_L/\partial\tau) b_L - (b_L \phi_\sigma^*) \partial b_L/\partial\tau}{b_L^2},$$

where the ND' term uses $N = b_L(\tau) \phi_\sigma^*$ (since $N/D = \phi_\sigma^*$ by equation (27)). Dividing through by b_L and applying the leading $1/A_c$:

$$\frac{\partial\omega_{\sigma,ct}^*}{\partial\tau_{ct}} = \frac{1}{b_L(\tau) A_c} B^\sigma, \quad B^\sigma = \frac{\partial\rho^\sigma}{\partial\tau} + \frac{\partial c_L}{\partial\tau} - \phi_\sigma^* \frac{\partial b_L}{\partial\tau}. \quad (29)$$

This is equation (10) of the main text, with $\partial\rho^\sigma/\partial\tau$ the response of the stratum's outside option – the development bid $\partial R_{ct}^{O,\text{press}}/\partial\tau$ in the fringe, the reversion response $\partial R^N/\partial\tau$ beyond it. The first two terms are the numerator effect: better access (τ falls) raises the outside option and moves costs ambiguously, both shifting the exit threshold ϕ_σ^* . The third is the denominator effect: under A1 the sign of $\partial b_L/\partial\tau$ is left to the data, so which way the livestock farm-gate pulls ω^* is itself an empirical question.

A highway lowers τ ($\Delta\tau < 0$), so the threshold's actual change $\Delta\omega^*$ runs *opposite* to $\partial\omega^*/\partial\tau$. On the value/capitalization margin the sign of the threshold response is set by development pressure, through the fringe development option of A2, not by soil: on the pressured subgroup, the development-bid term $\partial R_{ct}^{O,\text{press}}/\partial\tau$ is large and tends to dominate, while on the unpressured subgroup the outside option is the smaller reversion R^N and the net farm-gate slope governs. Suitability A_c enters this margin only as amplification:

1. The $1/A_c$ scaling in (28): for any given change in the bracket, the resulting movement in ω^* is larger on marginal (low- A_c) land, making marginal land more sensitive in

¹⁵The superscript σ on ρ , c , R^O , B , ω^* , and the primitives $(\pi_R^\sigma, \pi_R^\sigma)$ throughout this appendix indexes the development-pressure stratum, and is distinct from the CES substitution elasticity σ of Assumption A3. The two are different objects; the stratum superscript is a label, never an exponent, and the welfare component $W_{ct}^O \propto \rho_{ct}^\sigma$ of equation (62) is linear in ρ_{ct}^σ , not raised to a power.

magnitude.

2. The pressure stratum: a parcel in the conversion fringe faces the development bid $\partial R^{O,\text{press}}/\partial\tau$, larger than the reversion response beyond the fringe, soil-independently.

Together, these imply that on the value/capitalization margin the threshold response is governed by pressure and amplified on marginal land, ordering the value–revenue wedge of Proposition 1 on M_c . The survival margin runs on a different axis: the marginal farm use is uniformly livestock, but comparative advantage (Appendix B.16) makes low-suitability counties livestock-dominated, so the cross-suitability survival gradient is soil-ordered through composition and present at zero pressure (Lemma 1), not a projection of development pressure.

Why $1/A_c$ scaling alone cannot reverse the sign. The comparative static (29) has terms that both depend on A_c . To see that A_c affects only the *magnitude*, not the *sign*, of $\partial\omega^*/\partial\tau$, substitute the closed form for ω_σ^* from equation (27) into the denominator term:

$$\frac{\omega_\sigma^*}{b_L(\tau)} \cdot \frac{\partial b_L}{\partial\tau} = \frac{\rho^\sigma + c_L(\tau_{ct})}{b_L(\tau) A_c} \cdot \frac{\partial b_L}{\partial\tau}. \quad (30)$$

This term scales as $1/A_c$, identical to the scaling of the numerator term, so factoring $1/(b_L A_c)$ recovers (29):

$$\frac{\partial\omega_{\sigma,ct}^*}{\partial\tau} = \frac{1}{b_L(\tau) A_c} \left[\frac{\partial\rho^\sigma}{\partial\tau} + \frac{\partial c_L}{\partial\tau} - \frac{(\rho^\sigma + c_L) \partial b_L / \partial\tau}{b_L(\tau)} \right] = \frac{B^\sigma}{b_L(\tau) A_c}. \quad (31)$$

The bracket B^σ contains no A_c : the *sign* of $\partial\omega^*/\partial\tau$ is determined entirely by whether $|\partial\rho^\sigma/\partial\tau + \partial c_L/\partial\tau|$ exceeds $(\rho^\sigma + c_L) |\partial b_L/\partial\tau|/b_L(\tau)$ or falls short. The $1/A_c$ prefactor scales the magnitude but cannot change the sign. B^σ is common across suitability within a pressure stratum only when the channel responses entering it are themselves common across suitability – the within-stratum homogeneity condition (Maintained Condition 2 below). On the value/capitalization margin that condition holds, because the loading is the soil-independent outside option $R^{O,\sigma}$ of A2 (the development bid in the fringe, the reversion R^N beyond): B^σ is then common across suitability within a stratum and differs across strata, the pressured subgroup carrying the larger development response $|\partial R^{O,\text{press}}/\partial\tau|$. The survival heterogeneity runs instead through composition: the marginal use is uniformly livestock, and low-suitability counties are livestock-dominated, so the cross-suitability survival gradient is comparative advantage, present at $M_c = 0$ and not a projection of development pressure.

Proof of Proposition 1. Maintained Conditions. The proof invokes two conditions in addition to Assumptions A1 and A2:

1. *Interior solution (MC1):* $\omega_c^* \in (0, \infty)$, $f(\omega_c^*) > 0$, $b_L > 0$, $A_c > 0$, and the stratum outside option $\rho^\sigma > 0$ for parcels at the threshold.

2. *Within-stratum homogeneity of the channel responses (MC2)*: within a fixed development-pressure stratum σ , the channel responses $\partial R^{O,\sigma}/\partial\tau$, $\partial c_L/\partial\tau$, and the livestock farm-gate slope $\partial b_L/\partial\tau$, together with the cost baseline c_L and the conversion wedge κ , are common across suitability classes. This is consistent with the model specification in (27), which writes $c_L(\tau)$ and $b_L(\tau)$ without a suitability subscript; empirically, the structural GMM in Section 7.3 stratifies the primitives by development-pressure status rather than by suitability class, so common-within-stratum is the operative restriction. The pressure stratification of Section 7.3 (the strata {None, Low, High}, the median split of the positive development-pressure exposure M_c) operationalizes this condition, with suitability entering only as a first-stage trend control; holding the channel responses common across suitability within each development-pressure stratum is what the restriction requires.

MC2 holds on the value/capitalization margin. On the value margin the wedge loads on the outside option $R^{O,\sigma}$ – the development bid in the fringe, the reversion R^N beyond it – which is soil-independent under A2, so each term of the bracket is common across suitability within a stratum. The survival margin separates on a different axis: the marginal farm use is uniformly livestock, so the survival response carries no binding-activity switch, but comparative advantage makes low-suitability counties livestock-dominated (Appendix B.16), so the cross-suitability survival gradient is soil-ordered through composition. The proof below establishes the capitalization result under MC2 and records the composition-driven survival ordering.

The cross-county ordering of the *level* of the outside option – that the county-average R^O is higher on good land – is not maintained. The parcel-level outside option is soil-independent (A2); the county-average ordering is an endogenous consequence of development-pressure composition: good-land counties carry a larger pressured share through $\text{cor}(M_c, A_c) > 0$, so their area-weighted R^O is higher. No level ordering is imposed.

Proof strategy. Show the bracket governing the cutoff response is free of A_c on the value margin under MC2, identify the wedge as the capitalization loading on $R^{O,\sigma}$, and decompose the county-level wedge into a within-stratum scaling part and a between-stratum part carried by M_c .

Step 1: Bracket independence (value margin). Evaluated within pressure stratum σ , equation (31) gives

$$\frac{\partial \omega_{ct}^{*,\sigma}}{\partial \tau} = \frac{1}{b_L(\tau) A_c} B^\sigma, \quad B^\sigma \equiv \frac{\partial R_{ct}^{O,\sigma}}{\partial \tau} + \frac{\partial c_L}{\partial \tau} - \phi_\sigma^* \frac{\partial b_L}{\partial \tau}. \quad (32)$$

By MC2, which holds on the value margin, each term of B^σ is common across suitability within the stratum: $\partial R^{O,\sigma}/\partial\tau$ and $R^{O,\sigma}$ are soil-independent (A2), $\partial b_L/\partial\tau$ is common across parcels (A1), and c_L , κ , $\partial c_L/\partial\tau$ are common within the stratum (the threshold quality ϕ_σ^* is itself soil-free, the same across the suitability classes of a stratum). Hence $\partial B^\sigma/\partial A_c = 0$: suitability enters only through the $1/(b_L A_c)$ prefactor.

Step 2: The wedge isolates capitalization. Farmland value per acre records the best use, the maximum of agricultural and outside-use value, while agricultural revenue per acre records current gross output, the conditional-mean yield of surviving crop and livestock

parcels (Appendix B.8). The outside option $R_{ct}^{O,\sigma}$ enters value through both composition – which parcels survive, governed by $\omega^{*,\sigma}$ and hence by B^σ – and a direct capitalization term, whereas it enters revenue through composition alone. The wedge $\Delta \ln(\text{value/acre}) - \Delta \ln(\text{revenue/acre})$ therefore isolates the capitalization of the outside option, and its magnitude scales with $\partial R_{ct}^{O,\sigma} / \partial \tau$, the loading that A2 makes pressure-driven and soil-independent. The wedge is largest where the surviving net agricultural margin is thinnest relative to gross output, which the elevated $R_{ct}^{O,\text{press}}$ produces by raising the participation threshold $\omega^{*,\text{press}}$ above $\omega^{*,\text{none}}$.

Step 3: County wedge decomposition and ordering. Writing the county wedge as the area-weighted mixture $W_{ct} = (1 - M_c) W_{ct}^{\text{none}} + M_c W_{ct}^{\text{press}}$ with $W_{ct}^{\text{press}} > W_{ct}^{\text{none}}$ (Step 2), the suitability gradient of the wedge is

$$\frac{dW_{ct}}{dA_c} = \underbrace{(1 - M_c) \frac{\partial W_{ct}^{\text{none}}}{\partial A_c} + M_c \frac{\partial W_{ct}^{\text{press}}}{\partial A_c}}_{\text{(I) revenue-base amplification } (bA_c)} + \underbrace{(W_{ct}^{\text{press}} - W_{ct}^{\text{none}}) \frac{dM_c}{dA_c}}_{\text{(II) composition via } \text{cor}(M_c, A_c)}. \quad (33)$$

By Step 1, term (I) carries no suitability response in B^σ ; it is the secondary amplification of the wedge through the revenue base $b(\tau)A_c$ that scales the denominator. Term (II) is the channel by which “more pressure on better land” orders the wedge, governed by $dM_c/dA_c \propto \text{cor}(M_c, A_c)$. Conditioning on M_c holds the pressure stratum fixed and sweeps out term (II), leaving the secondary revenue-base amplification (I). The wedge is therefore ordered by development pressure, with suitability a secondary amplifier.

Step 4: The survival margin runs through composition, not a switch. The marginal farm use is uniformly livestock: it carries the lowest exit threshold, so the farm–exit cutoff is ω_σ^* in every county, with no binding-activity switch across soil. Survival therefore does not break MC2 through a suitability-dependent bracket; it separates from capitalization through a different outside option and a different force. The survival cutoff $\omega_{\text{none}}^* = (R_{ct}^N - \kappa_N + c_L)/(b_L A_c)$ runs against the *reversion* value R^N (the counterfactual to grazing beyond the fringe is conservation or idle reversion, not development), present at $M_c = 0$ and independent of pressure. The cross-suitability gradient enters through *composition*: the crop share of farmland $G(\phi_{CL})/G(\phi^*)$ is monotone increasing in A_c (Appendix B.16), so low-suitability counties are livestock-dominated and the marginal land access retains is livestock land. Soil orders *which counties* retain at the margin – the marginal class, livestock-dominated – while pressure orders *the exit destination*. The two margins are thus driven by distinct outside options, R^N (survival, soil-ordered through composition) and R^O (capitalization, pressure-ordered), so conditioning on M_c leaves the survival gradient intact.

The value/capitalization margin is unaffected. The wedge isolates the capitalization of the development option $R_{ct}^{O,\sigma}$, whose response $\partial R_{ct}^{O,\sigma} / \partial \tau$ is common across activities – the outside option is the same regardless of which activity a parcel would run – and soil-independent (A2). The activity-specific b_k and c_k enter value and revenue symmetrically through the surviving composition and net out of the wedge, leaving the common pressure-driven $R_{ct}^{O,\sigma}$ loading. MC2 holds on the wedge margin, the wedge orders on M_c , and Steps 1–3 stand.

Conclusion. On the value/capitalization margin MC2 holds, the bracket B^σ is free of A_c ,

and the value–revenue wedge is ordered by development pressure with suitability a secondary amplifier through the revenue base; conditioning on M_c attenuates the wedge toward that secondary scaling. The survival gradient runs instead through the soil-ordering of activity composition against the reversion counterfactual R^N (Step 4): it is present at $M_c = 0$ and conditioning on M_c does not remove it. $\square$

Proof of Lemma 1. The marginal farm use is livestock, the lowest-threshold use, so which land stays in farming is the livestock breakeven. Beyond the fringe this is the farm-versus-reversion margin $\omega_{\text{none},ct}^* = (R_{ct}^N - \kappa_N + c_L)/(b_L A_c)$, present at $M_c = 0$ and independent of development pressure. By Assumption A3 the crop share of farmland $G(\phi_{CL})/G(\phi^*)$ is monotone increasing in A_c (Appendix B.16), so low-suitability counties are livestock-dominated and high-suitability counties crop-dominated; retention of marginal land is therefore retention of livestock. The survival response inherits the livestock farm-gate slope: where reach raises b_L ($\partial b_L/\partial \tau < 0$), ω_{none}^* falls and marginal land that would revert stays in grazing (provided the farm-gate response dominates any access-induced movement in R^N). The cross-suitability gradient is then governed by the soil-ordering of activity composition rather than by development pressure, so it appears at $M_c = 0$ and conditioning on M_c – which sweeps out only the compositional pressure channel – does not collapse it. The reversion value R^N is level-higher in marginal counties (conservation targets poor land), which raises the baseline exit there; but the difference-in-differences identifies the response $\partial \omega^*/\partial \tau$, not the baseline level, and the higher baseline exit is why marginal counties carry the most land at risk to retain. $\square$

Appendix B.5 Falsification Tests

The empirical content of Proposition 1 and Lemma 1 is carried by four falsification tests, N1, N3, N4, and the survival test below; the acreage/extensive-margin falsification (N2) — that farmland acreage contracts under pressure and the displaced margin is extensive grazing rather than confinement — is the discriminating-test content of Appendix B.9, and is not restated here. Each test follows from the proposition or the lemma under its maintained conditions; none introduces a new assumption.

Corollary 1 (N1: capitalization mediation by development pressure – decisive) *Conditional on the development-pressure stratum M_c , the value–revenue capitalization wedge orders on pressure: it is larger on the pressured subgroup, and the partial association between suitability and the wedge attenuates to the secondary revenue-base amplification once M_c is held fixed. A wedge that does not order on M_c , or a partial suitability association that does not attenuate under conditioning on M_c , falsifies the capitalization-mediation claim.*

Proof of Corollary 1. By Steps 1–2 of the proof of Proposition 1, the wedge loads on the outside option $R^{O,\sigma}$ through capitalization, with $\partial B^\sigma/\partial A_c = 0$ on the value margin under MC2, so the only suitability dependence is the revenue-base amplification through term (I) of (33). By Step 3, the between-class ordering of the wedge operates through term (II), the compositional channel governed by $dM_c/dA_c \propto \text{cor}(M_c, A_c)$, and conditioning on M_c holds the stratum fixed and sweeps out term (II), leaving the secondary

revenue-base amplification. A wedge unordered by M_c , or a partial suitability association that fails to attenuate under conditioning, is inconsistent with the capitalization loading on the soil-independent $R^{O,\sigma}$. $\square$

Corollary 2 (N3: common-signed no-pressure value baseline) *Within the no-pressure stratum, the access effect on farmland value per acre is common in sign across the four suitability classes. The corresponding no-pressure baseline on the agricultural land share need not be common-signed: by Lemma 1 the survival baseline is soil-ordered, and a class that switches sign – the marginal class retained while a suitable class contracts – is the comparative-advantage signature, not a failure of the value baseline.*

Proof of Corollary 2. On the no-pressure stratum the value-margin bracket B^{none} is suitability-free (Step 1 of the proof of Proposition 1), so $\text{sign}(\partial\omega^{*,\text{none}}/\partial\tau) = \text{sign}(B^{\text{none}})$ for every class on the value margin, and the value response inherits this common sign with A_c entering only the magnitude through $1/A_c$. The survival baseline runs through composition (Step 4): low-suitability counties are livestock-dominated and retain the marginal land, so the no-pressure survival gradient is soil-ordered rather than common-signed, the comparative-advantage signature. $\square$

Corollary 3 (N4: the value–revenue wedge is pressure-first) *The value–revenue wedge*

$$\Delta \ln(\text{value/acre}) - \Delta \ln(\text{revenue/acre})$$

is primarily increasing in development pressure M_c , with suitability a secondary amplifier through the revenue base.

Proof of Corollary 3. By Appendix B.8 and Step 2 of the proof of Proposition 1, the value moment loads on the outside option R^O through both composition and direct capitalization, while the revenue moment loads on R^O through composition alone; the wedge therefore isolates capitalization. Capitalization is largest where the surviving net agricultural margin is thinnest relative to gross output, which the elevated outside option of A2 produces under development pressure. Since R^O varies with the parcel’s exposure to the fringe rather than with A_c (A2), the wedge orders on M_c ; suitability enters only through the revenue base $b_L(\tau)A_c$ that scales the denominator, a secondary amplifier. $\square$

Corollary 4 (Survival: soil-ordered, not pressure-mediated) *The cross-suitability gradient in the agricultural land share (survival) is soil-ordered: the marginal class is retained while the suitable classes turn on the net farm-gate incidence. The gradient is present in the no-pressure stratum ($M_c = 0$) and is not collapsed by conditioning on M_c ; it is the comparative-advantage margin, distinct from and not mediated by the pressure-driven capitalization wedge.*

Proof of Corollary 4. By Lemma 1 and Step 4 of the proof of Proposition 1, the marginal farm use is uniformly livestock, and under Assumption A3 the crop share of farmland $G(\phi_{CL})/G(\phi^*)$ is monotone increasing in A_c : low-suitability counties are livestock-dominated,

so the marginal land access retains is livestock land. The survival gradient is then ordered by the soil-ordering of activity composition against the reversion counterfactual R^N rather than by the development pressure M_c , so it appears at $M_c = 0$ and conditioning on M_c (which sweeps out only the compositional pressure channel) leaves it intact. The survival gradient is therefore comparative-advantage driven. $\square$

Agricultural land share under Fréchet. The fraction of parcels remaining in agriculture is:

$$S_{ct}^A = 1 - F(\omega_{ct}^*) = 1 - \exp(-(\omega_{ct}^*)^{-\theta}), \quad (34)$$

as in equation (7) of the main text. The response to access:

$$\frac{\partial S_{ct}^A}{\partial \tau} = -f(\omega_{ct}^*) \cdot \frac{\partial \omega_{ct}^*}{\partial \tau} = -\theta (\omega_{ct}^*)^{-(\theta+1)} \exp(-(\omega_{ct}^*)^{-\theta}) \cdot \frac{\partial \omega_{ct}^*}{\partial \tau}. \quad (35)$$

On the unpressured land the bracket B^{none} is set mainly by the livestock farm-gate slope $\partial b_L / \partial \tau$: when reach wins, $\partial S^A / \partial \tau < 0$ and better access adds farmland, holding marginal land in grazing against the reversion counterfactual R^N . On the pressured land the fringe development response of A2 drives B^{press} the other way, so $\partial S^A / \partial \tau > 0$ and better access shrinks the farm base as the development option takes over. The dividing line is pressure, not soil; soil only scales the size of the move through $1/A_c$.

The four predictions, as originally stated in prose.

1. **Marginal land is retained in farming, against a reversion counterfactual.** On marginal land, where crops cannot pay, livestock holds the land in farming; 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, present where there is no development pressure ($M_c = 0$), and not washed out by conditioning on M_c , because the counterfactual to retained marginal land is reversion, not development. *Falsified if* marginal retention vanishes at $M_c = 0$ or collapses under M_c -conditioning.
2. **Where development pressure binds, farmland acreage falls and revenue per acre falls as the land-rent bar prices out extensive grazing, even as the surviving mix shifts toward input-intensive crops; value per acre is flat.** A highway at the fringe raises the development outside option R^O — a rising *land-rent bar* every farm use must clear to stay in agriculture. The bar enters the farm-exit cutoff ϕ^* , not the technology-only ϕ_{CL} , so it climbs through the grazing band from below and prices out extensive grazing first: acreage contracts and livestock revenue per acre collapses, while surviving crop parcels substitute non-land inputs for the now-dearer land, so input use per acre rises with M_c . Value per acre is not signed by the model — it is the sum of the revenue decline of Prediction 2 and the capitalization wedge of Prediction 4, so its empirical flatness is a joint over-identifying restriction. The mechanism 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 $M_c = 0$). *Falsified if* acreage does not fall under pressure, if revenue

per acre *rises* with pressure, if input use rises under access alone rather than scaling with pressure, or if the displaced animal margin is confinement rather than extensive grazing. [Appendix B.8](#) and [Appendix B.9](#) derive the result, 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 reads Prediction 2 one layer finer.

3. **Adjacent counties show attenuated, same-direction effects, by channel.** Adjacent counties receive the access improvement attenuated, $\lambda \Delta\tau$ with $\lambda < 1$, on the transport channel, so revenue and input margins, which load on b_k , propagate to neighbors. The development-capitalization channel is local: the builder's bid R^O is a direct-county phenomenon, so it does not spill. The model predicts channel-specific pass-through — a substantial λ on the revenue effect and a much smaller one on the value effect — not a single common ratio.
4. **Farmland value and agricultural revenue diverge where the outside option capitalizes (the capitalization signature).** Farmland value per acre records the parcel's best use, including the development option; agricultural revenue per acre records only current output. The development option touches revenue only through which parcels survive, never directly. 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 the outside option is the lower reversion value $R^N \ll R^O$, with soil only a secondary amplifier (Proposition 1, falsification test N4). Prediction 4 is distinct from Prediction 2, not a relabeling: they differ in comparison object — acreage and revenue per acre versus the value–revenue gap — and bind in the same place. The crop mix-shift enters value and revenue in the same direction and so does not open the wedge; the development option capitalizes into value alone, and so does. On a no-pressure marginal county neither pressure statement operates: Prediction 1 holds while the Prediction 4 wedge is near zero and the Prediction 2 contraction is absent. Under pressure the signed prediction is the capitalization wedge (equation (55)), not a compositional value gain.

Appendix B.6 Direct Evidence on Assumption A2

Where the displacement lands. The A2 correlation $\text{Corr}(A_c, M_c) > 0$ is meaningful but not primitive. The displacement *mechanism* — the exit cutoff truncating an occupied band under a high developer bar — is native to the two-regime structure. *Where* it lands — the shed grazing band across good and middle soil, and the Regime II crop-exit band on middle soil where the best soil still capitalizes — 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.

Assumption A2 maintains that the outside-use return responds more strongly to falling transport costs for parcels in the conversion fringe of growing cities than for distant parcels, and that the developer's bid at the parcel is soil-blind: within a fixed exposure to that fringe, the bid does not tilt toward better cropland. The soil-blindness clause is the

empirically contestable one, and we test it directly on observed nonfarm development, since the outside option is an unobservable in the model. The natural worry is that development is observed disproportionately on good land, which would appear to contradict a soil-blind bid; the test isolates whether that raw pattern survives holding true fringe exposure fixed, or is instead the mechanical consequence of cities sitting on good land. We measure development from the HISDAC-US built-up record (Leyk and Uhl, 2018), which reconstructs the cumulative footprint of built-up structures on a 250-meter raster grid at 5-year intervals. Because the raw structure count mixes non-farm building with agricultural structures the record cannot label, and because logging a count with heterogeneous baselines manufactures spurious across-cell contrasts, we use a base-robust developed-share outcome – the fraction of county pixels built-up – estimated with the canonical DID2s specification of the main analysis (all controls year-interacted and partitioned in both stages, break-ground treatment, Conley 100 km standard errors). Development pressure enters as the continuous moderator M_c rather than a coarse metro-proximity split, and the suitability partition is the same $N/S3/S2S1$ classes used throughout.

The discriminating test holds continuous M_c fixed and asks whether, within a level of fringe exposure, the developed-share response tilts toward better cropland; it does not. The good-land tilt visible under a coarse binary pressure split ($b_{\text{press}} \approx +0.014$ on the higher-suitability classes) collapses once the continuous dose absorbs true exposure: the within-exposure suitability slope is $b_{\text{trend}} = +0.002$ ($p = 0.57$) and the omnibus test of equality across suitability classes cannot reject ($F_{\text{equal}}, p = 0.17$), while the pressure dose itself carries the development response ($+0.054$, significant). The raw good-land pattern is thus cities-on-good-land exposure, not a soil-graded bid: once true exposure is held fixed, the developer’s bid is soil-blind, consistent with A2. This test speaks to the extent margin the developed-share outcome records; because HISDAC carries no building-use labels and cannot net out agricultural structures, it is informative for A2’s soil-blindness clause without settling every channel of the outside option.

Appendix B.7 County-Level Aggregation

Cross-sectional analog: size–productivity by activity. The livestock implication has a directly observable cross-sectional analog: cropping operations show positive size–productivity correlation and livestock operations the opposite, larger ranches on poorer rangeland reflecting carrying-capacity scaling, a pattern consistent with federal land-grant institutions that granted larger parcels for grazing (Appendix B.7).

We aggregate from individual parcels to county-level outcomes observed in the Agricultural Census. All integrals are taken with respect to the Fréchet density

$$f(\omega) = \theta \omega^{-(\theta+1)} \exp(-\omega^{-\theta}).$$

Let $h_c(\omega, s)$ denote the joint density of the idiosyncratic productivity draw ω_{jc} and parcel size s_{jc} in county c .

Activity-conditional dependence between size and productivity. Parcel size and productivity are not assumed independent. Cross-sectionally, $\text{cov}(s, \omega)$ exhibits opposite

signs across activities. Within cropping, economies of scale and fixed managerial overhead favor larger farms on better cropland, so $\text{cov}(s, \omega \mid \text{cropping}) > 0$. Within livestock and extensive grazing, carrying capacity scales inversely with land quality and operations require larger acreage on poorer rangeland to reach financial viability, so $\text{cov}(s, \omega \mid \text{livestock}) < 0$. This pattern is consistent with the federal land-grant institutions of the expansion era, which explicitly compensated for low ω with larger acreage to deliver financially viable units: the Homestead Act of 1862 allocated 160 acres on relatively good farmland; the Kinkaid Act of 1904 allocated 640 acres in the Nebraska Sand Hills; the Stock-Raising Homestead Act of 1916 allocated 640 acres for grazing on rangeland; and the Desert Land Act of 1877 allowed up to 640 acres in arid western states. These institutions left an empirical artifact in mid-twentieth-century farm sizes that is consistent with negative cross-sectional $\text{cov}(s, \omega)$ where livestock dominates and positive $\text{cov}(s, \omega)$ where cropping dominates.

Aggregation under activity-conditional dependence. The closed-form welfare expressions in [Appendix B.10](#) use $\bar{s}_c$ as a county-level aggregation weight, replacing the joint density $h_c(\omega, s)$ with the marginal $f(\omega)$ multiplied by mean parcel acreage $\bar{s}_c$. This is exact under activity-composition stability within the county and serves as a benchmark approximation otherwise. Ratios of welfare components are robust to this approximation: in W_{ct}^O/W_{ct}^A and W_{ct}^A/W_{ct} , the $\bar{s}_c$ factor outside the integrals cancels between numerator and denominator. The lower-bound application in [Section 7.3](#) uses such ratios; the bound argument is therefore well-grounded under activity-conditional dependence even when the absolute levels of W_{ct}^A and W_{ct}^O depend on the size-productivity mapping. Activity-specific modeling of $g_c(s \mid \omega, \text{activity})$ is a natural extension reserved for future work.

Number of active farms.

$$S_{ct} = \iint_{\omega \geq \omega_{ct}^*} h_c(\omega, s) d\omega ds = 1 - F(\omega_{ct}^*) = 1 - \exp(-(\omega_{ct}^*)^{-\theta}). \quad (36)$$

Under Fréchet, this is the closed-form expression from [\(34\)](#).

Land in farms. Total acreage in agricultural use:

$$N_{ct}^A = \iint s_{jc} \cdot \mathbf{1}\{\omega_{jc} \geq \omega_{ct}^*\} h_c(\omega, s) d\omega ds. \quad (37)$$

Average productivity among survivors. The conditional mean productivity among surviving parcels uses the first incomplete moment in [\(21\)](#):

$$\int_{\omega^*}^{\infty} \omega f(\omega) d\omega = \gamma \left(1 - \frac{1}{\theta}, (\omega^*)^{-\theta}\right), \quad (38)$$

so that the average productivity among operating parcels is:

$$\bar{\phi}_{ct}^S = A_c \cdot \frac{\gamma \left(1 - \frac{1}{\theta}, (\omega_{ct}^*)^{-\theta}\right)}{S_{ct}^A}. \quad (39)$$

Equation (39) defines $\bar{\phi}_{ct}^S$ as the conditional mean of farm productivity over the surviving base; it is the survivor-composition input the welfare and θ -identification expressions below consume (Appendix B.10 and the θ -GMM system), where $m_{ct} = \bar{\phi}_{ct}^S / A_c$ enters the moment functions. The conditional-mean construction follows Sotelo (2020). It is a book-keeping object, not a channel: the land-value results of Appendix B.8 do not run through a single-index claim that tighter participation raises survivor quality and thereby lifts value or revenue; under development pressure the per-acre revenue movement is the displacement of the extensive-grazing band and the crop-mix tilt (eq. 52), not a selection gain in $\bar{\phi}_{ct}^S$.

Net revenue per acre (model). Net revenue per acre of farmland is the gross output of the average surviving parcel, averaged over the crop and livestock bands, less per-acre cost. Writing Φ_k for the parcels running activity k and N_{ct}^A for total farmed acreage,

$$\frac{\text{Revenue}_{ct}}{N_{ct}^A} = \frac{1}{N_{ct}^A} \left[\sum_{j \in \Phi_C} b_C(\tau_{ct}) A_c \omega_{jc} + \sum_{j \in \Phi_L} b_L(\tau_{ct}) A_c \omega_{jc} \right] - \bar{c}_{ct} = \bar{b}_{ct} A_c m_{ct} - \bar{c}_{ct}, \quad (40)$$

where $m_{ct} = \mathbb{E}[\omega \mid \omega \geq \omega_{ct}^*]$ is the conditional-mean quality of surviving parcels (from equation (39), $\bar{\phi}_{ct}^S = A_c m_{ct}$) and $\bar{b}_{ct}$, $\bar{c}_{ct}$ are the activity-share-weighted farm-gate return and cost. The second equality normalizes parcel acreage to unity; with heterogeneous parcels the mean-acreage weight $\bar{s}_c$ enters numerator and denominator identically and cancels, as in the welfare ratios of Appendix B.10. Equation (40) is the per-acre revenue identity; the conditional mean m_{ct} is well defined whenever $\theta > 1$ (a finite first moment). Under development pressure this identity does not move through a selection gain in m_{ct} : as derived in Appendix B.8, the operative revenue statement resolves equation (40) into its crop and extensive-grazing activities (eq. 52), and per-acre revenue falls because the land-rent bar prices off the marginal extensive-grazing band, not because participation raises survivor quality. The same $\bar{b}_{ct} A_c m_{ct}$ gross-revenue term appears in the value–revenue wedge below. Revenue per acre in the Agricultural Census corresponds to gross receipts per acre, which uses the received price rather than the net return $\bar{b}$ (which subtracts delivery costs), so reported receipts, profits, and land values need not move together.

Appendix B.8 Land Values

Why the bar sheds extensive grazing first. Comparative advantage fixes not just which use occupies each parcel but its *position* in the quality distribution, and that placement is all the rising land-rent bar of A2 needs: extensive grazing earns the lowest farm rent and sits on the marginal band $\phi^* \leq \phi < \phi_{CL}$, so a rising development bid, which lifts

the farm–exit cutoff ϕ^* but not the technology-only ϕ_{CL} , climbs *through the grazing band from below* and prices grazing out first – because comparative advantage placed it lowest, not because it is the most land-intensive use. Acreage falls, extensive revenue falls, and the surviving mix tilts toward input-intensive cropping, whose per-acre input use rises through a mix shift and a within-parcel CES land–input substitution the model folds together (both driven by the same bar, so not separately identified). 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 (feedlots, hogs), whose parcel rent is near-insensitive to the land price; A3 carves it out of the flat-return grazing technology as the benchmark use the bar leaves untouched, which is what makes the displaced margin a test – the mechanism sheds extensive grazing, not confinement ([Appendix B.8](#)).

Reading the capitalization Proposition. Proposition 1 ties the capitalization gap to pressure, not to soil. Because a builder bids only in the fringe, the wedge is present where pressure is – B^{press} carries the development bid’s response $\partial R^{O,\text{press}}/\partial\tau < 0$, the bid rising as the wedge falls – and small beyond it, 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. Holding the pressure stratum fixed removes the between-stratum channel and leaves only the soil scaling, so the wedge attenuates once pressure is controlled. The pressure ordering and the attenuation under M_c -conditioning are the falsifiable content (falsification test N1, [Appendix B.5](#)).

Reading the survival Lemma. Lemma 1 keeps the survival margin apart from the capitalization margin. What survives is tied to soil – crops on good land, livestock on poor – so the cross-suitability survival pattern is a comparative-advantage result that holds wherever development pressure happens to sit. 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 revert stays in grazing, provided the farm-gate response dominates any access-induced movement in R^N . The reversion value R^N is level-higher in marginal counties, since conservation targets poor land, which raises the baseline exit there; but the difference-in-differences identifies the treatment *response* $\partial\omega^*/\partial\tau$, not the baseline level, and the higher baseline exit is precisely why marginal counties have the most land at risk to retain.

Parcel-level values. The value of a parcel under each use, capitalized at the common rate $r > 0$:

$$V_{ct}^A(\phi_{jc}) = \frac{\max\{R^C(\phi_{jc}, \tau_{ct}), R^L(\phi_{jc}, \tau_{ct})\}}{r}, \quad (41)$$

$$V_{ct}^{\text{nf}} = \frac{\max\{(R_{ct}^O - \kappa_{ct})\mathbf{1}\{\text{fringe}\}, R_{ct}^N - \kappa_N\}}{r}, \quad (42)$$

the best farm value V^A (crop or livestock) and the best non-farm value V^{nf} – the development option $(R^O - \kappa)/r$ in the fringe, the reversion $(R^N - \kappa_N)/r$ beyond it. All options share the perpetuity rate $1/r$, so the value max ranks the same as the flow allocation of the main text. The parcel owner takes its best use:

$$V_{ct}(\phi_{jc}) = \max\{V_{ct}^A(\phi_{jc}), V_{ct}^{\text{nf}}\}. \quad (43)$$

County-level farmland value per acre. What the Agricultural Census measures as farmland value per acre is the acreage-weighted average over parcels classified as farmland:

$$\bar{V}_{ct} = \frac{\iint V_{ct}(\phi_{jc}) \cdot s_{jc} \cdot \mathbf{1}\{\omega_{jc} \geq \omega_{ct}^*\} h_c(\omega, s) d\omega ds}{N_{ct}^A}. \quad (44)$$

When access improves, $\bar{V}_{ct}$ changes through two channels: (1) the value per parcel rises, because the best farm return $\max\{R^C, R^L\}$ rises for high-productivity parcels (under A1) and R^O rises for all parcels; and (2) the composition of farmland changes, because the set of parcels in the denominator N_{ct}^A shifts as ω_{ct}^* moves.

The development bid is a proximity-driven option; its soil-dependence is measured, not assumed. Assumption A2 fixes which regime a parcel faces by proximity: a fringe parcel that leaves farming faces the capitalized development option $R_{ct}^O - \kappa_{ct}$, raised by a highway, $\partial R^O / \partial \tau > 0$, and beyond the fringe no builder bids and land reverts to $R_{ct}^N - \kappa_N \geq 0$. A2 makes no claim about whether the bid depends on the parcel's farm productivity ω_{jc} : the derivation below does not lean on any sign of $\text{Cov}(R_{jc}^O, \omega_{jc})$, and whether a developer's willingness to pay tilts toward good farmland within a location is settled empirically – holding continuous exposure M_c fixed leaves no detectable monotone good-land tilt in the value-only leg ([Appendix B.6](#)). A higher development option does one thing with many consequences: it raises the price of land, the *land-rent bar* a farm use must clear to stay in agriculture, and $V^A(\omega)$ of (41) is the parcel's farm capitalization against that bar. The single driver of everything that follows is this rising bar; the activities sort against it by their value per acre.

Comparative advantage orders the activities against the land-rent bar (A3). Comparative advantage (A3) places each activity at a *position* in the quality distribution, and the von Thünen logic sorts those positions against the land-rent bar by value per acre ([von Thünen, 1826](#)). Extensive grazing – range beef – is the marginal use: land *is* the input, with no margin on which to raise output per acre, so its land rent is the extensive return

$$r^L(\omega, \tau) = b_L(\tau) A_c \omega - c_L, \quad (45)$$

the lowest farm rent on any given parcel (A3, $b_L < b_C$). Cropping earns a higher land rent per acre, $r^C(\omega, \tau) = b_C(\tau) A_c \omega - c_C$ with $b_C > b_L$; achieving that higher value per acre takes more chemicals and fuel per acre, the input-heaviness A3 already records. Land-

light confinement – dairy belts, feedlots, hog operations – earns its return on housed animals and feed, using land only as a footprint, so its land rent is near-insensitive to the bar; A3 carves it out of the flat-return grazing technology, and it enters here only as the activity the land-rent bar does *not* touch. Order the three by value per acre against the bar: extensive grazing (marginal, no value-per-acre margin) $\prec$ crops (input-intensive, high value per acre) $\prec$ confinement (land-light, off the land-rent margin). A rising land-rent bar sorts activities up this order.

The land-rent bar drives per-acre intensification through a CES land-input choice. Per-acre input use rises through two channels: the surviving mix tilts toward the input-intensive crop technology as grazing is priced out, and surviving crop parcels substitute non-land inputs for land as land grows dear. To carry the second channel the crop parcel chooses non-land inputs against a *variable* land quantity. A crop operator on land of quality ϕ facing land price $R = R^O - \kappa$ and non-land input price w_τ produces with a constant-returns CES aggregate of land services ℓ (acres of quality ϕ) and a non-land input bundle x ,

$$y(\ell, x; \phi) = \phi b_C(\tau) \left[\alpha \ell^{\frac{\sigma-1}{\sigma}} + (1-\alpha) x^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 0, \alpha \in (0, 1), \quad (46)$$

and maximizes $p y - R \ell - w_\tau x$ over ℓ and x . Land is a chosen quantity, so the land price enters the input choice through the two interior conditions $p \partial y / \partial \ell = R$ and $p \partial y / \partial x = w_\tau$. The cost-minimizing input-land ratio depends only on relative prices,

$$\frac{x^*}{\ell^*} = \left(\frac{1-\alpha}{\alpha} \right)^\sigma \left(\frac{R}{w_\tau} \right)^\sigma \implies \frac{\partial \ln(x^*/\ell^*)}{\partial \ln R} = \sigma > 0, \quad (47)$$

so per-acre input intensity – input per cultivated acre, the object the fertilizer-per-harvested-acre outcome measures – rises in the land price with elasticity σ , unconditionally. This is the substitution law: as the development bid raises R , the parcel farms fewer acres harder. Because $R = R^O - \kappa$ in the fringe, the per-acre response *scales with the bar* – with pressure M_c – and is flat-to-negative where no development bid operates, matching Prediction 2. The canonical estimates bear this out, with fertilizer per harvested acre rising +0.72 per unit of M_c while its pure-access level is negative (−0.28), so access without pressure does not intensify – a level the dedicated paragraph below (eq. 54) signs as the dilution footprint of retained grazing on the farmland-acre measure. The Leontief limit $\sigma = 0$ switches the substitution channel off, $\partial(x^*/\ell^*)/\partial \ln R = 0$, and the per-acre intensity reverts to the fixed activity constant of the share-mix identity (eq. 48) below; the extension is the single feature $\sigma > 0$, and its null is the prior model.¹⁶

A shift in the operator pool is likewise ruled out as a rival to within-parcel intensification: farm count is stable-to-up and average farm size stable-to-down in the fringe crop classes under pressure, and the off-farm-work operator share is flat under pressure across

¹⁶The mechanism reads on the per-harvested-acre ratio x^*/ℓ^* together with survivor composition, not on literal within-parcel idling; the chosen land footprint of (46) is the formal device that generates the substitution law. The harvested share of cropland is flat under pressure, so per-harvested and per-cultivated intensity coincide in the data and the idling margin carries no empirical weight.

the post-treatment years the census records it, so the per-acre input rise is not an artifact of fewer, larger farms or of a tilt toward part-time operators farming intensively on the side.

Lemma: a rising land-rent bar prices out the marginal (extensive-grazing) band first. Pressure raises the bar, $\partial(R_{ct}^O - \kappa_{ct})/\partial\tau > 0$. A parcel is developed when the bar exceeds its best farm rent. Because extensive grazing earns the *lowest* farm rent (it is the marginal use, A3) and has *no* value-per-acre margin to lift that rent – r^L in (45) is fixed-intensity, land already its only input – the parcels whose best farm rent first falls below the rising bar are the grazing parcels. The bar enters the *exit* cutoff $\phi_{\text{press},ct}^*$ of equation (8), which rises with $R_{ct}^O - \kappa_{ct}$ and climbs *through the grazing band from below*, shrinking the grazing band on the ϕ^* side and converting range land to development. The crop–grazing composition cutoff $\phi_{CL} = (c_C - c_L)/(b_C - b_L)$ of equation (5) is a comparison of the two farm technologies alone and is invariant to R^O ; the bar does not move it. The contraction therefore runs entirely through the rising exit cutoff ϕ^* , not through ϕ_{CL} : the marginal grazing band is shed first (cattle inventory per acre -0.40 , livestock revenue per pasture acre -0.65 per unit of M_c) while land-light confinement is untouched (hogs -0.14 , not significant, consistent with the A3 carve-out). These are the same marginal/extensive parcels as the livestock corollary of Prediction 1: retained in grazing where there is no pressure, priced out by development where there is.

Lemma: the surviving activity mix shifts toward input-intensive crops, raising input use per acre. The same rising exit cutoff ϕ^* that sheds grazing tilts the surviving farm composition toward the input-intensive use that clears the bar. Let s_t^L be the extensive-grazing share of surviving fringe farmland and $s_t^C = 1 - s_t^L$ the crop share. As ϕ^* rises through the grazing band, s_t^L falls and s_t^C rises; within the crop band, the von Thünen logic selects the higher-value-per-acre, more input-intensive cropping that can pay the higher rent. Average input use *per acre* is the share-weighted intensity

$$\bar{x}_t = s_t^C \bar{x}_t^C + s_t^L \bar{x}_t^L, \quad \bar{x}_t^C \gg \bar{x}_t^L, \quad (48)$$

with $\bar{x}_t^C$ the (high) crop intensity and $\bar{x}_t^L \approx 0$ the (negligible) grazing intensity. As s_t^C rises the mix tilts toward the input-heavy activity, so the share-weighted *per-acre* intensity $\bar{x}_t$ rises – fertilizer per farmland acre (+0.64) and crop value per cropland acre (+0.17 to +0.22 per unit of M_c).

Total input use rises under a magnitude condition. Equation (48) is a per-acre share-weighted average; it does not by itself sign total input use, since total use is per-acre intensity times surviving acreage and acreage is falling. Total input use is

$$X_t = \bar{x}_t \cdot N_t^A, \quad \frac{d \ln X_t}{dt} = \underbrace{\frac{d \ln \bar{x}_t}{dt}}_{> 0 \text{ (mix-shift to crops)}} + \underbrace{\frac{d \ln N_t^A}{dt}}_{< 0 \text{ (acreage contracts)}}, \quad (49)$$

a sum of opposite signs. Total input use rises only under the *maintained magnitude condition* – stated explicitly, as for the revenue and wedge legs – that the crop intensive- margin and activity expansion dominate the acreage contraction, $d \ln \bar{x}_t / dt > -d \ln N_t^A / dt$. The canonical estimates place the data in this regime: fertilizer use per farmland acre rises with pressure (+0.64 per unit of M_c), so the intensification is genuine input-land substitution at the activity-mix level and not an artifact of dividing a fixed input bill by a shrinking harvested-acre denominator. The totals result is therefore an empirical magnitude statement, not a corollary of the per-acre identity.

Acreage falls under pressure (robust, signed). Let f denote the within-county Fréchet density of ω and define the per- ω retention probability

$$\pi_t(\omega) = \Pr \left[\max_k r^k(\omega, \tau_t) \geq R_{ct}^O - \kappa_{ct} \mid \omega \right], \quad (50)$$

the chance the parcel's best farm rent clears the (pressure-shifted) land-rent bar. By A3, $\max_k r^k$ is increasing in ω , so π_t is increasing in ω ; a higher bar lowers π_t pointwise. The fringe farm-acreage share is the retention-weighted mass $S_t^A = \int_0^\infty \pi_t(\omega) f(\omega) d\omega$. Because the bar shift is a first-order stochastic worsening of the farm-versus-develop comparison at every ω ,

$$\frac{\partial S_t^A}{\partial t} = \int_0^\infty \frac{\partial \pi_t(\omega)}{\partial t} f(\omega) d\omega \leq 0, \quad (51)$$

with strict inequality whenever the bar shift has positive mass on the survivor margin. The interchange of differentiation and integration is justified by dominated convergence: the difference quotients of π_t are bounded in absolute value by an integrable envelope, since $\sup_\omega |\partial_t \pi_t(\omega)| \leq g_0 < \infty$ for the bounded reverse-hazard family governing the exit comparison, and $g_0 f(\omega)$ is integrable. Farm acreage falls under pressure *regardless* of any soil-bid dependence: the bar shifts up at every ω , so every parcel is weakly more likely to be priced out, and the marginal extensive-grazing band is shed first. This is the signed acreage half of Prediction 2 (−0.24 in the canonical full-sample estimates), the leg the model signs without any further condition.

Revenue per acre falls because the priced-out extensive use dominates the affected land. Total agricultural revenue per farmland acre is the acreage-weighted average of the crop and grazing revenues over surviving parcels. With per-acre revenues $\bar{R}_t^C$ (crop) and $\bar{R}_t^L$ (extensive grazing), the revenue identity (40), resolved into the two activities and net of per-acre cost $\bar{c}_{ct}$, is

$$\frac{\text{Revenue}_{ct}}{N_{ct}^A} = s_t^C \underbrace{b_C(\tau) A_c m_t^C}_{\bar{R}_t^C} + s_t^L \underbrace{b_L(\tau) A_c m_t^L}_{\bar{R}_t^L} - \bar{c}_{ct}, \quad (52)$$

with m_t^k the conditional-mean quality of activity- k survivors. Equation (52) is an accounting identity for the two revenue legs, and the reduced form measures the legs $s_t^k \bar{R}_t^k$ directly; it does not identify the slopes b_C and b_L separately. The two conditional means m_t^C and m_t^L are taken over different, non-overlapping quality bands – crops on the high

tail, extensive grazing on the middle band – so recovering $b_k = \bar{R}_t^k / (A_c m_t^k)$ from observed revenue divides each activity's revenue by its own survivor mean and compares returns evaluated at different quality points, which does not sign $b_C - b_L$. The A3 ordering is identified by the crop-share allocation gradient of [Appendix B.16](#) and by the revenue co-movement below, not by these level ratios; the activity-specific level split is a selection-deflated robustness object, not a slope estimate. Pressure moves the two legs in *opposite* directions through the single sorting force. The grazing leg collapses: the land-rent bar prices off the marginal grazing parcels, and per surviving pasture acre the extensive operation winds down as conversion approaches, so $s_t^L \bar{R}_t^L$ falls sharply (livestock revenue per pasture acre -0.65). The crop leg rises as the mix tilts toward higher-value-per-acre cropping (crop revenue per cropland acre $+0.17$ to $+0.22$). Because extensive grazing is the marginal use occupying the bulk of the *affected* fringe acreage, the negative grazing leg dominates the acreage-weighted average:

$$\frac{d}{dt} \frac{\text{Revenue}_{ct}}{N_{ct}^A} = \underbrace{\frac{d}{dt} (s_t^L \bar{R}_t^L)}_{< 0 \text{ (extensive priced out, large)}} + \underbrace{\frac{d}{dt} (s_t^C \bar{R}_t^C)}_{> 0 \text{ (crop mix-shift, small)}} - \frac{d\bar{c}_{ct}}{dt} < 0 \quad (53)$$

whenever the grazing contraction exceeds the crop gain on the affected land – the maintained magnitude regime the data select, in which the -0.65 grazing leg outweighs the $+0.17$ crop leg net of any cost movement. We retain the cost term $\bar{c}_{ct}$: it is common to the level and nets out of the proportional statement but is not dropped from the identity. This is the revenue half of Prediction 2 (total agricultural revenue per acre -0.18 in the canonical full-sample estimates), resting on the land-rent displacement of extensive grazing rather than on cropland-quality selection.

The mechanism is fringe-localized. The land-rent bar binds only where a builder is bidding, so the entire pattern is a *fringe* statement: in the rural core, where $R^O \ll R^N$ and the bar is the reversion value rather than a developer's option, the bar does not rise with pressure and the displacement does not operate. The rural-core pattern in which crops contract is a separate access and relative-price phenomenon outside this fringe-localized mechanism, leaving the net core response near zero; any nonzero core component is a distinct channel and not the land-rent sorting derived here.

The data fix the case, and the alternatives cannot jointly match it. Neither the displacement nor the mix-shift is assumed for convenience; they are the data-selected branch of a single sorting mechanism, and the falsifiable content is the *joint* co-movement and its scaling with pressure. A pure composition story – development skimming the agriculturally worst land – would push revenue per acre *up* and predict no rise in input use; the data show revenue down with total fertilizer up. An access-driven intensification story would predict input use rising at $M_c = 0$ (under pure access); the data show the pure-access level *negative* (-0.28), with the rise appearing only as pressure M_c rises ($+0.72$ slope), so the driver is the land-rent bar, not access. A confinement story would shed the land-light animal margin; the data leave it untouched to expanding (hogs -0.14 , not

significant; the cattle-on-feed leg rises under pressure) and shed extensive grazing instead (range cattle per acre -0.40). Only the unified land-rent sorting here – the bar prices out the marginal, lowest-value-per-acre use first and tilts the surviving mix toward input-intensive crops – reproduces the observed pattern, scaling with M_c : acreage down, extensive-grazing revenue per acre down, crop input use up in totals, value per acre flat, and the value–revenue wedge up. The model does not “predict” the revenue decline in isolation; it offers a signed comparative static in $\partial R^O/\partial\tau$ – one driver – and the data discriminate this branch from the alternatives by its sign pattern and its scaling with pressure.

The pure-access level on intensity is negative as a feature of Prediction 1. The identifying content of the intensification result is the *positive pressure slope* on input use ($+0.72$ per unit of M_c on fertilizer per harvested acre), which is robust across the totals (eq. 49) and across every per-acre denominator. The role of the *negative* pure-access level is narrower: it rules out an access-driven reading of the intensification, since a market-access story would put the intensity rise at $M_c = 0$. On the *farmland-acre* measure that negative level is itself a feature of Prediction 1. Where access improves without development pressure, ω_{ct}^* falls and marginal land is *retained* in extensive grazing rather than reverting (Prediction 1, the livestock corollary). The retained acres are low-fertilizer grazing acres added to the county *farmland* base, so fertilizer per farmland acre is diluted *downward* through the county-wide share-weighted average $\bar{x}_{ct}$ of (48):

$$\left. \frac{\partial}{\partial\tau} \right|_{M_c=0} \ln \bar{x}_{ct} = \underbrace{\left. \frac{\partial s_t^L}{\partial\tau} \right|_{M_c=0}}_{>0 \text{ (grazing retained, P1)}} \cdot \underbrace{(\bar{x}_t^L - \bar{x}_t^C)}_{<0 \text{ (grazing is input-light)}} \leq 0, \quad (54)$$

where $\partial s_t^L/\partial\tau|_{M_c=0} > 0$ because the marginal band retained under Prediction 1 is extensive grazing by A3, so the marginal entrants under access are grazing acres and the grazing share of the farmland base rises. The dilution operates mechanically only on the farmland-acre basis, since the retained grazing land enters that denominator; it is not a statement about the harvested- or cropland-acre measures, where retained grazing does not enter the denominator and the model does not separately sign the pure-access level (any level there may reflect a crop-mix or input-price effect the model leaves unsigned). On the farmland-acre measure the same retention that signs Prediction 1 thus signs the pure-access intensity level, and the intensification appears only when pressure raises the bar and tilts the surviving mix toward crops. The level and slope on fertilizer per farmland acre are two readings of one unified story: retention dilutes intensity under access, the bar concentrates it under pressure.

Value per acre: a lost extensive use against a capitalized option, hence unsigned. Per-acre value $\bar{V}_{ct}$ of (44) moves under pressure through two channels of opposite sign. The first is a composition channel running through the *loss of the low-value extensive use*: when grazing parcels valued near the bottom of the farm distribution are converted, the surviving farmland base loses its cheapest acres, which by itself pulls average value in an

ambiguous direction depending on what replaces them in the denominator. The second is direct capitalization: near-margin survivors *in the fringe* are valued at essentially $(R_{ct}^O - \kappa_{ct})/r$, so a higher development option capitalizes *up* into the value of surviving parcels (Prediction 4). The crop mix-shift does *not* sign the value level either: tilting toward higher-value-per-acre cropping raises both current crop output and, through the higher land rent it generates, the capitalized farm value V^A , so it enters value and revenue in the same direction and does not by itself open or close the value–revenue wedge. The level change in log value,

$$\frac{d}{dt} \ln \bar{V}_{ct} = \underbrace{(\text{extensive-use composition})}_{\text{ambiguous sign}} + \underbrace{\frac{d}{dt} \ln \Theta_{ct}}_{> 0 \text{ (capitalization)}}, \quad (55)$$

is the offset of an ambiguous composition term and a positive capitalization term and is *not* signed by the model. Its empirical flatness (≈ 0) is not a convenient offset but an *over-identifying check*: total agricultural revenue per acre (-0.18 , the grazing-minus-crop net) and the wedge ($+0.19$, capitalization) are two independently estimated signed legs on the same full sample, and a value-per-acre level near zero is their arithmetic implication, $\Delta \ln \bar{V} = \Delta \ln(\text{revenue}) + \Delta(\text{wedge}) \approx -0.18 + 0.19 \approx 0$ – matching the directly estimated farm value per acre ($+0.032$, not significant, i.e. flat), a consistency restriction the data pass, not a parameter we tune. This over-identifying check is the meaningful value-flatness statement, and it holds however hard the surviving crop parcels are farmed: the CES substitution of (47) raises inframarginal crop rents, but because the two legs are measured separately, whatever those rents contribute already sits inside the estimated composition leg, and their near-cancellation with the capitalization leg is what the data verify.

Why the two legs offset: fringe land value tracks the development option, not agronomy. The check states *that* the legs offset; the development bid explains *why*. Near cities land is priced by the development option, not by agricultural productivity: a crop parcel’s market value is the larger of its agricultural residual and the builder’s bid, $V(\phi) = \max\{V_{\text{res}}^A(\phi), (R^O - \kappa)/r\}$. For the parcels the builder out-bids, the land sells at the bid, so however intensively they are farmed the rent that capitalizes is set by R^O , not by farm product – intensification underneath the ceiling is value-neutral. This is why the capitalization leg is a *bid* effect that offsets the composition/revenue leg rather than compounding it: both are the same rising R^O read on two margins, and (55) is the statement that they cancel. The bid’s reach over the suitability distribution is the same fact seen a second way. The agricultural residual rises only shallowly in land quality on the surviving fringe crop band – an order of magnitude of land quality buys well under an order of magnitude of agricultural value per acre – so the flat development bid dominates the agricultural residual across the observed quality range: the crossing at which farming out-values the bid lies beyond observed land quality, and the escaping high- ϕ tail whose rents *do* capitalize carries essentially zero acreage weight within support. Its weight is a bounded, calibrated quantity in the corroborated-bound style: on the Fréchet survival G of (78) the escaping acreage share of the fringe crop band is $(\phi_C^*/\phi^\dagger)^\theta$, with ϕ_C^* the co-equal crop-exit cutoff (eq. 77) and $\phi^\dagger$ the quality at which the agricultural residual overtakes the bid, falling be-

low a fifth of the band only once the bid reaches the upper decile at the estimated $\theta \approx 2$ – and the measured schedule places the crossing above that reach. Whatever residual survives is one contributor to the unsigned composition leg the check already measures, so the escaping tail is accounted for by the data, not assumed away. This is the [Capozza and Helsley \(1989, 1990\)](#) development-option capitalization the paper invokes in the co-equal crop-exit margin (eq. 77) and in Prediction 4, supplying the economic content of the offset without adding a claim the consistency check does not already carry.

Lemma: the value–revenue wedge is signed positive (capitalization). The capitalization loading is the survivor-weighted ratio of capitalized parcel value to its current-output base,

$$\Theta_{ct} \equiv \mathbb{E}_t \left[\frac{V_{ct}(\phi_{jc})}{A_c \omega_{jc}} \right] \quad (\text{ratio of capitalized parcel value to its current-output base}). \quad (56)$$

The current-output base $A_c \omega$ is common to value and revenue and differences out of the wedge; only after it nets out does $d \ln \Theta_{ct} / dt$ isolate the capitalization of the development option, net of any movement in the farm-gate return:

$$\frac{d}{dt} \left[\ln \bar{V}_{ct} - \ln (\text{Revenue}_{ct} / N_{ct}^A) \right] = \underbrace{\frac{d}{dt} \ln \Theta_{ct}}_{> 0 \text{ (capitalization)}} - \frac{d}{dt} \ln \bar{b}_{ct} > 0. \quad (57)$$

The development option capitalizes into value but not into current agricultural output: a parcel that stays in farming sells the same harvest whether or not a builder covets it, so the option touches revenue only through which parcels survive, never directly. The wedge is signed positive under the *maintained magnitude condition* – stated explicitly, not parenthetically – that the capitalization loading dominates any movement in the farm-gate return induced by the activity mix-shift:

$$\frac{d}{dt} \ln \Theta_{ct} > \frac{d}{dt} \ln \bar{b}_{ct}. \quad (58)$$

The *same* treatment of the farm-gate term – it moves with the activity mix but is dominated by capitalization – governs both the revenue identity (53) and the wedge (57); we do not hold it fixed in one and let it move in the other. This is the object of Prediction 4 and Proposition 1: signed positive, pressure-first, with suitability a secondary amplifier through the revenue base. The wedge corroborates that a development option exists and is capitalized; it is not an input to the welfare-omission magnitude, which rides the direct π_R leg. Reading the wedge into the omission size would double-count the capitalization the value moment $\ln V^A$ already carries.

The wedge nets out farm fundamentals only up to the capitalization multiplier Θ_{ct} of (56). Two features of Θ_{ct} keep it from being a clean structural zero. First, it is soil-correlated through survival: the surviving farmland base sheds its cheapest acres as the bar climbs, so the composition term in (55) carries an *ambiguous* sign rather than a null. Second, Θ_{ct} is the ratio of capitalized parcel value to its current-output base, so it houses a frontage- and

amenity-value component – the value a developer prices into a parcel that never enters current farm revenue – which the current-output normalization does not remove. Holding continuous development pressure M_c fixed conditions out the exposure confound and leaves no *detectable monotone* good-land tilt in this value-only component, but it does not provably drive it to zero: it does not clear the non-monotone significant $S3$ hump, nor the GAEZ terrain and climate (slope, frontage-amenity) channel, and a structurally small but nonzero value-only tilt cannot be excluded. The defensible claim is the weaker one – conditional on continuous exposure, no monotone good-land tilt in the value-only leg is detectable – so the wedge’s pressure ordering is a *tested implication*, not an assumed consequence. It is never asserted to be a structural zero.

On the *unpressured subgroup* under net reach ($\partial b_L / \partial \tau < 0$), ω_{ct}^* falls when access improves: previously marginal parcels (with ω just below the old threshold) enter agriculture as livestock. These new entrants have per-acre value approximately equal to the reversion value $V_{ct}^{nf} = (R_{ct}^N - \kappa_N) / r$, since beyond the fringe they are near the farm-versus-reversion margin, and this is below the average value of the existing agricultural base. Adding these below-average parcels to the denominator of (44) exerts downward pressure on $\bar{V}_{ct}$, partially or fully offsetting the direct price effect on incumbents. The retention itself (ω_{ct}^* falls beyond the fringe, so marginal land is held in grazing rather than reverting) is Prediction 1, driven by the livestock farm-gate response $\partial b_L / \partial \tau$ and unchanged from the main-text derivation.

What Prediction 2 reduces to, and where the value movement comes from. With a rising land-rent bar that prices out the marginal extensive use first and an activity mix that tilts toward input-intensive crops, Prediction 2 reduces on the pressured subgroup to two signed statements – *acreage down* (eq. 51) and *revenue per acre down* through the extensive displacement net of the crop mix-shift (eq. 53) – and one explicitly unsigned one, the per-acre value level (eq. 55), whose empirical flatness is the over-identifying residual of the two signed legs. The value–revenue wedge of Proposition 1 isolates capitalization (eq. 57). $\bar{V}_{ct}$ and the aggregate W_{ct} of Appendix B.10 coincide only when the composition of farmland is fixed; under pressure S_t^A contracts and the extensive band is displaced, so they diverge, and pressure-stratified treatment effects on farmland value per acre identify the capitalization wedge rather than a composition gain.

Appendix B.9 Distinguishing the Land-Rent Mechanism from Its Alternatives

The two poles of the animal margin, and the prime end. no builder bidding, displaced into development where there is. The confinement carve-out is directly separated in the data, not merely assumed: the land-rent bar sheds extensive grazing (cattle and calves per acre -0.40 per unit of M_c , livestock revenue per pasture acre -0.65) while land-light confinement is untouched to expanding, hog inventory flat on the marginal dose (-0.14 , not significant) and rising where pressure binds hardest as the feedlot leg of cattle rises where range cattle collapse (Appendix B.9). The two poles of the animal margin move in opposite directions on the same land – extensive grazing displaced, confinement held to

expanding – so the A3 carve-out is shown rather than posited. At the prime end the two forces run together: good land carries the largest fringe share, so the conversion pull is strongest there, but holding continuous exposure M_c fixed collapses the within-location good-land tilt – good land converts more through greater exposure, with no monotone soil tilt in the development bid detectable once exposure is held fixed (Appendix B.6), the exposure phenomenon consistent with the two-regime option and subject only to the residual value-only premium in the capitalization multiplier (Appendix B.8).

The land-rent sorting mechanism is one driver, and it is identified by three signatures that the leading alternatives cannot jointly produce. Each was pre-registered and confirmed in the canonical estimates (pooled treatment $\times M_c$, full sample, with the level reading the pure-access effect at $M_c = 0$ and the slope reading the per-unit pressure effect).

Signature 1: intensification scales with pressure, not access. If the rising input use were an access response – crops farmed harder because shipping is cheaper – it would appear under access alone, at $M_c = 0$; it does not. Fertilizer per harvested acre has a *negative* pure-access level (-0.28) and a strongly *positive* pressure slope ($+0.72$): input use rises only as the land-rent bar rises with M_c . The identifying content is the positive slope; the negative level rules out the access reading, and on the farmland-acre measure it is the dilution footprint of the Prediction 1 retention of low-fertilizer grazing (eq. 54). The driver is the land price, not market access, which is the von Thünen logic the mechanism formalizes and the discriminating fact against an access-intensification reading.

Signature 2: the input rise is genuine, not a denominator artifact. A rise in input use *per harvested acre* could in principle reflect a shrinking denominator rather than more inputs. It does not: the rise holds against the full-farmland denominator, fertilizer per farmland acre rising with pressure ($+0.64$ slope, the magnitude condition of eq. 49 satisfied in the data), and crop value per cropland acre rises with it ($+0.22$). The activity mix-shift toward input-intensive cropping raises input use in absolute terms, the genuine land–input substitution the sorting mechanism implies.

Signature 3: the displaced margin is extensive grazing, not confinement. The mechanism prices out the marginal, lowest-value-per-acre activity first. The displaced margin is extensive grazing – cattle and calves per acre fall sharply with pressure (-0.40 slope), and livestock revenue per pasture acre falls (-0.65) – while land-light confinement is untouched to expanding: hog and pig inventory shows no pressure response on the continuous dose (-0.14 , not significant), and the split runs *within* cattle. Separating the range and feedlot legs of the cattle bundle in the raw County Agriculture Census records point-identifies the carve-out rather than assuming it: at high pressure the extensive leg collapses (total cattle and calves per acre ≈ -0.5 ; pastureland ≈ -0.25) while the confinement leg expands (hogs $+0.36$ to $+0.47$ across GAEZ surfaces, and the cattle-on-feed share and head rise under pressure – a levels estimate on the covered window whose sign is decisive). The A3 carve-out of dairy and feedlots is thus an empirically discriminated fact, not only a maintained assumption: the bar sheds range beef, for which land is

the input, and leaves land-light confinement, which meets the land-rent margin only as a footprint, untouched to expanding.

What the three signatures jointly rule out. A pure composition story (development skimming the worst land) implies revenue per acre rises and input use is flat; it fails Signatures 1 and 2. An access-intensification story implies input use rises at $M_c = 0$; it fails Signature 1. A story in which any animal agriculture is displaced implies confinement falls with grazing; it fails Signature 3. Only the unified land-rent sorting – one driver, the rising bar set by R^O , with activities sorted by value per acre against the bar – produces all three signatures together, which is the model’s falsifiable content.

Appendix B.10 Welfare Decomposition

This section derives the county-level welfare decomposition implied by the model and connects it to the farmland value aggregate used by [Donaldson and Hornbeck \(2016\)](#) to measure the welfare gains from market integration. The welfare object is the one-sector aggregate of [Donaldson and Hornbeck \(2016\)](#), so within this section $b(\tau)$ and $c(\tau)$ denote the activity-aggregated farm-gate return and cost of the surviving base; at the exit margin the marginal use is livestock, so the indifference condition is $b_L A_c \omega^* - c_L = \rho^\sigma$, where ρ^σ is the stratum’s best non-farm option – the development bid $R_{ct}^O - \kappa_{ct}$ in the fringe and the reversion $R_{ct}^N - \kappa_N$ beyond it. The exiting parcel is valued at ρ^σ , so the county outside-use component is the M_c mixture of the developed value R^O (fringe) and the reversion value R^N (beyond). Each per-parcel outside option is non-negative, but the aggregate omission ΔW^O is an estimated object whose sign is left to the data: it nets the outside-use legs against the identified in-agriculture value change, and its confidence interval can span zero.

Aggregate county welfare. Aggregate county-level land value, the welfare object in this paper, is the integral of parcel values over the within-county productivity distribution:

$$W_{ct} = \int_0^\infty V_{ct}(\phi_{jc}) \bar{s}_c dF(\omega), \quad (59)$$

where $\bar{s}_c$ is mean parcel acreage (held fixed for clarity; the acreage-weighted version follows analogously) and $V_{ct}(\phi_{jc})$ is the *capitalized* parcel value of equation (43) (each option already divided by r), so no additional discount factor multiplies the integral. Substituting equation (43) and splitting at the participation threshold — each capitalized option carries its own $1/r$, which factors out as the $\bar{s}_c/r$ prefactor below:

$$W_{ct} = \underbrace{\frac{\bar{s}_c}{r} \int_{\omega_{ct}^*}^\infty [b(\tau_{ct}) A_c \omega - c(\tau_{ct})] f(\omega) d\omega}_{W_{ct}^A} + \underbrace{\frac{\bar{s}_c}{r} \int_0^{\omega_{ct}^*} \rho_{ct}^\sigma f(\omega) d\omega}_{W_{ct}^O}. \quad (60)$$

Equation (60) is written for a single development-pressure stratum σ , whose exiting parcels take the best non-farm option ρ_{ct}^σ (the development bid $R_{ct}^O - \kappa_{ct}$ in the fringe,

the reversion $R_{ct}^N - \kappa_N$ beyond it); the county object is the area-weighted mixture $W_{ct}^O = M_c W_{ct}^{O,\text{press}} + (1 - M_c) W_{ct}^{O,\text{none}}$, with the fringe stratum carrying developed value and the non-fringe stratum the reversion value; each stratum's outside option is non-negative per parcel, while the aggregate omission is estimated with a confidence interval that can span zero. The first term W_{ct}^A is the agricultural component: the capitalized agricultural returns earned by parcels remaining in production. The second term W_{ct}^O is the outside-use component: the capitalized outside-use returns (net of conversion costs) earned by parcels that have exited agriculture. The [Donaldson and Hornbeck \(2016\)](#) farmland value aggregate sums capitalized values over parcels classified as farmland; under the model this aggregate equals W_{ct}^A rather than the full welfare object W_{ct} . The two coincide ($W_{ct}^A \approx W_{ct}$) when the outside-use share W_{ct}^O/W_{ct} is small, as in the open-frontier setting of [Donaldson and Hornbeck \(2016\)](#); in the settled-frontier setting where W_{ct}^O is non-trivial, W_{ct}^A is a downward-biased proxy for W_{ct} .

Assumption A4 (Exit-valuation pricing). The inframarginal exited mass on $[0, \omega_{ct}^*)$ in W_{ct}^O is valued at the common stratum outside rent ρ_{ct}^σ , as equation (60) writes it; within-stratum heterogeneity in the developer bids that inframarginal parcels command is set to zero.

A4 governs the inframarginal mass only. The *boundary* parcel at ω_{ct}^* is priced at ρ_{ct}^σ by the indifference identity $b_L A_c \omega_{ct}^* - c_L = \rho_{ct}^\sigma$ that appears in the Leibniz boundary term of equation (65) — that valuation is forced, not assumed. The assumption extends the boundary price to the mass strictly interior to the exit region. The alternative, in which inframarginal exited parcels command heterogeneous bids above the marginal rent, adds a non-negative level term to W_{ct}^O : the gap is the per-parcel developer wedge Δ^O integrated over the held fringe slice $\eta \in [0, 1]$, and is bounded above by Δ^O/η rather than governed by the soil-flat outside-option structure of A2. A4 sets that gap to zero; relaxing it moves the omission up by a non-negative amount within the stated bound. This is the maintained assumption behind Proposition 1's capitalization reading and the fringe level channel of Section 7.3.

Fréchet closed forms. Under the Fréchet specification (19), both components have explicit expressions. Using the tail moment (21):

$$W_{ct}^A = \frac{\bar{s}_c}{r} \left[b(\tau_{ct}) A_c \gamma\left(1 - \frac{1}{\theta}, (\omega_{ct}^*)^{-\theta}\right) - c(\tau_{ct}) S_{ct}^A \right], \quad (61)$$

$$W_{ct}^O = \frac{\bar{s}_c}{r} \rho_{ct}^\sigma (1 - S_{ct}^A), \quad (62)$$

where $\gamma(a, x) = \int_0^x t^{a-1} e^{-t} dt$ is the lower incomplete gamma function and $S_{ct}^A = 1 - \exp(-(\omega_{ct}^*)^{-\theta})$ from equation (34).

Response to market access. Differentiating the decomposition with respect to τ :

$$\frac{\partial W_{ct}}{\partial \tau} = \frac{\partial W_{ct}^A}{\partial \tau} + \frac{\partial W_{ct}^O}{\partial \tau}. \quad (63)$$

The outside-use component responds:¹⁷

$$\frac{\partial W_{ct}^O}{\partial \tau} = \frac{\bar{s}_c}{r} \left[\frac{\partial \rho_{ct}^\sigma}{\partial \tau} (1 - S_{ct}^A) + \rho_{ct}^\sigma f(\omega_{ct}^*) \frac{\partial \omega_{ct}^*}{\partial \tau} \right]. \quad (64)$$

On the pressured subgroup: when access improves (τ falls), $\partial R^{O, \text{press}} / \partial \tau < 0$ (first term, W^O rises from higher outside returns) and $\partial \omega^* / \partial \tau < 0$ (second term, W^O rises further because more parcels convert to outside use). Both contributions to $\partial W^O / \partial \tau$ are negative, so W_{ct}^O rises unambiguously on the pressured subgroup when access improves.

The agricultural component responds:

$$\frac{\partial W_{ct}^A}{\partial \tau} = \frac{\bar{s}_c}{r} \left[\underbrace{\frac{\partial b}{\partial \tau} A_c \gamma \left(1 - \frac{1}{\theta}, (\omega_{ct}^*)^{-\theta}\right)}_{\text{revenue}} - \underbrace{\frac{\partial c}{\partial \tau} S_{ct}^A}_{\text{cost}} - \underbrace{\rho_{ct}^\sigma f(\omega_{ct}^*) \frac{\partial \omega_{ct}^*}{\partial \tau}}_{\text{composition (threshold movement)}} \right]. \quad (65)$$

The first two terms are the direct revenue and cost effects on parcels remaining in agriculture. The third term is the Leibniz boundary contribution from the moving participation cutoff: a marginal (livestock) parcel transitioning between agricultural and outside use contributes its net return at the threshold, $b_L(\tau_{ct}) A_c \omega_{ct}^* - c_L(\tau_{ct}) = \rho_{ct}^\sigma$, weighted by the density $f(\omega_{ct}^*)$ and the threshold movement $\partial \omega_{ct}^* / \partial \tau$. This boundary term is generally nonzero in the interior region.¹⁸

The envelope identity that cancels threshold-movement contributions operates at the level of *total* welfare $W_{ct} = W_{ct}^A + W_{ct}^O$, not within $\partial W_{ct}^A / \partial \tau$ alone. The boundary term $-\rho_{ct}^\sigma f(\omega_{ct}^*) \partial \omega_{ct}^* / \partial \tau$ in equation (65) cancels exactly against the matching boundary term $+\rho_{ct}^\sigma f(\omega_{ct}^*) \partial \omega_{ct}^* / \partial \tau$ in equation (64), because the integrands of W_{ct}^A and W_{ct}^O coincide at the participation cutoff by indifference. The aggregate $\partial W_{ct} / \partial \tau$ therefore has no threshold-movement piece; the component-level derivatives retain their composition contributions.

The sign of $\partial W_{ct}^A / \partial \tau$ depends on the balance of three effects: the revenue effect (under net reach, $\partial b / \partial \tau < 0$ raises W^A via higher agricultural returns; the net sign is the A1 incidence object), the cost effect ($\partial c / \partial \tau$ ambiguous), and the composition effect (the sign of $-\rho_{ct}^\sigma f(\omega_{ct}^*) \partial \omega_{ct}^* / \partial \tau$ tracks the direction of conversion at the participation margin). On the unpressured subgroup, $f(\omega_{ct}^*)$ is small because parcels sit far from the participation threshold, so the composition effect is small and the revenue effect governs: under net reach $\partial W_{ct}^A / \partial \tau < 0$, so W_{ct}^A rises when access improves. On the pressured subgroup, $f(\omega_{ct}^*)$ is non-negligible because more parcels sit near the margin; the composition term is larger and W_{ct}^A can fall or rise depending on which of the three effects dominates.

¹⁷The conversion wedge κ_{ct} is treated as exogenous to τ_{ct} throughout: $\partial \kappa / \partial \tau = 0$. This is appropriate when κ represents time-invariant institutional frictions (zoning, transaction costs). If highway access reduces physical development costs, $\partial \kappa / \partial \tau > 0$ would add a term $-(1/r)(1 - S^A)(\partial \kappa / \partial \tau)$ to $\partial W^O / \partial \tau$, reinforcing the outside-use channel in the direction of the main results.

¹⁸Interior solution requires $0 < \omega_{ct}^* < \infty$, $0 < S_{ct}^A < 1$ (some but not all parcels are in agriculture), and $\rho_{ct}^\sigma > 0$ (the outside option net of conversion costs is positive). The model predictions are stated for this interior region, which corresponds empirically to counties where the Agricultural Census records both farmed and non-farmed land.

What the farmland value aggregate reveals and conceals. The total welfare gain in the model is $\Delta W_{ct} = \Delta W_{ct}^A + \Delta W_{ct}^O$. The DH2016 farmland value aggregate measures ΔW_{ct}^A alone. An access improvement raises each subgroup's per-parcel outside option ($\partial \rho^\sigma / \partial \tau < 0$) as parcels at the margin convert, and the omitted value runs on a *different* outside option in each stratum:

- On the unpressured subgroup: the marginal land that leaves farming reverts to the non-development use R^N , and this reverted land is the omission ΔW_{ct}^O the Census does not record. Marginal land that access instead holds in grazing (Prediction 1) does not leave farming and is not part of this sum – it enters W_{ct}^A as retained farmland. The reverted-land omission carries the reversion value $R^N \geq 0$ per parcel and is the level channel of Section 7.3. The difference-in-differences identifies the treatment *response* and nets out the level, but the *level* of this omission inherits the period average of R^N across the time-varying reversion regime (Soil Bank, the near-floor interprogram years, and the Conservation Reserve Program); the calibrated no-pressure omission is reported against an observed reversion value rather than a freely estimated one, and is disclosed as such.
- On the pressured subgroup: exiting marginal land develops at the bid R^O , but only the slice that actually converts in the census wave carries the full developed value. A parcel priced out of grazing but held at the option – crop shed, land coasting on low-input cover, conversion realized only later – carries a held fraction $\eta \in [0, 1]$ of the discounted development option, $\eta(R^O - \kappa)$, so its contribution to the omission is the interval $[0, R^O - \kappa]$, parallel to the reversion floor $R^N \geq 0$ rather than a point. The census five-year snapshot cannot separate η from R^O on the held slice, so η is carried as a bound and not estimated. The land value gain reflects outside-use conversion, agriculture contracts even as total welfare rises, and the farmland value aggregate understates the true gain on the converting slice. The omission counts development value *incremental* to the option value already capitalized into the observed prices of surviving near-fringe farmland (Prediction 4 and Proposition 1): the difference-in-differences nets out that pre-existing capitalized level, so ΔW^O is the treatment response of the exit-valuation level rather than the full developed rent R^O counted a second time.

Gross conversion versus net farmland change. The difference-in-differences reads the *net* farmland change $\Delta S_{net}^A = \Omega_{ent} - \Omega_{dev} - \Omega_{rev}$, which nets farm entry and re-entry (Ω_{ent} , reverted land pulled back in as reach raises b_L) against gross farm-to-development conversion (Ω_{dev}) and farm-to-reversion exit (Ω_{rev}). The welfare omission needs the gross object Ω_{dev} alone, priced at ρ^{High} . On the net-contraction cells — the fringe conversion cells where farmland sheds — entry and reversion partially offset a larger gross exit, so $\Omega_{dev} = |\Delta S_{net}^A| + \Omega_{ent} + \Omega_{rev} \geq |\Delta S_{net}^A|$: pricing the omission off net acreage *understates* gross conversion there, a conservative reading. The sign flips on the retention cells: on the marginal no-pressure stratum farmland net *rises* ($\Delta S_{net}^A > 0$, entry dominates the negligible fringe conversion), so net is not the relevant object and the omission is approximately

zero. The county-net masks the simultaneity by mixing a rising-farmland retention stratum with a shedding conversion stratum; net acreage is a lower input on Ω_{dev} , and a direct development-extent measure is the loose upper input, over-counting by the share of access-induced agricultural structures it cannot label.

The farmland value aggregate therefore understates ΔW_{ct} on *both* strata – through development R^O in the fringe and through reversion R^N beyond it – and the split of the omission between the two is the empirical question the lower-bound result of Section 7.3 resolves. Where the frontier is still open and the outside-use term W_{ct}^O is small, the aggregate tracks total welfare well; on settled, converting land it leaves out a sizable outside-use piece. Splitting the treatment effects by pressure recovers what the single aggregate hides: it reads the farming component ΔW_{ct}^A off crop revenue and farm counts and the outside-use component ΔW_{ct}^O off land conversion and non-farm activity, separately. The conclusion applies this decomposition to the paper’s ATT estimates.

From the decomposition to the reported bound. The omission ΔW^O is not reported as a point. Its per-acre price is the outside rent $\rho^\sigma = R^O - \kappa$ on the fringe, and the design identifies only the sum $X_{\text{can}} = R^O - \kappa + c$ (the near-singular c/R block of Appendix B.11), not $R^O - \kappa$ alone. The GMM primitives alone therefore do not separate the fringe cost from the outside-option response, and the omission is set-identified rather than pinned to a point. Its magnitude enters only through the external, disclosed development-value anchor: the observed developer bid R^O read off fringe property values re-prices the converted exit stock, which places the omission at as much as \$195.0–210.1 billion across surfaces. The standard farmland-value aggregate is a lower bound on the welfare gain, and pinning the omission’s precise magnitude is the subject of ongoing work.

Appendix B.11 Structural Identification: A Minimum-Distance GMM Specification

This subsection formalizes the structural identification step invoked in Section 7.3. It defines the channel primitives, states the ATT-to-derivative bridge, writes out the moment system explicitly, demonstrates linearity, and provides the closed-form minimum-distance GMM estimator together with the Hansen J over-identification test.

Specification note. The canonical structural estimator is the *pooled two-activity moment system*, estimated on a common sample. The four outcomes ($\ln V^A, \ln S^A, \ln \overline{\text{Rev}}^C, \ln \overline{\text{Rev}}^L$) — farmland value per acre, farmland acreage, crop revenue, and livestock revenue — are stacked across K pressure strata to form the moment vector, and the channel primitives — the two farm-gate elasticities (π_{b_C}, π_{b_L}) held common across strata, together with stratum-specific cost and outside-option elasticities ($\pi_c^\sigma, \pi_R^\sigma$) for each pressure stratum σ — are estimated by minimum distance. Pressure is resolved into $K = 3$ strata by a median split of positive development-pressure exposure — {None, Low, High}, where *None* collects counties with no city-reach exposure and *Low/High* split the counties with positive exposure at the median of M_c . The per-cell (suitability $\times$ pressure) cellwise imputation grid lacks identifying power in the thin prime (S1) cells (Appendix A.4) and enters only as demoted robustness (the pooling pools pressure over

suitability, using suitability as a first-stage trend control rather than a second-stage interaction). The moment vector is read off one common estimation sample — every county-year enters all four base moments non-missing — so the four ATTs share a single sample and a disclosure-suppressed crop-revenue subset cannot leak leverage into one moment alone. The weight matrix is a stable-county panel bootstrap with $\hat{\Omega}^{-1}$ dense across moments; the pressure moderator M_c is the predetermined city-reach measure of Section 4.1 at the 25,000-population cutoff, and inference is Conley (1999) 100 km spatially corrected. The specific numerical results reported below — the Hansen J , the calibration table, and the omitted-share band — are estimated under this specification with the winsorized bootstrap (1,000 resamples); the per-cell suitability $\times$ pressure leave-out is a demoted robustness left to ongoing work.

Channel primitives. Let $\sigma \in \{\text{None}, \text{Low}, \text{High}\}$ index development-pressure status, the $K = 3$ median split of the predetermined city-reach moderator M_c : *None* collects counties with no city-reach exposure, and *Low/High* split the counties with positive exposure at the median of M_c . The structural identification works at the level of semi-elasticities of the bundled access shock, with the farm-gate return split by activity into its crop and livestock components:

$$\pi_{b_C} \equiv \frac{\partial b_C}{\partial \tau}, \quad \pi_{b_L} \equiv \frac{\partial b_L}{\partial \tau}, \quad \pi_c^\sigma \equiv \frac{\partial c^\sigma}{\partial \tau}, \quad \pi_R^\sigma \equiv \frac{\partial R^{O,\sigma}}{\partial \tau}, \quad (66)$$

collected in the primitive vector $\pi = (\pi_{b_C}, \pi_{b_L}, \{\pi_c^\sigma\}_\sigma, \{\pi_R^\sigma\}_\sigma)'$. Output prices for agricultural commodities clear in national markets, so the two farm-gate elasticities π_{b_C}, π_{b_L} are held common across strata; input markets for labor and land services clear locally, so π_c varies across development-pressure status; outside-use returns respond to residential and commercial demand where conversion pressure is high, so π_R varies across development-pressure status. An alternative specification that relaxes the homogeneity restriction on the farm-gate elasticities across development-pressure status is described in the robustness paragraph below; the welfare bound is identified consistently across both specifications.

In the canonical (base) specification π_c^σ is an expenditure object, not an input-price elasticity. Cost is calibrated as the residual $c^\sigma = \overline{\text{Rev}} - V_{\text{flow}}^A$ between gross revenue per acre and the value flow, so any force that raises gross output per surviving acre without a matching rise in the capitalized value flow lands in c^σ . Intensification is such a force: a rising land-rent bar raises inputs per acre and gross output while value per acre stays flat, so the intensification quantity response leaks into π_c^σ , and it does so unevenly across strata because its magnitude scales with development pressure. The estimated π_c^σ is therefore a per-acre-cost response — price times quantity — rather than a materials-price primitive. The price reading is recovered only under a σ -augmented specification that gives intensification its own input-intensity moment — a calibrated CES land-input margin (the $\partial \ln(x^*/\ell^*)/\partial \ln R = \sigma$ clause of A3) carrying the pressure-varying quantity response — and leaves the residual cost response as a common materials price. That augmentation is a robustness register, not the canonical estimator; the welfare bound reads off the reduced-form π_c^σ regardless of its price/quantity decomposition.

From ATTs to derivatives. The DID2s estimator of Section 5 produces, for each cell k and each outcome $Y \in \{\ln V^A, \ln S^A, \ln(V^A S^A)\}$, an average post-treatment percent change $\hat{\beta}_k^Y = \mathbb{E}[\Delta \ln Y_{ct} \mid \text{cell } k, \text{ post-IHS}]$. We map these empirical objects to the model’s comparative-static derivatives via a first-order linearization,

$$\Delta \ln Y_{ct} \approx \frac{\partial \ln Y_{ct}}{\partial \tau} \cdot \Delta \tau, \quad (67)$$

where $\Delta \tau$ is the average access-wedge change induced by the interstate opening, treated as a constant common across cells. The identification target is the bound ratio $\Delta W^O / \Delta W$. Because ΔW^O and $\Delta W = \Delta W^A + \Delta W^O$ are both linear combinations of the same primitives, $\Delta \tau$ enters the numerator and denominator multiplicatively and cancels. The bound is invariant to the normalization of $\Delta \tau$; only π is identified up to this constant scale. We therefore set $\Delta \tau = 1$ in implementation and report π in units such that the bound ratio is dimensionless.

Pooled identification of the structural moment vector on a common sample. The pooled moment system in equation (68) maps *stratum-specific* ATTs to the primitive vector π through a design matrix whose rows are indexed by the stratum-level baseline configuration $(A_k, S_k^A, \omega_k^*, f_k, b, c^{\sigma(k)}, R_k^O - \kappa_k)$, with the strata resolved by the $K = 3$ median split of development pressure {None, Low, High} and suitability entering as a first-stage trend control rather than a second-stage interaction. Following Gardner et al. (2024) and Borusyak et al. (2024), the empirical counterpart that delivers $\hat{\beta}_k^Y$ is a two-stage imputation estimator whose first-stage projection is fit on the untreated subsample and whose counterfactual trajectory $\hat{Y}_{j,t}^*$ is evaluated at the treated covariate values. Maintained Condition 2 of Appendix B.4 maintains within-stratum homogeneity of the primitives across suitability class. All four base outcomes are read off one common estimation sample — each 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; requiring the harvested-cropland denominator that defines crop revenue removes the disclosure-suppressed cells whose leverage otherwise inflates the crop moment on the high-suitability surface. The identifying support for these stratum moments — the effective number of treated counties behind each share-weighted ATT — is documented in Appendix A.4. The pooled two-activity system is the canonical structural specification; the per-cell suitability $\times$ pressure grid is a demoted robustness because it lacks identifying support in the thin prime (S1) cells (Appendix A.4; the pooling pools pressure over suitability by the dimension the DID2s second stage varies). The Hansen J and the omission band are estimated under the pooled winsorized specification with the final bootstrap: $J = 14.7$ ($p = 0.0055$, df 4) on the rainfed-low surface, which rejects the over-identifying restrictions at the 5 % level, and $J = 9.2$ ($p = 0.056$, df 4) on the rainfed-high surface, which is borderline under the χ_4^2 reference. The misfit is carried by the high-pressure stratum: dropping either the farmland-acreage or the livestock-revenue moment in that stratum collapses J symmetrically, and a bootstrap reference distribution sharpens rather than softens the rejection (detailed below, with Tables A16 and A15). The rejection therefore impugns point identification of the eight primitives rather than the lower-bound reading the decomposition commits to (Section 7.3),

which is invariant to the moment dropped and governed by the thick no-pressure stratum that fits. The θ -heterogeneity robustness paragraph below probes whether the bound itself is sensitive to the pooled- θ assumption.

Baseline calibration of V^A : net of input costs. The structural object V^A is the per-acre rent on parcels in agriculture at the cell threshold, with flow $V_{\text{flow},k}^A = bA_k m_k - c^{\sigma(k)}$ – where $m_k = \mathbb{E}[\omega \mid \omega \geq \omega_k^*]$ is the conditional mean productivity above the threshold, not the threshold itself – and stock $V_{\text{stock},k}^A = V_{\text{flow},k}^A / r$ by the parcel-level capitalization in equation (41). We anchor the calibration by setting $V_{\text{stock},k}^A$ to the cell-level mean real 2025-USD farmland value per acre and inverting at $r = 0.05$, which gives a cell-level flow of \$30 to \$91 per acre per year, rising from the marginal no-pressure cell to the suitable pressure cell. We anchor V_{flow}^A to the capitalized farmland rent rather than to gross agricultural product value per acre, which would conflate output with input costs and inflate the structural level. Under the recalibration the M2-step constants (b, c^σ) and the elements of $\hat{\pi}$ rescale to their structural units; the model’s empirical fit is unchanged and only the dollar-level interpretation of π is affected. The captured in-agriculture value change aggregated across all IHS-treated counties on the rainfed-low surface is \$54.3 billion (stock, 2025-USD capitalized at $r = 0.05$; winsorized bootstrap). Stratum baseline calibrations are reported in Table A12.

The omission *share* of welfare is scaled against the per-farmland-acre base that matches the capitalized V^A stock, not a per-county-acre base; the absolute omission dollars are unaffected by the denominator choice. The GMM primitives alone do not pin a positive omission magnitude on either surface, so the development-value magnitude enters only through the external, disclosed anchor. The aggregate omission is a corroborated bound rather than a point estimate, and the conditions under which the bound holds are stated explicitly here rather than folded into a fit claim. The distinction between what is imposed and what emerges is exact and worth stating plainly: the *component* sign restrictions of Assumption A5 below — $\pi_R^{\text{High}} \leq 0$ and $\pi_c^{\text{High}} \geq 0$ — are imposed, while the *aggregate* non-negativity of ΔW^O is not imposed and emerges (or fails to emerge) from the estimated system, so its confidence interval can span zero. A non-negative aggregate is corroboration precisely because it is produced by the estimated primitives rather than read off the accounting map.

The aggregate omission is $\Delta W^O = \sum_k dW_k^O$, and the fringe cost and outside-option primitives $(\pi_c^{\text{High}}, \pi_R^{\text{High}})$ do not separately point-identify (Appendix A.4 and the near-singular c/R separating lever below). Four features make the bound honest without that separation. First, the None stratum self-identifies: it is thick, its value flow carries no development premium, and its omission rides on a calibrated reversion value R^N , for which $R^N = 0$ is a guaranteed lower bound since the reversion use returns a non-negative flow. Second, the fringe stratum enters ΔW^O through a *direct* leg $\pi_R^{\text{High}}(1 - S^A)$ that carries π_R^{High} alone — not the identified $(\pi_c + \pi_R)^{\text{High}}$ sum, which enters only the composition leg. Two sign restrictions, stated as an assumption because they are imposed rather than tested, place a finite bound on that direct leg.

Assumption A5 (Fringe sign restrictions). In the high-pressure fringe the development

response is non-positive, $\pi_R^{\text{High}} \leq 0$, and the operating-cost response is non-negative, $\pi_c^{\text{High}} \geq 0$.

The first restriction reads improved access as raising the developer’s bid on fringe land; the second reads it as weakly raising per-acre operating cost under the intensification-dominance reading of A3, in which access raises input use per surviving acre. Combined with the identified acreage sum $(\pi_c + \pi_R)^{\text{High}}$, A5 places π_R^{High} in the finite interval $[(\pi_c + \pi_R)^{\text{High}}, 0]$, so the fringe contribution is set-identified with finite width even though its primitives are not. The $\pi_c^{\text{High}} \geq 0$ jaw is the weaker of the two: a highway that lowers the delivered price of purchased inputs or the haulage cost of moving output would push $\pi_c^{\text{High}} < 0$, widening the interval and raising the omission bound. The bound is therefore reported under π_c^{High} unrestricted and under a small negative value as a sensitivity. Third, the shock size $\delta\tau$ is anchored on the flat $\ln V^A$ moment, so the omission level inherits that anchor rather than a free scale. The remaining condition is numerical and confirmed on the estimated system: the no-pressure stratum contributes *zero* to ΔW^O at the reported bound — absent development pressure there is no option value to capitalize, and its reversion is held at the guaranteed lower bound $R^N = 0$ — so the thick no-pressure mass cannot inflate the aggregate. The direct leg is what the difference-in-differences identifies, and it carries π_R^{High} alone; because π_R^{High} is weakly identified (its t -statistic is below one, the near-singular c/R block below), the identified direct leg is near zero and its point sign is fragile. The GMM alone therefore does not pin a positive omission magnitude; the development-value magnitude enters only through the external, disclosed anchor. The per-stratum direct-leg contributions are correspondingly near zero. Interior-stratum intensification enters the composition leg rather than the direct leg, and the crop-revenue disclosure artifact that would inflate it is removed on the estimation sample, so the fringe bound does not ride on the intensive margin. Under the two sign restrictions and this confirmation the aggregate omission is a corroborated bound, consistent with the reading of Section 7.3.

The bound rests on how much of the omission the design identifies versus an external property anchor. Priced at $\eta = 0$ from the difference-in-differences alone, the direct exit-valuation leg uses the GMM primitives and delivers the [Donaldson and Hornbeck \(2016\)](#) lower bound; the composition (flux) leg does not enter, because by the envelope identity it cancels in total welfare (the Leibniz boundary term of the welfare decomposition above; the three-piece footnote below) and is a conserved transfer, not an omitted level. Because the direct leg carries the weakly identified π_R^{High} , the GMM alone does not pin a positive omission magnitude on either surface. The magnitude enters only through the external, disclosed development-value anchor: the developer bid R^O ($\approx 8.6\times$ on the fringe), read off observed fringe property values ([Haines et al., 2018](#)), re-prices the converted exit stock as a level term with no boundary piece — not the cancelling flux leg re-priced. That anchor is not point-identified by the design, because a spillover-clean never-treated county resembling a treated fringe county essentially does not exist and the matched comparator is degenerate, so it is disclosed as an anchored “at most” rather than an estimate; the observed fringe premium is flat before and after the highway, so valuing induced conversion at observed prices does not substantially double-count the highway’s own contribution,

and the maximum double-count is bounded by the small post-minus-pre gap. Valued at the observed developer bid, the omission reaches as much as \$200.7–210.1 billion on the rainfed-low surface (PRE/POST) and \$195.0–199.0 billion on the rainfed-high surface. Pinning its precise magnitude within these bounds is the subject of ongoing work.

Table A12: Per-Cell Calibration Summary

Cell	n	$V_{\text{stock}}^A/\text{acre}$	S^A	A_k	ω^*	m_k	f_k
Marginal (N), none	518	\$600	0.386	0.48	1.400	2.849	0.455
Marginal (N), pressure	276	\$1,033	0.555	0.48	1.105	2.361	0.694
S3, none	906	\$515	0.638	0.96	0.993	2.190	0.788
S3, pressure	998	\$843	0.665	0.96	0.958	2.139	0.814
Suitable (S2+S1), none	646	\$913	0.883	1.56	0.698	1.816	0.764
Suitable (S2+S1), pressure	748	\$1,333	0.844	1.56	0.747	1.867	0.825

Notes. Per-stratum baseline calibrations at the data-anchored $\hat{\theta} = 2.13$ (rainfed-low surface) for the $K = 3$ pressure strata {None, Low, High}, the median split of positive development-pressure exposure; suitability enters as a first-stage trend control rather than a stratum interaction. n is the number of counties in each stratum; $V_{\text{stock}}^A/\text{acre}$ is the per-acre real 2025-USD farmland value at the cell-level mean (the anchor for the flow-stock calibration above); S^A is the cell-level mean farmland share of county area; A_k is the model suitability scale parameter, increasing across the three classes; ω^* , m_k , and f_k are the Fréchet participation threshold, conditional mean productivity above the threshold, and density at the threshold, all functions of $\hat{\theta}$ and the cell-level baseline. The calibrated ω^* inverts each cell's *observed* S^A through equation (7) and so embeds cross-cell composition and selection differences; it therefore need not preserve the within-county comparative-static ordering $\omega_{\text{press}}^* > \omega_{\text{none}}^*$ of equation (8), which compares the two subgroups of a *single* county at fixed suitability rather than cells that differ in baseline composition.

The pooled two-activity moment system. Index the $K = 3$ pressure strata by

$$k \in \{\text{None, Low, High}\},$$

the median split of positive development-pressure exposure, and the base outcomes by $Y \in \{V, F, \text{Rv}^C, \text{Rv}^L\}$ for $\ln V^A$, $\ln S^A$, crop revenue, and livestock revenue respectively. The model implies a pooled moment vector $g_k^Y(\pi)$ linear in π , stacking the four outcomes across the three strata; under the final $K = 3$ two-activity specification the design matrix has twelve rows — four outcomes stacked across the three pressure strata. The loadings below are written for a single agricultural-value activity, in the $\{V, F, T\}$ notation of the value, farmland-share, and product ($T = \ln V^A S^A$) rows, to fix the algebra; this single-activity presentation is pedagogical, and the estimated system carries the four base outcomes ($\ln V^A, \ln S^A, \ln \overline{\text{Rev}}^C, \ln \overline{\text{Rev}}^L$) once the livestock row is added. The two-activity extension adds a livestock-revenue row with its own composition kernel and does not change the linearity of the system in π . Writing $\omega_k^*, S_k^A, T_k = \gamma(1 - 1/\theta, (\omega_k^*)^{-\theta})$, $m_k = T_k/S_k^A$, $f_k = f(\omega_k^*)$, $V_k^A = (1/r)[bA_k m_k - c^{\sigma(k)}]$, and $D_k = bA_k T_k - c^{\sigma(k)} S_k^A$, and using

$X_{\text{can},k} \equiv R_k^O - \kappa_k + c^{\sigma(k)}$ for the canonical threshold numerator,

$$g_k^Y(\pi) = a_{b,k}^Y \pi_b + a_{c,k}^Y \pi_c^{\sigma(k)} + a_{R,k}^Y \pi_R^{\sigma(k)}, \quad k \in \{\text{None, Low, High}\}; Y \in \{V, F, T\}. \quad (68)$$

The coefficients $a_{*,k}^Y$ are determined by the cell-level baseline calibration $(A_k, S_k^A, \omega_k^*, b, c^{\sigma(k)}, R_k^O - \kappa_k, \theta)$ and follow directly from differentiating equations (61)–(62). Their closed forms appear in Table A13.

Table A13: Design Matrix Coefficients for the Pooled GMM System (Single-Activity Rows)

Outcome Y	Loading	$a_{b,k}^Y$	$a_{c,k}^Y$	$a_{R,k}^Y$
$V (\ln V^A)$	direct	$\frac{A_k m_k}{r V_k^A} - \frac{A_k f_k (m_k - \omega_k^*) \omega_k^*}{r V_k^A S_k^A}$	$\frac{b A_k f_k (m_k - \omega_k^*) \omega_k^*}{r V_k^A S_k^A X_{\text{can},k}} - \frac{1}{r V_k^A}$	$\frac{b A_k f_k (m_k - \omega_k^*) \omega_k^*}{r V_k^A S_k^A X_{\text{can},k}}$
$F (\ln S^A)$	threshold	$\frac{f_k \omega_k^*}{S_k^A b}$	$-\frac{f_k \omega_k^*}{S_k^A X_{\text{can},k}}$	$-\frac{f_k \omega_k^*}{S_k^A X_{\text{can},k}}$
$T (\ln V^A S^A)$	sum	$\frac{A_k T_k}{D_k} + \frac{f_k \omega_k^* (R_k^O - \kappa_k)}{D_k b}$	$-\frac{S_k^A}{D_k} - \frac{f_k \omega_k^* (R_k^O - \kappa_k)}{D_k X_{\text{can},k}}$	$-\frac{f_k \omega_k^* (R_k^O - \kappa_k)}{D_k X_{\text{can},k}}$

Notes. The design matrix stacks one row per stratum-outcome moment; entries $a_{c,k}^Y, a_{R,k}^Y$ load the row onto the stratum-appropriate primitive $\pi_c^{\sigma(k)}$ or $\pi_R^{\sigma(k)}$, with zero coefficient on the other stratum. The rows shown are the single-activity value, threshold, and sum loadings; the pooled two-activity system replaces the single revenue row with *two* activity-resolved revenue rows — crop revenue (equation (69)), a clean π_b moment with $a_c^{\text{Rv}^C} = a_R^{\text{Rv}^C} = 0$) and livestock revenue (equation (70), loading the (c, R) sum through the exit cutoff) — and stacks across the three pressure strata, so the realized dimension of X under the final specification is 12×8 (twelve moments, eight primitives). The within-stratum algebraic identity $\partial \ln V^A / \partial \tau + \partial \ln S^A / \partial \tau = \partial \ln(V^A S^A) / \partial \tau$ implies $a_{*,k}^V + a_{*,k}^F = a_{*,k}^T$ for each primitive index $*$.

Two activity-resolved revenue moments. The realized system carries *two* revenue rows, one per farm activity, and they load on the structural primitives in qualitatively different ways — a distinction that is important for identification and that a single pooled revenue row conceals. Both are gross-output flows on surviving parcels, but crop revenue is governed by the *crop–livestock* cutoff ϕ_{CL} (equation (26)) while livestock revenue is governed jointly by ϕ_{CL} and the *farm-exit* cutoff ϕ^* . Write $\omega_{CL,k} = \phi_{CL} / A_k$ for the crop–livestock threshold in quality units, $T_{C,k}$ and $T_{L,k}$ for the crop and livestock survival masses, and $f_{CL,k} = f(\omega_{CL,k})$.

Crop revenue per acre is $\text{Rv}_k^C = b_C A_k T_{C,k} / S_k^A$, the gross output of the crop survivors $\omega \geq \omega_{CL,k}$. Its response to τ runs through the crop cutoff $\omega_{CL,k}$ alone, and $\phi_{CL} = (c_C - c_L) / (b_C - b_L)$ is invariant to the outside option R^O and, with the cost gap access-invariant under A1, moves only through the farm-gate return gap $\pi_{b_C} - \pi_{b_L}$. The three loadings of the crop-revenue outcome Rv_k^C are therefore

$$a_{b_C,k}^{\text{Rv}^C} = \frac{1}{b_C} - \frac{\omega_{CL,k} f_{CL,k}}{T_{C,k}} \cdot \frac{\omega_{CL,k}}{b_C - b_L}, \quad a_{b_L,k}^{\text{Rv}^C} = \frac{\omega_{CL,k} f_{CL,k}}{T_{C,k}} \cdot \frac{\omega_{CL,k}}{b_C - b_L}, \quad \boxed{a_{c,k}^{\text{Rv}^C} = a_{R,k}^{\text{Rv}^C} = 0}. \quad (69)$$

The boxed identity is the structurally important fact: *crop revenue carries zero loading on both π_c^σ and π_R^σ* . Crop survival is a return-gap phenomenon — a parcel drops out of crops when the crop–livestock frontier ϕ_{CL} moves, which depends on the return gap $b_C - b_L$ and not on the cost or outside-option responses — so the crop-revenue moment is a clean π_b instrument that identifies the farm-gate slopes without contamination from the fringe c/R split. This is why crop revenue does *not* separate π_c from π_R and why the c/R separation must come from the value moment alone.

Livestock revenue per acre is $Rv_k^L = b_L A_k T_{L,k} / S_k^A$, the gross output of the grazing band $[\omega_k^*, \omega_{CL,k}]$. It responds to *both* cutoffs: the farm-exit cutoff ω_k^* trims the band from below, and $\omega_{CL,k}$ trims it from above. The farm-exit cutoff runs on c and R through the canonical numerator $X_{\text{can},k} = R_k^O - \kappa_k + c^{\sigma(k)}$ ($\partial \omega_k^* / \partial c = \partial \omega_k^* / \partial R = \omega_k^* / X_{\text{can},k}$), so livestock revenue inherits a nonzero, symmetric (c, R) loading:

$$a_{b_L,k}^{Rv^L} = \frac{1}{b_L} + h_{*,k}^{Rv^L} \frac{-\omega_k^*}{b_L} + h_{CL,k}^{Rv^L} \frac{\omega_{CL,k}}{b_C - b_L}, \quad a_{c,k}^{Rv^L} = a_{R,k}^{Rv^L} = h_{*,k}^{Rv^L} \frac{\omega_k^*}{X_{\text{can},k}}, \quad (70)$$

with the composition kernels $h_{*,k}^{Rv^L} = -\omega_k^* f_k / T_{L,k}$ (from the exit cutoff) and $h_{CL,k}^{Rv^L} = \omega_{CL,k} f_{CL,k} / T_{L,k}$ (from the crop cutoff). As in the value row, c and R enter ω_k^* symmetrically through $X_{\text{can},k}$, so $a_{c,k}^{Rv^L} = a_{R,k}^{Rv^L}$: the livestock-revenue moment loads the (c, R) sum but does not itself split the pair. Neither revenue row carries the direct cost term $-1/(rV_k^A)$ that the value row does, since gross output nets no $-c$ to differentiate; that direct term is what makes the value moment — and only the value moment — separate π_c from π_R . The two revenue rows are not scalar copies of one another: crop revenue loads only (π_{b_C}, π_{b_L}) , while livestock revenue additionally loads the (c, R) sum through the exit cutoff, so the pair over-identifies the farm-gate slopes and the acreage sum jointly. The empirical comparison of the moment systems with and without each revenue block is reported separately.

Stratum ATT inputs. The empirical moment vector $\hat{\beta}$ stacks the pooled Gardner DID2s ATTs on the four base outcomes across the three pressure strata, read off the common estimation sample. Point estimates, standard errors, p -values, and sample sizes are reported in Table A14.

Linearity. The moment system is linear in π by construction. Every coefficient $a_{*,k}^Y$ in Table A13 is a function of the cell baseline $(A_k, S_k^A, \omega_k^*, b, c^{\sigma(k)}, R_k^O - \kappa_k, \theta)$ alone; no entry depends on π . Linearity follows from the linearity of the welfare derivatives (65)–(64) in the three semi-elasticities under the maintained Fréchet specification and the cell-level homogeneity restriction (Maintained Condition 2 of Appendix B.4: within each development-pressure stratum, primitives are common across suitability class).

Closed-form minimum-distance GMM. Stacking the four base outcomes across the three pressure strata yields the empirical moment condition

$$g(\pi) = \hat{\beta} - X\pi, \quad (71)$$

Table A14: Pooled Gardner DID2s ATTs: The Empirical Moment Vector $\hat{\beta}$

Stratum	Outcome	ATT	(SE)	p -value	Counties
Marginal (N), none	$\ln V^A$	0.126	(0.052)	0.015	518
Marginal (N), pressure	$\ln V^A$	-0.094	(0.139)	0.503	276
S3, none	$\ln V^A$	0.007	(0.049)	0.880	906
S3, pressure	$\ln V^A$	0.074	(0.051)	0.148	998
Suitable (S2+S1), none	$\ln V^A$	0.028	(0.048)	0.560	646
Suitable (S2+S1), pressure	$\ln V^A$	0.094	(0.032)	0.004	748
Marginal (N), none	$\ln S^A$	0.181	(0.047)	< 0.001	518
Marginal (N), pressure	$\ln S^A$	0.192	(0.126)	0.128	276
S3, none	$\ln S^A$	-0.015	(0.046)	0.748	906
S3, pressure	$\ln S^A$	-0.054	(0.046)	0.238	998
Suitable (S2+S1), none	$\ln S^A$	-0.059	(0.021)	0.005	646
Suitable (S2+S1), pressure	$\ln S^A$	0.074	(0.025)	0.003	748
Marginal (N), none	$\ln(V^A S^A)$	0.303	(0.057)	< 0.001	518
Marginal (N), pressure	$\ln(V^A S^A)$	0.089	(0.070)	0.205	276
S3, none	$\ln(V^A S^A)$	-0.008	(0.065)	0.904	906
S3, pressure	$\ln(V^A S^A)$	0.018	(0.034)	0.597	998
Suitable (S2+S1), none	$\ln(V^A S^A)$	-0.030	(0.051)	0.551	646
Suitable (S2+S1), pressure	$\ln(V^A S^A)$	0.166	(0.023)	< 0.001	748

Notes. Pooled Gardner two-stage difference-in-differences estimates ([Gardner et al., 2024](#)), one per stratum-outcome combination on the common estimation sample. Outcomes: $V = \ln V^A$ (farmland value per acre), $F = \ln S^A$ (farmland acreage), Rv^C (crop revenue), and Rv^L (livestock revenue). Strata: {None, Low, High}, the median split of positive development-pressure exposure M_c ; suitability enters as a first-stage trend control. “Counties” is the number of counties in the stratum. p -values from [Conley \(1999\)](#) 100 km spatial-correlation-corrected standard errors.

with $\hat{\beta}$ the empirical vector of stratum-level ATTs and X the pooled design matrix; the realized dimensions are $\hat{\beta} \in \mathbb{R}^{12}$ and $X \in \mathbb{R}^{12 \times 8}$. The minimum-distance estimator with weight matrix $W = \hat{\Omega}^{-1}$ admits the closed form

$$\hat{\pi} = (X' \hat{\Omega}^{-1} X)^{-1} X' \hat{\Omega}^{-1} \hat{\beta}, \quad \widehat{\text{Var}}(\hat{\pi}) = (X' \hat{\Omega}^{-1} X)^{-1}. \quad (72)$$

The closed form in equation (72) requires $X' \hat{\Omega}^{-1} X$ invertible. That inverse exists, but the $(\pi_c^{\text{High}}, \pi_R^{\text{High}})$ block is near-singular. The three moments that could in principle separate c from R do not: the farmland-share and livestock-revenue moments load c and R *identically* (both enter the exit cutoff ω^* symmetrically through X_{can}), and the crop-revenue moment loads *neither* — $a_c^{\text{Rv}^C} = a_R^{\text{Rv}^C} = 0$ by equation (69), since crop survival runs off the return-gap cutoff ϕ_{CL} , not the exit cutoff. The value moment alone carries the direct cost term $-1/(rV^A)$ that breaks the symmetry, so it is the sole separator of the pair, and it is the noisiest of the four base moments.¹⁹ The equation therefore point-identifies the farm-gate slopes and the acreage sums, but not the fringe c/R split; its stratum-by-stratum π_c^{High} and π_R^{High} point estimates are not interpreted. The identified objects on the fringe are the sum $(\pi_c + \pi_R)^{\text{High}}$ and, under Assumption A5, the set-identified bound on the direct leg; the point split is set aside without a regularization step, and no ridge correction is applied on the reported specification.

The one internal moment that would break the fusion. Because a single pooled value moment supplies only one separating equation, the c/R split stays near-singular. Resolving *value by activity* would supply a second. Let crop-land and pasture-land value per acre be $V_C = [b_C A_k m_C - c_C]/r$ and $V_L = [b_L A_k m_L - c_L]/r$, with m_C, m_L the conditional-mean qualities of the crop and grazing survivor bands. Each activity value moment carries an activity-specific direct cost term $-1/(rV_C)$ or $-1/(rV_L)$ on π_c , plus a symmetric composition loading through the exit cutoff, written with the activity-specific kernel $\Phi_k \equiv b A_k f_k(m_k - \omega_k^*) \omega_k^* / (r V_k S_k^A X_{\text{can},k})$ for $k \in \{C, L\}$. The (π_c, π_R) sub-block of the two rows is

$$\begin{pmatrix} \Phi_C - \frac{1}{rV_C} & \Phi_C \\ \Phi_L - \frac{1}{rV_L} & \Phi_L \end{pmatrix}, \quad \det = \frac{1}{r} \left[\frac{\Phi_C}{V_L} - \frac{\Phi_L}{V_C} \right] = \frac{1}{r} \cdot \frac{\Phi_C V_C - \Phi_L V_L}{V_C V_L}. \quad (73)$$

The determinant is nonzero iff $\Phi_C V_C \neq \Phi_L V_L$, and the value *level* cancels: substituting the kernel gives $\Phi_k V_k = b A_k f_k(m_k - \omega_k^*) \omega_k^* / (r S_k^A X_{\text{can},k})$, in which V_k has disappeared. *Separation therefore rests not on $V_C > V_L$ — the levels cancel identically — but on the crop and grazing survivor-band composition kernels differing: $m_C \neq m_L$, $\omega_C^* \neq \omega_L^*$, $f_C \neq f_L$, $S_C^A \neq S_L^A$, and $X_{\text{can},C} = R^O - \kappa + c_C \neq R^O - \kappa + c_L = X_{\text{can},L}$.* Crops are selected on the high tail and grazing on the middle band, so these objects differ generically, $\Phi_C V_C \neq \Phi_L V_L$, and the crop-value plus pasture-value pair is full-rank in (π_c, π_R) exactly where the pooled value moment — the value-weighted sum of the two rows — is rank-1. This is the one internal moment that would break the fusion the design otherwise leaves standing.

The rank-adder is not executable on the available data. The Agricultural Census

¹⁹Dropping the value moments collapses the rank of $X' \hat{\Omega}^{-1} X$ on the $(\pi_c, \pi_R)^{\text{High}}$ block.

records one operator-estimated market value of all land and buildings operated, owned and rented together, as a single figure per farm; it does not report cropland value separately from pastureland value. Only the pooled $V^A = w_C V_C + w_L V_L$ is observed, and constructing the split requires imputing an activity value allocation — but any such imputation imposes the V_C/V_L relationship the determinant needs to be free, re-importing the very cost information it claims to identify, and is isomorphic to an external cost cap rather than a new identifying moment. The rank-adder is therefore mechanically valid but not executable on Census data; recovering the split, as opposed to merely establishing that rank exists, additionally requires the levels V_C, V_L (which enter the loadings even though they cancel in the determinant), so a response differential alone would not close it. The omission accordingly stays a bound.

The two farm-gate levels: one object, two estimators. The exit-cutoff Jacobian carries the livestock farm-gate return b_L as a *forced slot* — the derivative of the participation threshold at the model’s own calibrated point, $\partial \ln \omega^* / \partial \ln b_L = -1$ and $\partial \omega^* / \partial b_C = 0$, so the marginal exiting parcel is the livestock parcel and its indifference identity fixes the slot. Substituting the activity-aggregate farm-gate $\bar{b}$ into that slot would misprice the marginal parcel’s indifference and break the envelope cancellation of the welfare decomposition; the fit-side b_L is not discretionary there. The one discretionary object is the *level* estimator of the exit numerator that enters the development-value re-pricing, for which $\bar{b}$ — the ordering-consistent, weakly lower estimator — and b_L are two estimators of the *same* structural level. We report the $\{\bar{b}, b_L\}$ pair as a full-pipeline sensitivity band. The choice moves $\hat{\pi}$ only along the non-identified c/R direction and shifts the external development-value figure by roughly one to three percent; the GMM alone still does not pin a positive magnitude under either estimator, since the direct leg is cost-driven rather than bar-driven. The band is therefore a disclosed robustness register on the external anchor, not a source of movement in what the GMM identifies.

The weight matrix $\hat{\Omega}$ is a stable-county panel bootstrap. The bootstrap resamples stable counties with replacement, block-clustering on the original county identifier and preserving the within-county year panel, re-estimates the Gardner first stage and all pooled DID2s on each resample, and recovers the empirical cross-moment covariance from the 1,000-iteration bootstrap distribution. Re-estimating the first stage on every resample carries the first-stage imputation uncertainty into $\hat{\Omega}$ rather than conditioning on a single fitted first stage. The resulting $\hat{\Omega}^{-1}$ is dense across moments, so a contaminated moment propagates into every primitive estimate through the off-diagonal weights — the reason the common-sample crop-revenue fix moves the outside-option estimate even though crop revenue carries no direct loading on π_R . The bootstrap captures the within-stratum cross-outcome correlation generated by the algebraic identity $\ln V^A + \ln S^A = \ln(V^A S^A)$, which ties the value and acreage outcomes estimated on the same common sample; cross-stratum correlations are near zero, consistent with the disjoint stable-county samples that define the strata. As a robustness check we also report results under the simpler diagonal $\hat{\Omega}$ built from stratum-level Conley 100 km variances; the comparison appears in the robustness paragraph below.

Hansen J over-identification. The pooled two-activity system over-identifies the eight primitives. The Hansen J statistic

$$J = g(\hat{\pi})' \hat{\Omega}^{-1} g(\hat{\pi}) \sim \chi_{\text{df}}^2 \quad \text{under correct specification,} \quad (74)$$

where df is the over-identification count, tests whether the implied cross-stratum restrictions on the ATT vector are consistent with the model's homogeneity assumptions. No separate sample-size factor multiplies J : $\hat{\Omega}$ is the bootstrap covariance of the ATT vector itself (below), so the panel size is already absorbed into the weight metric. The benchmark specification is the pooled winsorized system; its Hansen J is 14.7 ($p = 0.0055$, df 4) on the rainfed-low surface, rejecting the over-identifying restrictions at the 5 % level, and 9.2 ($p = 0.056$, df 4) on the rainfed-high surface. The reported J is computed on an un-repaired $\hat{\Omega}$: its condition number is 663 on the rainfed-low surface and 1,060 on the rainfed-high surface, three orders of magnitude below the nearest-positive-definite repair gate, so the χ_4^2 approximation is not compromised by a covariance repair. Replacing that approximation with a bootstrap reference distribution — J recomputed at the fixed $\hat{\Omega}$ across the 1,000 resample draws, each recentered on the fitted model — does not soften the verdict: the bootstrap distribution is tighter than χ_4^2 , so the observed statistic sits further in the tail and the over-identification rejection is sharper, with bootstrap $p = 0.019$ on the rainfed-low surface and 0.013 on the rainfed-high surface. Table A15 reports the two references side by side. The high surface, borderline under χ_4^2 , rejects under the correct reference.

The misfit is carried by the high-pressure stratum. Leave-one-moment-out re-estimation, dropping each of the twelve moments in turn and refitting the closed-form system on the remaining eleven, localizes the rejection: dropping either the farmland-acreage or the livestock-revenue moment in the high-pressure stratum collapses J symmetrically — the two drops are numerically indistinguishable — to $p = 0.23$ on the rainfed-low surface and $p = 0.93$ on the rainfed-high surface, while dropping the low-pressure livestock-revenue moment rescues the rainfed-low surface to $p = 0.12$. Table A16 reports the two representative high-pressure drops. The rejection therefore impugns the point identification of the eight primitives, not the lower-bound reading the decomposition commits to. Across all twelve drops the reported omission is invariant: the GMM alone does not pin a positive magnitude on either surface, and the external development-value figure is unchanged (Table A16, final column), since neither reads off the rejected point estimates. The identification-robust confidence set for the omission ΔW^O includes zero on both surfaces (Table A17, below), consistent with the set-identified reading. The over-identification rejection is real and, under the correct reference, sharper; it is also orthogonal to the welfare object the decomposition reports.

Identification-robust confidence sets. The fringe c/R block is near-singular, so the fixed- $\hat{\Omega}$ percentile bootstrap on those primitives is not valid under weak identification. We therefore report Anderson–Rubin confidence sets, inverting the statistic $\text{AR}(\pi) = g(\pi)' \hat{\Omega}^{-1} g(\pi)$ over the fringe primitives against the χ_{12}^2 projection reference and mapping the accepted region through the welfare decomposition. The identified sum $(\pi_c + \pi_R)^{\text{High}}$ excludes

zero on both surfaces, and the identification-robust set for the omission ΔW^O includes zero on both surfaces; the latter is roughly 2.4 to 2.9 times as wide as the percentile band it replaces. Table A17 reports both objects. The set thus locates the identified quantity — the acreage-and-cost sum on the fringe — as tight and significantly positive, while leaving the level of the outside-use omission unidentified in the way the set-identified reading already records. The sharper χ^2_{m-k} reference set, which conditions on the over-identifying restrictions, is empty on the rainfed-low surface: with the fitted system rejected, no primitive value satisfies both the moments and the restrictions simultaneously, an over-identification signal that reproduces the Hansen rejection from the confidence-set side.

The A5 sign restriction and the reported bound. Assumption A5 restricts the high-pressure input-cost response to be non-negative, $\pi_c^{\text{High}} \geq 0$, which is the sign restriction that gives the fringe outside-option response a finite interval. We report the omission across three settings of that restriction: the A5 baseline, the unrestricted fitted point, and a large-negative $\pi_c^{\text{High}} = -0.13$ calibrated to the observed per-acre operating-expenditure fall. Table A18 reports the result. The external development-value figure is invariant across all three settings on both surfaces — it re-prices the converted stock at the developer bid and never touches π_c — and the GMM-identified direct leg stays at approximately \$0 at both the baseline and the unrestricted fitted point, since the fitted π_c^{High} is itself non-negative and A5 does not bind there. The direct leg turns negative only under the -0.13 setting, so the A5 restriction is meaningful for it in principle even though the width it opens at the reported bound is zero.

Two pieces of evidence reject the -0.13 setting as a movement in the cost primitive. The input-cost response π_c is a residual on the product of input prices and quantities, and the observed per-acre expenditure fall is a quantity contraction rather than a price fall: the one clean price measure — the fertilizer unit price on the pressured fringe — is a precise null (-0.009 , $p = 0.31$), and the external ERS estimate of the grazing-margin exit cost gives $\partial c_L / \partial \tau \approx 0$, the cost margin that A5 governs. A disinvestment contraction in input use does not sign the price primitive negative, so the setting that would move the direct leg is one the data do not support; the grid displays the robustness rather than resting on the assumption.

Three-piece decomposition of the captured component. The model-implied ΔV^A admits a three-piece decomposition: a direct revenue-and-cost margin on infra-marginal parcels, a composition margin driven by output-price and input-cost changes (π_b/π_c shifting the participation threshold), and a composition margin driven by outside-use returns (π_R pulling marginal parcels out of agriculture). Stratum-level and aggregate decomposition shares are reported in Table A19. The three pieces sum to 100 percent by construction; the two composition pieces are individually large and offsetting, so they cancel in net and the direct revenue-and-cost margin on infra-marginal parcels survives the cancellation. The all-treated aggregate shares are 96 % (direct), -136 % (composition π_b/π_c), and $+139$ % (composition π_R) on the rainfed-low surface. The indifference identity at the participation threshold forces this pattern: in equilibrium, threshold-crossing terms enter

Table A15: Bootstrap Hansen- J Reference Distribution

Surface	Hansen J (df)	p_{χ^2}	p_{boot}	$\text{cond}(\hat{\Omega})$	$\hat{\Omega}$ repaired?
Low-suitability	14.7 (4)	0.0055	0.019	663	No
High-suitability	9.2 (4)	0.056	0.013	1060	No

Notes. Hansen J over-identification statistic (four degrees of freedom) on each GAEZ surface, with the χ^2_4 p -value and the bootstrap- J p -value side by side. The bootstrap reference recomputes $J = g'\hat{\Omega}^{-1}g$ at the fixed point $\hat{\Omega}$ across the 1,000 resample draws, each recentered on the fitted model to form the null over-identification distribution. $\text{cond}(\hat{\Omega})$ is the condition number of the bootstrap covariance; $\hat{\Omega}$ is not nearest-positive-definite-repaired on either surface, so the reported J is computed on the un-repaired covariance.

Table A16: Leave-One-Moment-Out: Invariance of the Omission Ladder

	Moments	Hansen J (df)	p	Pooled ΔW^O (\$B)	R^2 ceiling (\$B)
<i>Low-suitability surface</i>					
Full pooled system	12	14.7 (4)	0.0055	47.7	200.7
Drop ln farmland $\times$ High	11	4.3 (3)	0.228	141.0	200.7
<i>High-suitability surface</i>					
Full pooled system	12	9.2 (4)	0.056	92.6	195.0
Drop ln farmland $\times$ High	11	0.4 (3)	0.931	110.5	195.0

Notes. Each of the twelve moments is dropped in turn and the closed-form system re-estimated on the remaining eleven. The rows report the full pooled system and the representative high-pressure drop (dropping the farmland-acreage or the livestock-revenue moment in the high-pressure stratum is numerically indistinguishable and collapses J symmetrically). The GMM alone does not pin a positive omission magnitude across any drop; the external development-value figure and the set-identification are invariant to which moment is dropped (final column). Dollar magnitudes in real 2025-USD billions.

Table A17: Identification-Robust Confidence Sets

Surface	Sum $(\pi_c + \pi_R)^{\text{High}}$		Omission ΔW^O (\$B)	
	Point	AR 95% set	Point	AR 95% set
Low-suitability	0.0335	[0.0050, 0.0620]	47.7	[−77.9, 149.6]
High-suitability	0.0585	[0.0152, 0.1018]	92.6	[−7.2, 192.5]
Percentile-bootstrap ΔW^O band: low [7.7, 103.5], high [62.6, 132.1].				

Notes. Anderson–Rubin confidence sets, inverting $\text{AR}(\pi) = g(\pi)'\hat{\Omega}^{-1}g(\pi)$ over the fringe primitives against the χ^2_{12} projection reference and mapping the accepted region through the welfare decomposition at the point-calibrated shock size. The identified sum $(\pi_c + \pi_R)^{\text{High}}$ excludes zero on both surfaces; the omission ΔW^O set includes zero on both. The final row reports the fixed- $\hat{\Omega}$ percentile-bootstrap ΔW^O band for comparison; the AR set is roughly 2.4 to 2.9 times as wide. On the rainfed-low surface the sharper χ^2_{m-k} reference set, which conditions on the over-identifying restrictions, is empty. Dollar magnitudes in real 2025-USD billions.

Table A18: A5 Sign Grid on the High-Pressure Cost Response

Surface	A5 setting on π_c^{High}	R0 (\$B)	R1 (\$B)	R2 ceiling (\$B)
Low	A5 baseline ($\pi_c^{\text{High}} \geq 0$)	0.0	0.0	200.7
	Unrestricted (fitted point)	0.0	0.0	200.7
	$\pi_c^{\text{High}} = -0.13$ (expenditure-fall setting)	0.0	-51.0	200.7
High	A5 baseline ($\pi_c^{\text{High}} \geq 0$)	0.0	0.0	195.0
	Unrestricted (fitted point)	0.0	0.0	195.0
	$\pi_c^{\text{High}} = -0.13$ (expenditure-fall setting)	0.0	-17.4	195.0

Notes. The omission under three settings of the Assumption A5 restriction on the high-pressure input-cost response π_c^{High} : the A5 baseline ($\pi_c^{\text{High}} \geq 0$), the unrestricted fitted point, and a large-negative setting $\pi_c^{\text{High}} = -0.13$ calibrated to the observed per-acre operating-expenditure fall. The external development-value figure is invariant across all three settings on both surfaces; the GMM-identified direct leg turns negative only under the -0.13 setting, which the fertilizer-price null and the ERS grazing-cost anchor reject as a price movement. Dollar magnitudes in real 2025-USD billions.

agricultural value with offsetting signs as composition shifts both into and out of agriculture in response to the same access change.²⁰

Table A19: Three-Piece Decomposition of ΔV^A : By Cell and Aggregate

Cell	Direct (\$B)	π_b/π_c (\$B)	π_R (\$B)	Net ΔV^A (\$B)	Direct (%)	π_b/π_c (%)	π_R (%)
Marginal (N)	5.4	-7.8	7.7	5.3	102	-148	146
S3	5.7	-7.3	7.1	5.6	103	-130	127
Suitable (S2+S1)	6.0	-5.5	5.3	5.9	103	-93	90
All IHS-treated	17.2	-20.5	20.1	16.7	103	-123	120

Notes. Three-piece decomposition of the captured component ΔV^A at the production calibration. Dollar magnitudes in real 2025-USD billions, capitalized at $r = 0.05$. Percentage shares sum to 100% by construction; individual components can exceed 100% in absolute value when offsetting (the indifference identity at the participation threshold forces threshold-crossing terms to enter V^A with offsetting signs, so the two composition pieces cancel in net and the direct margin survives the cancellation). Each row aggregates the component over its pressure stratum; the All IHS-treated row aggregates all three strata. Dollar and percentage entries are reported in the table body (rainfed-low surface; winsorized bootstrap).

²⁰The composition flux that cancels here — the Leibniz boundary term of equations (65)–(64), evaluated at the marginal rent ρ_{ct}^σ — is a distinct object from the re-pricing of the inframarginal exited *stock* that sets the width of the omission bound. The cancellation applies to the marginal parcel crossing the threshold; the bound width comes from valuing the mass strictly interior to the exit region above the marginal rent when Assumption A4 is relaxed, an $\eta\Delta^O$ level term on the held slice. The external development-value figure of Section 7.3 is that inframarginal-stock re-pricing carried up to the external developer bid R^O , not a reintroduction of the flux that cancels in the aggregate.

Robustness across alternative specifications. The headline bound is robust to relaxation of the homogeneity restriction on the output-price elasticity π_b . We report results from two specifications, both estimated under the same baseline calibration (bootstrap $\hat{\Omega}$, the data-anchored $\hat{\theta}$, V^A stock baseline).

(i) *Baseline: the pooled two-activity specification, π_b common across strata.* The setup defined above. On the rainfed-low surface the Hansen $J = 14.7$ ($p = 0.0055$), the GMM alone does not pin a positive omission magnitude, and the identification-robust confidence set for ΔW^O (Table A17) spans zero. The omission concentrates in the high-pressure fringe rather than the thick no-pressure stratum: the no-pressure stratum contributes zero at the reported bound (its reversion is held at $R^N = 0$), while the fringe carries the omission through the sign-restricted direct leg. The high-pressure stratum is also where the over-identification residual concentrates, so the residual misfit and the omission share the same fringe; the leave-one-moment-out battery (Table A16) shows that the external development-value figure and the set-identification are invariant to which moment is dropped, so the bound survives the Hansen rejection because the fringe contribution is set-identified by sign restrictions rather than by the rejected point estimates. The Hansen J reported here and throughout is computed on an un-repaired $\hat{\Omega}$.

(ii) *Development-pressure-stratified output-price elasticity.* A relaxation that allows π_b to vary across development-pressure status, motivated by Prediction 4’s observation that the value–revenue wedge is pressure-first near urban markets. The relaxed specification leaves the all-treated omitted share essentially unchanged; the relaxation does not alter the headline bound.

The welfare bound is therefore identified by a linear combination of channel primitives that the data pin down robustly; the omitted share is not sensitive to the particular way the output-price heterogeneity is parameterized. The baseline pooled specification is reported as the paper headline; the development-pressure-stratified specification is a confirming robustness check.

θ -sensitivity of the welfare bound. The welfare bound depends on the input $\hat{\theta}$ through the Fréchet quantities (ω_k^*, m_k, f_k) that enter the stratum-level baseline calibration and the design matrix. The bound’s sensitivity across the θ range and the two closest literature comparators (Sotelo (2020) Peruvian district at 1.66; Costinot et al. (2016) developed-country subsample at 2.29) is reported in Section 7.3, Table 5. The paper’s headline brackets this sensitivity through the dual- θ band spanning the rainfed-low and rainfed-high surfaces: the captured ΔV^A is \$54.3B on the rainfed-low surface and \$65.9B on the rainfed-high surface, while the omission the aggregate misses enters only through the external development-value anchor (Section 7.3), since the GMM alone does not separate the cost and outside-option responses on the fringe.

Robustness to crop-level θ_k heterogeneity. Appendix B.15 documents that crop-specific θ_k values within the SUI-9 bundle vary materially across the eight non-numeraire crops; the pooled $\hat{\theta} = 2.13$ is a variance-weighted average rather than a structural common parameter. We test whether the welfare bound is sensitive to this heterogeneity by replacing the pooled $\hat{\theta}$ with cell-specific effective values $\theta_{\text{eff}}(c)$ computed as acreage-weighted

means of the per-crop θ_k within each pressure stratum, excluding crops with $\theta_k < 0.3$. Running the pooled GMM at the stratum-specific θ_{eff} vector leaves the all-treated omission share essentially unchanged relative to the pooled- θ headline, and a Monte Carlo over the per-crop θ_k sampling distribution gives a structural-uncertainty width an order of magnitude below the ATT-sampling bootstrap CW that already enters the headline. The single- θ welfare bound is robust to crop-level θ_k heterogeneity at the order of magnitude that matters for the paper’s headline claims.²¹

Predetermined stratum assignment. The pressure partition uses a predetermined county attribute. The city-reach moderator M_c is built from 1940-and-earlier geography and absorbs no information from the highway treatment itself, so stratum assignment is orthogonal to IHS-induced reallocation by construction. Market potential and the other market-access controls enter the Gardner first stage, so the stratum-level DID2s identifies within-stratum residual variation orthogonal to the continental market-access gradient the first stage already projects out. Suitability enters the first stage as a year-interacted trend control rather than a second-stage interaction, so the pooled strata preserve the outside-use margin that Assumption A2 operates on while delivering enough within-stratum counties for the ATT-level Conley variance to stabilize. Section 7.3 reports sensitivity to the development-pressure threshold.

The data reject development-pressure primitive pooling. The baseline pooled specification estimates the eight primitives $(\pi_{b_C}, \pi_{b_L}, \pi_c^{\text{None}}, \pi_c^{\text{Low}}, \pi_c^{\text{High}}, \pi_R^{\text{None}}, \pi_R^{\text{Low}}, \pi_R^{\text{High}})$ against the stratum-level moments. Imposing $\pi_c^{\text{None}} = \pi_c^{\text{Low}} = \pi_c^{\text{High}}$ and $\pi_R^{\text{None}} = \pi_R^{\text{Low}} = \pi_R^{\text{High}}$ collapses the outside-option and cost responses across the development-pressure strata. The difference in the Hansen J tests the pooling restrictions directly and rejects them. The identified exit-cutoff response $(\pi_c + \pi_R)^\sigma$ is larger under development pressure — consistent with the reduced-form acreage and value gradients, and not requiring the within-stratum split of π_c from π_R — so the strata cannot be pooled. Conditional on the Fréchet structure, development-pressure primitive heterogeneity is not a refinement the data can dispense with, and the headline bound is robust to a restriction the data reject.

Aggregate moments cannot recover the wedge. A second alternative replaces the stratum-level moments with three aggregate moments from a pooled Gardner DID2s with no heterogeneity interactions (Gardner et al., 2024). The three outcomes satisfy the identity $\ln(V^A S^A) = \ln V^A + \ln S^A$, so the $\ln(V^A S^A)$ design row is the sum of the $\ln V^A$ and $\ln S^A$ rows, $a^T = a^V + a^F$, and the aggregate design matrix has rank two: π_c and π_R are not separately identified from aggregate moments alone. Imposing $\pi_c = \pi_R$ to restore identification produces an all-treated omission share far below the stratified estimate. The aggregate construction discards the stratum-level baselines $(A_k, S_k^A, \omega_k^*, V_k^A)$ through which the stratified system separates the composition and outside-option channels: aggregate moments cannot recover the welfare wedge.

²¹The Monte Carlo propagates only the per-crop θ_k sampling layer rather than jointly with the ATT bootstrap; the two layers are nearly orthogonal because per-crop θ_k are estimated from the cross-crop ratio regression (Appendix B.15), while stratum-level Gardner DID2s ATTs identify against treatment timing.

The two pooled alternatives answer distinct questions. The pooled-primitive system tests whether development-pressure variation in the composition primitives is informative; the difference-in- J test rejects pooling, as the estimated outside-option response is larger under pressure ($\pi_R^{\text{press}} > \pi_R^{\text{none}}$). The aggregate-pooled system asks whether the welfare wedge survives aggregation of the moments; its rank-two design shows it does not. Both results are conditional on the maintained Fréchet structure and the stratum-level pooling.

Appendix B.12 Three Farm Decisions

The framework accommodates three margins of adjustment that appear in the Agricultural Census. Throughout this subsection write $R^A(\phi_{jc}, \tau_{ct}) \equiv \max\{R^C(\phi_{jc}, \tau_{ct}), R^L(\phi_{jc}, \tau_{ct})\}$ for the parcel's best agricultural return.

Exit to outside use. This is the primary margin: parcels with $\omega_{jc} < \omega_{ct}^*$ leave agriculture. The number of farms falls, land in farms contracts, and remaining agricultural activity concentrates on higher-productivity parcels. The exit margin is governed by equation (25).

Tenancy. Land can remain in agriculture even when the current owner-operator exits, if a separate tenant can profitably farm it. The tenancy margin involves two conditions:

- *Operator's condition:* a tenant operates parcel j if $R^A(\phi_{jc}, \tau_{ct}) \geq \rho_{ct}^{\text{lease}}$, where ρ_{ct}^{lease} is the lease rate (distinct from the outside-option ρ_{ct}^σ used elsewhere).
- *Owner's condition:* the owner leases rather than self-operates when the burden of self-operation ζ_{ct} exceeds the leasing friction χ_{ct}^L :

$$\Pi_{ct}^{\text{lease}} = s R^A(\phi_{jc}, \tau_{ct}) - s \chi_{ct}^L \geq \Pi_{ct}^{\text{own}} = s R^A(\phi_{jc}, \tau_{ct}) - s \zeta_{ct}, \quad (75)$$

which simplifies to $\zeta_{ct} > \chi_{ct}^L$.

An interstate opening raises ζ_{ct} (nonfarm employment raises the opportunity cost of farm management) and may lower χ_{ct}^L (thicker local markets reduce search costs). Both changes push toward greater tenancy.

Consolidation. Surviving farms with $\omega_{jc} \geq \omega_{ct}^*$ can expand by absorbing neighboring parcels that exit owner-operation. Whether exiting parcels are absorbed or converted depends on κ_{ct} : high κ keeps exiting land in the farming sector under tenant operators, generating consolidation; low κ triggers conversion to outside use.

Appendix B.13 Mapping Model Objects to Empirical Outcomes

The mapping to the Census carries three measurement caveats. First, revenue per acre in the data corresponds to gross receipts per acre using the received price rather than the

Table A20: Mapping Model Objects to Agricultural Census Outcomes

Model Object	Equation	Census Counterpart
S_{ct}^A	(34)	Agricultural land share; number of farms
N_{ct}^A	(37)	Land in farms; cropland harvested
$\bar{V}_{ct}$	(44)	Farmland value per acre
W_{ct}	(60)	Aggregate county land value
W_{ct}^A	(61)	Agricultural component of land value
W_{ct}^O	(62)	Outside-use component of land value
$\text{Revenue}_{ct}/N_{ct}^A$	(40)	Crop revenue per acre
$\bar{\phi}_{ct}^S$	(39)	Average productivity (endogenous)
Tenancy share	(75)	Share of farms/acreage tenant-operated
A_c	(18)	GAEZ land suitability index
ω_{ct}^*	(27)	Not directly observed; identified via S_{ct}^A
R_{ct}^O	(17)	Development bid / land-rent bar; not directly observed, proxied by fringe built-up structures (BUPL)
M_c	§4.1	County share of convertible 1940 land in the pre-interstate city-reach fringe (development-pressure moderator)
ϕ_{CL}	(26)	Crop–livestock composition cutoff; crop share of farmland (cropland / farmland)
$\Delta \ln \bar{V} - \Delta \ln(\text{Rev}/N^A)$	(57)	Value–revenue capitalization wedge ($\Delta \ln$ farmland value/acre less $\Delta \ln$ revenue/acre)
$\omega_{\text{press}}^*, \omega_{\text{none}}^*$	(8)	Pressure-stratified exit cutoffs; not directly observed, read off the M_c -stratified S_{ct}^A response

net return $b(\tau)$ (which subtracts delivery costs). Second, farmland value in the data is an average over parcels classified as farmland, as represented in equation (44). Third, W_{ct}^A and W_{ct}^O are not separately observed in any single data source; the welfare decomposition in Appendix B.10 is applied to the paper’s ATT estimates in the conclusion using the suitability-stratified results to bound each component.

Appendix B.14 GAEZ Suitability as a Proxy for A_c

The empirical analysis measures A_c using FAO GAEZ crop suitability indices, which combine soil quality, climate, and terrain into a single index of agronomic potential under rainfed low-input conditions. For GAEZ to be a valid proxy for A_c , the GAEZ index must be monotonically related to A_c up to classical measurement error uncorrelated with the interstate treatment variable and with the within-county distribution $F(\omega)$.

We defend this assumption on two grounds. First, GAEZ indices measure agronomic fundamentals (soil texture, drainage, organic matter, climate suitability for major crops) that are time-invariant and not affected by the highway network. Second, although GAEZ is calibrated to low-input farming technology, the ordinal ranking of counties by soil and climate suitability is robust to technology: the relevant condition is rank invariance, not level equivalence.

Under the Fréchet specification, A_c enters as the scale parameter of the composite productivity distribution $\phi_{jc} \sim \text{Fréchet}(A_c^\theta, \theta)$. A higher GAEZ index stochastically shifts the distribution of ϕ_{jc} upward, which is the condition required for the comparative statics in Appendix B.4 to hold. The four FAO suitability classes used in the empirical analysis — N, S3, S2, and S1, collapsed to the three identification classes $N/S3/S2S1$ where the thin top class is merged — approximate this continuous ordering.

As a robustness check, we verify that results are qualitatively unchanged when the suitability classes are constructed using pre-treatment crop revenue per acre from the 1940 or 1950 Agricultural Census rather than GAEZ indices.

Appendix B.15 Direct US Estimation of the Tail Parameter

Section 7.2 reports the tail-parameter estimate $\hat{\theta} = 2.13$ (asymptotic two-way 95% CI [1.79, 2.61]) recovered from US Census of Agriculture data. This subsection documents the estimating equation, the bundle construction, the sample and specification, the inference ladder, the per-crop heterogeneity diagnostic, and the literature comparators.

Identification. The tail parameter θ is identified from the parcel-productivity Fréchet shape alone and is untouched by the A1/A2 reframe: it depends on the dispersion of suitability-driven yields, not on the slopes of the farm-gate return b or the outside option R^O . Under the Sotelo (2020)-style single- θ Fréchet land-allocation model with constant input-cost shares at the parcel level (the Sotelo constant-cost-share condition, distinct from the competing-channels Assumption A1 of Section 3), the conditional land share

of crop k in county c at year t satisfies

$$\pi_{k,ct} = \frac{(\lambda_{ct} A_{c,k})^\theta}{\sum_j (\lambda_{ct} A_{c,j})^\theta},$$

where $A_{c,k}$ is the county–crop GAEZ rain-fed-low suitability and λ_{ct} absorbs county-time price levels, conversion costs, and other county–time scalars. Taking the log-ratio with respect to a chosen numeraire crop k_0 (corn) eliminates λ_{ct} and delivers [Sotelo \(2020\)](#)’s equation (4):

$$\ln\left(\frac{\text{acres}_{c,k,t}}{\text{acres}_{c,k_0,t}}\right) = \theta \cdot \ln\left(\frac{A_{c,k}}{A_{c,k_0}}\right) + \alpha_c + \delta_{k,t} + u_{c,k,t}, \quad (76)$$

where α_c is a county fixed effect (absorbing λ_{ct} at the county-uniform level) and $\delta_{k,t}$ is a crop $\times$ year fixed effect (absorbing the national relative-price and weather components of $\ln \lambda_{ct} - \ln \lambda_{ct}^{k_0}$). The slope on relative log suitability is θ directly – verified by independent algebraic audit against the [Sotelo \(2020\)](#) derivation; not $1/\theta$.

Bundle (SUIT-9). The bundle aggregates the nine Ricardian crops in the paper’s FAO–Rain_Mean suitability index: corn, wheat, soybean, sorghum, barley, oats, cotton, white potato, wetland rice. Alfalfa, the tenth crop in the index, is excluded from the M1 estimation bundle on three jointly sufficient grounds. (i) The GAEZ alfalfa surface is empirically collinear with the grain bundle (correlation with wheat = 0.885; mean correlation with the other nine SUIT-10 crops = 0.718; PCA loading on PC1 = 0.341, indistinguishable from grains). (ii) The Agricultural Census `all_hay` aggregate covers eight hay species against an alfalfa-only GAEZ surface, an LHS/RHS aggregation mismatch. (iii) Alfalfa locates primarily by livestock proximity rather than by Ricardian land-suitability comparative advantage. The cost of the SUIT-9 versus SUIT-10 stratification choice is empirically trivial: 99.1% of counties receive identical Q1–3/Q4–5 binary assignment (26 of 3,034 counties flip at the Q3/Q4 boundary).

Sample and specification. The empirical sample stacks county-crop-year cells for the eight non-numeraire SUIT-9 crops across the 1940–1974 panel of the Agricultural Census, restricted to the 48 contiguous states and to county-crop-year cells with positive harvested acres. The pooled estimating sample is $N = 115,887$ observations over 2,923 counties and eight Census stops (1940, 1945, 1950, 1954, 1959, 1964, 1969, 1974). Right-hand-side suitability is the area-weighted-with-zeros GAEZ rain-fed-low surface (the convention that includes county-crop cells with zero positive GAEZ pixels at zero suitability, rather than dropping them). No balance requirement is imposed across years; the specification admits any county-crop-year cell that satisfies the harvested-acres condition. Standard errors are asymptotic two-way (county, year) cluster-robust under the Cameron-Gelbach-Miller (2011) variance estimator ([Cameron et al., 2011](#)); we adopt this as primary under the panel architecture and report the full inference ladder in [Table A22](#).

Window choice. The production window is 1940–1974. The panel-architecture window grid produces $\hat{\theta}$ trajectories that are flat across three consecutive Census stops at

Table A22: Inference Ladder for the Production Panel Estimate $\hat{\theta} = 2.13$

Inference	SE	95% CI
Two-way cluster (county, year)	0.164	[1.81, 2.45]
County-only cluster	0.155	[1.83, 2.43]
Conley 50km	0.176	[1.78, 2.47]
State cluster ($G = 48$)	0.429	[1.29, 2.97]

Notes. Four-rung inference ladder for the production estimate $\hat{\theta} = 2.13$. The asymptotic two-way (county, year) cluster (Cameron et al., 2011) is the conventional inference under the panel architecture; we adopt it as primary despite the small year-cluster count ($G_{\text{year}} = 8$) because the small- G correction in the cluster-robust variance estimator inflates the standard error conservatively in this regime. The county-only and Conley 50 km rows are reported as transparency about which residual dimension drives the standard error; both ignore at least one correlation direction and so understate uncertainty relative to two-way. The state-cluster row ($G = 48$) imposes within-state correlation that the county fixed effect has already absorbed at the state-uniform level and is reported as a conservative upper bound on standard errors. Table generated by `welfare_analysis/R/15-paper-exhibits.R` from `output/01_theta/clustering_ladder.csv`.

$\theta \in [2.18, 2.23]$ in 1969, 1974, and 1978, decline mildly to $\hat{\theta} = 2.12$ at 1982, and drop materially to $\hat{\theta} = 1.91$ by 2002. The 1940–1974 production window sits inside the empirical plateau. Multiple candidate post-1980 contaminations would explain the decline outside the plateau: the IHS-induced production-reallocation response (visible only after ~ 1980 in the panel), accumulation of Class 2 time-varying (c, k, t) confounders such as differential technology adoption and federal crop-policy shocks, or genuine structural drift in the Fréchet shape. The panel-architecture window grid alone cannot separate which contamination is operating; the conservative choice is to stop the window inside the plateau before any candidate bites materially.

Per-crop heterogeneity. A saturated regression interacting the slope on $\ln(A_{c,k}/A_{c,\text{corn}})$ with crop fixed effects estimates eight crop-specific θ_k values for the SUIT-9 non-numeraire crops. The empirical θ_k values range from -0.06 (wetland rice; an irrigation-bound crop where the GAEZ rain-fed surface does not track US production geography) to 3.29 (white potato). The pooled $\hat{\theta} = 2.13$ is a variance-weighted average of these crop-specific values rather than a structural common parameter; the same maintained pooling assumption underlies both Sotelo (2020) (20 crops, single θ) and Costinot et al. (2016) (10 crops, single v). Per-crop dispersion is the natural axis for a multi-group extension and is the subject of a separate investigation; the single- θ welfare bound reported here is robust to crop-level θ_k heterogeneity at the order of magnitude relevant for the paper’s headline claims (see the robustness paragraph in Appendix B.11).

Literature comparators. The most relevant published estimates in the structural Ricardian agriculture literature bracket our yield-based dual- θ band [1.61, 2.13] on both sides. [Sotelo \(2020, p. 2715\)](#) reports $\hat{\theta} = 1.658$ at the district level using data for 4 of Peru’s 24 departments (his finest spatial unit, on the subsample with the strongest GAEZ match); our rainfed-high estimate of 1.61 sits just below the Sotelo Peruvian-district headline. [Costinot et al. \(2016, Tables 2, 5\)](#) report a baseline $\hat{\theta}^{\text{CDS}} = 2.46$ on cross-country crop output, with their 29-country developed-country subsample (including the US) yielding $\hat{\theta}_{\text{Developed}}^{\text{CDS}} = 2.29$ and their 21-country developing-country subsample yielding $\hat{\theta}_{\text{Developing}}^{\text{CDS}} = 2.71$. Our rainfed-low estimate of 2.13 sits 0.16 below the CDS developed-country subsample, so the band falls comfortably between the two comparators most relevant to the US application.

Maintained identifying assumption. The estimating equation absorbs $\log(P_{k,t}/P_{k_0,t})$ via the crop $\times$ year fixed effect $\delta_{k,t}$ under the county fixed effect α_c . To make explicit what survives, decompose the county-crop-year farm-gate price as $\log p_{c,k,t} = \mu_{k,t} + \lambda_c + \nu_{c,k} + \xi_{c,t} + \pi_{c,k,t}$ – the sum of a national crop-year level, a county-uniform level, a county-crop time-invariant intercept, a county-year level, and a county-crop-year residual. Taking the relative price with respect to the numeraire k_0 drops the county-uniform terms (λ_c and $\xi_{c,t}$ hit both crops equally) and leaves three components: the national crop-year ratio $\mu_{k,t} - \mu_{k_0,t}$, absorbed by $\delta_{k,t}$; a county-crop time-invariant intercept $\nu_{c,k} - \nu_{c,k_0}$; and a county-crop-year residual $\pi_{c,k,t} - \pi_{c,k_0,t}$. Neither of the last two is absorbed by the fixed-effect structure.

Assumption θ -1 (within-county orthogonality of relative crop prices to relative GAEZ suitability). The residual relative price $(\nu_{c,k} - \nu_{c,k_0}) + (\pi_{c,k,t} - \pi_{c,k_0,t})$ is mean-independent of the relative GAEZ suitability $\ln(A_{c,k}/A_{c,k_0})$ conditional on county and crop $\times$ year fixed effects. This is the maintained assumption common to [Sotelo \(2020\)](#) and [Costinot et al. \(2016\)](#); we treat it as defensible on the SUIT-9 bundle because pre-1940 transport-infrastructure differences across grains had compressed substantially by 1940 (the wheat-elevator advantage stabilized by ~ 1920), and within-bundle time-varying forces that could decouple acreage from suitability (irrigation, fertilizer adoption, mechanization, federal crop policy) are bounded within the production window at approximately 2 percentage points by the plateau stability $\hat{\theta} \in [2.18, 2.23]$ across 1969–1978.

Appendix B.16 Within-Agriculture Comparative Advantage

The orderings, coexistence, and the crop-share gradient. The slope ordering is agonomic – crop yields respond sharply to soil while extensive grazing is near quality-flat – and the cost ordering is what lets livestock exist: were crops both steeper and cheaper they would dominate every parcel and none would graze, so the coexistence the data show requires crops to carry the higher cost. The two orderings place the single crossing at ϕ_{CL} , crops above and livestock in the band below, with grazing’s low cost holding its exit threshold low enough to keep marginal land in farming where crops cannot pay – marginal-land retention is livestock holding the low-cost extensive margin. The livestock technology here is extensive grazing (range beef); intensive peri-urban animal agriculture

(dairy belts, market gardening), which improved access historically favored, is a distinct access-sensitive activity outside the single flat-return technology. A3's content is testable through the allocation it implies: the closed form makes the crop share of farmland monotone in suitability if and only if $b_C > b_L$ (Appendix B.16). The crop share 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, not a pressure or feasibility artifact.²²

The held, not-yet-converted band on middle land. One layer finer, the acreage response and the grazing-revenue response part company on middle land under pressure, and the gap is itself a test. A parcel priced out of grazing near the fringe need not convert in the same census wave: with conversion irreversible and the developer bid R^O rising in access, the optimal-stopping threshold sits above the myopic farm-exit breakeven, opening a held band of farm-marginal, not-yet-converted parcels that shed intensive crop revenue but keep their acres counted as farmland or idle grass (Capozza and Helsley, 1989, 1990). Such a parcel sheds grazing *revenue* faster than *acres* – the middle-suitability, high-pressure cells show acreage falling modestly while livestock revenue per acre collapses – so the acreage-versus-revenue wedge is the observable signature of this modeled held state, a Prediction 2 quantity read one layer finer rather than a separate prediction. The fertilizer rise is not part of the signature: it is the land-input substitution on the surviving crop parcels above the cutoff, a distinct set that clears the bar. The held slice is valued between 0 and $R^O - \kappa$, the option weight not point-identified by a five-year census snapshot, so it enters value only within the disclosed bound the welfare accounting carries (Section 7). The held state is distinguished from marginal-land retention by a same-parcel test – middle-land parcels going crop-to-idle-grass should later convert to developed use, whereas retained marginal land stays in grazing indefinitely (Appendix B.16).

This section derives the two-technology sorting on a single land-quality index that the main text uses: crops (C) and livestock (L) read the same quality $\phi = A_c \omega$ but reward it at different slopes, and the sort against two non-farm outside options (development R^O in the fringe, reversion R^N beyond it) delivers the survival, composition, and capitalization results. The cross-county comparative-advantage ordering follows from the crop-livestock cutoff; the composition shift toward livestock on the retained marginal base is a selection consequence of which land is held in farming, not a crop-versus-livestock revenue-elasticity differential. The mechanism is related to Pellegrina and Sotelo (2025), who document specialization toward comparative-advantage crops following Brazil's agricultural frontier expansion: infrastructure access reshapes within-agriculture comparative advantage across activities.

²²Crop share of farmland by canonical suitability class ($N \rightarrow S3 \rightarrow S2S1$): $0.183 \rightarrow 0.381 \rightarrow 0.605$; the suitability-level slope is strictly positive and highly significant (HC1, $p < 0.01$), and statistically indistinguishable when restricted to no-pressure counties. This is the allocation analog of the land-share identification of θ (Sotelo, 2020, and Section 7.2); a share above a cutoff is a headcount, not a return on selected land, so the gradient is free of the activity-selection that contaminates any comparison of per-acre returns.

Setup. Parcel quality is $\phi_{jc} = A_c \omega_{jc}$ with $\omega \sim \text{Fréchet}(\theta)$, $\Pr(\omega \leq w) = e^{-w^{-\theta}}$. Per-acre returns are $R^C = b_C \phi - c_C$ and $R^L = b_L \phi - c_L$, with the develop option $R_{ct}^O - \kappa$ (fringe) and the non-development reversion $R_{ct}^N - \kappa_N \geq 0$ (everywhere), and $b_C > b_L > 0$, $c_C > c_L > 0$ (Assumption A3).²³ A parcel takes $V_{jc} = \max\{R^C, R^L, (R^O - \kappa)1\{\text{fringe}\}, R^N - \kappa_N, 0\}$, the amenity floor 0 the residual; write ρ for the best non-farm option ($R^O - \kappa$ in the fringe, $R^N - \kappa_N$ beyond).

Two cutoffs, coexistence, and squeeze-out. Crops beat livestock iff $(b_C - b_L)\phi \geq c_C - c_L$, i.e. $\phi \geq \phi_{CL} = (c_C - c_L)/(b_C - b_L)$ (equation (26)). Livestock, the marginal farm use, beats the best non-farm option iff $b_L \phi - c_L \geq \rho$, i.e. $\phi \geq \phi^* = (\rho + c_L)/b_L$. Because crops are steeper, $\phi \geq \phi_{CL} > \phi^*$ implies crops also beat the outside option, so the ordering is clean: **crops** on $\phi \geq \phi_{CL}$, **livestock** on $\phi^* \leq \phi < \phi_{CL}$, **exit** below ϕ^* (to development in the fringe, to reversion R^N beyond it).

Coexistence requires $\phi^* < \phi_{CL}$; beyond the fringe ($\rho = R^N - \kappa_N$) this is $(R^N - \kappa_N + c_L)/b_L < (c_C - c_L)/(b_C - b_L)$, which holds when livestock is the low-cost extensive use and the reversion value R^N is modest. Two development-pressure regimes follow, divided by the pivot $\bar{R} \equiv b_L \phi_{CL} - c_L$, the level of the outside option at which the farm-exit cutoff ϕ^* reaches ϕ_{CL} (by the definition of ϕ_{CL} , $b_C \phi_{CL} - c_C \equiv b_L \phi_{CL} - c_L = \bar{R}$, so both farm returns equal the outside option at ϕ_{CL} at once). *Regime I* ($R^O - \kappa < \bar{R}$, low-to-moderate pressure): the crop-exit cutoff $\phi_C^* = (\rho + c_C)/b_C$ lies strictly between ϕ^* and ϕ_{CL} — the ordering $\phi^* < \phi_C^* < \phi_{CL}$ holds because $\partial \phi_C^* / \partial \rho = 1/b_C < 1/b_L = \partial \phi^* / \partial \rho$ — so every crop parcel clears its own exit cutoff, crops never bind, and the livestock band is the only shed farm margin. This nests the single-margin sort exactly, with ϕ_C^* slack inside the retained farm range. *Regime II* ($R^O - \kappa \geq \bar{R}$, high pressure): ϕ^* has risen past ϕ_{CL} , the pure grazing band is gone, and the binding farm-exit cutoff on the crop range is the *co-equal crop-exit cutoff*

$$\phi_C^* = \frac{R^O - \kappa + c_C}{b_C} \geq \phi_{CL}, \quad (77)$$

which climbs through the crop band from below and develops the lowest-value crop parcels on $[\phi_{CL}, \phi_C^*]$: displacement is no longer a grazing-only phenomenon. The crop-exit cutoff is a modeled co-equal margin, not a corner case; it is the object promoted to Regime II in the conceptual model, and it drives the middle-and-good-land displacement there. The displacement *mechanism* is native to the two-margin sort; *where* it lands — disproportionately on good and middle soil — is set by the maintained correlation $\text{Corr}(A_c, M_c) > 0$ (Assumption A2), which places high-pressure counties on good soil.²⁴

²³The single draw ω_{jc} is common across activities, consistent with ω_{jc} representing overall parcel quality (drainage, slope, microclimate) rather than crop-specific endowments. This differs from Costinot et al. (2016), who assign independent Fréchet draws per crop per parcel; independent draws would generate a gravity-style within-county composition formula, while the single draw yields cleaner threshold comparative statics. The main qualitative predictions hold under either specification.

²⁴Under unbounded Fréchet the very highest- ϕ crops are never out-bid; with bounded real crop returns a high-enough R^O takes even prime land, so the model under-captures prime-land conversion at the largest metros and is scoped to the marginal-to-middle of the suitability distribution. The co-equal cutoff is the ϕ_C^* of equation (77) above.

Closed-form county shares. Write the survival function of ϕ as $G(x) \equiv \Pr(\phi \geq x) = 1 - \exp(-(A_c/x)^\theta)$, decreasing in x . For a homogeneous (single fringe-status) group,

$$S^{\text{crop}} = G(\phi_{CL}), \quad S^{\text{farm}} = G(\phi^*), \quad S^{\text{live}} = G(\phi^*) - G(\phi_{CL}), \quad S^{\text{exit}} = 1 - G(\phi^*), \quad (78)$$

with the exit going to development in the fringe and to reversion R^N otherwise. The farmland-share equation of the main text is $S^{\text{farm}} = 1 - \exp(-(\omega^*)^{-\theta})$, $\omega^* = \phi^*/A_c$.

Equation (78) is the Regime-I sort. Under Regime II the pure grazing band is empty and the crop band is truncated from below at ϕ_C^* , so the shares become

$$S^{\text{crop}} = G(\phi_C^*), \quad S^{\text{live}} = 0, \quad S^{\text{farm}} = G(\phi_C^*), \quad S^{\text{exit}} = 1 - G(\phi_C^*), \quad (79)$$

with the exiting mass on $[0, \phi_C^*)$ — which now includes the displaced low crop band $[\phi_{CL}, \phi_C^*)$ — converting to development. The crop share of farmland $S^{\text{crop}}/S^{\text{farm}}$ jumps to 1 on the Regime-II fringe, reversing the soil-ordered composition of the Regime-I shares below. The county object is the M_c mixture over fringe status (Appendix B.17): the fringe slice of a high-pressure county can be in Regime II while its non-fringe slice, facing only the reversion option, stays in Regime I — so retention and displacement coexist within one county.

The held band: an optimal-stopping state. A parcel priced out of its intensive crop use but sitting near the fringe faces a real-option timing choice (Capozza and Helsley, 1989, 1990): convert now at $R^O - \kappa$, or *hold* in a low-input cover use that earns a floor flow $R^H \geq 0$ while retaining the discounted development option. With conversion irreversible and R^O rising in M_c , the optimal-stopping threshold sits above the myopic breakeven, opening a held band $\phi \in [\phi_{\text{hold}}, \phi_{\text{bind}}^*)$ below the binding farm-exit cutoff (ϕ^* in Regime I, ϕ_C^* in Regime II). A parcel in the held band sheds its intensive crop revenue (cropland and crop revenue fall), keeps its acres counted as farmland or idle grass (pasture held), and winds its extensive operation down (livestock revenue per pasture acre falls) — the acreage-versus-revenue wedge that signs the presence of parking. Its held slice enters the outside-use omission as the bounded interval $[0, R^O - \kappa]$ derived in Appendix B.10. The held band delivers the cropland-down, pasture-held, livestock-revenue-down signature; it does *not* deliver the fertilizer rise, which is the separate CES land-price intensification of Prediction 2 on the *surviving* crop band above ϕ_{bind}^* — a distinct set of parcels that clears the bar. The held state is distinguished from ordinary peri-urban change and from marginal-land retention by a same-parcel test: middle-land parcels that go crop-to-idle-grass under pressure should convert to developed use in a later census wave, whereas retained marginal land stays in grazing indefinitely. Where the linked panel cannot support that test, a within-class gradient in the crop-to-idle-grass response toward the built-up edge stands in for it; the acreage-versus-revenue wedge is the observable signature of the modeled band, not a substitute for it.

Composition is soil-ordered. The crop share of farmland is $S^{\text{crop}}/S^{\text{farm}} = G(\phi_{CL})/G(\phi^*)$, **monotone increasing in A_c :** writing $u = (A_c/\phi^*)^\theta$ and $s = (\phi^*/\phi_{CL})^\theta \in (0, 1)$, the ratio is $(1 - e^{-su})/(1 - e^{-u})$, increasing in u hence in A_c . It rises from a floor $(\phi^*/\phi_{CL})^\theta$ as

$A_c \rightarrow 0$ (marginal counties livestock-dominated but not exclusively) to 1 as $A_c \rightarrow \infty$ (good counties crop-dominated). This is Assumption A3's comparative-advantage ordering, recovered from a single suitability index. The gradient is the identifying content of A3, and it is selection-free: the crop share is a headcount of parcels above a fixed cutoff, not a return averaged over the parcels that chose crops, so it carries none of the activity-specific truncation that would contaminate a comparison of per-acre returns. In the data the crop share of farmland rises monotonically across the suitability classes, from 0.183 on marginal land to 0.381 and 0.605 on the two suitable classes, with a strictly positive suitability-level slope; the same slope holds when the sample is restricted to counties with no development pressure, so the ordering is comparative advantage rather than a pressure or feasibility artifact. This is the allocation analog of the land-share identification of θ (Section 7.2). The activity-specific levels b_C and b_L are a separate matter: because each activity's per-acre revenue is observed only over its own quality band – crops on the high tail, extensive grazing on the middle band – dividing observed revenue per acre by the band's conditional-mean quality recovers an average product over a selected truncation, not the slope of equation (24), and cannot be compared across activities to sign $b_C - b_L$. A3 is identified by the allocation gradient above and by the reduced-form co-movement (extensive grazing and its revenue fall, the crop mix rises), not by the activity-specific level ratios.

Comparative statics. With $G'(x) = -F(x) \theta A_c^\theta x^{-\theta-1} < 0$ and $F = 1 - G$:

- *Retention (Prediction 1).* Beyond the fringe $\phi_{\text{none}}^* = (R^N - \kappa_N + c_L)/b_L$, so $\partial S^{\text{farm}}/\partial\tau = G'(\phi^*) \partial\phi^*/\partial\tau$. If reach raises the livestock farm-gate ($\partial b_L/\partial\tau < 0$), ϕ^* falls and S^{farm} rises (provided the farm-gate response dominates any access-induced movement in R^N): marginal land that would revert stays in grazing. Present at $M_c = 0$, driven by $\partial b_L/\partial\tau$.
- *Livestock composition (the livestock implication).* $S^{\text{live}} = G(\phi^*) - G(\phi_{CL})$, and crop count is $S^{\text{crop}} = G(\phi_{CL})$ with $\partial S^{\text{crop}}/\partial\tau = G'(\phi_{CL}) \partial\phi_{CL}/\partial\tau$. The common cost response fixes the gap $c_C - c_L$, but A1 leaves the price legs $\partial b_C/\partial\tau$ and $\partial b_L/\partial\tau$ unordered, so ϕ_{CL} – and hence S^{crop} – may move with access. The movement is good-land-loaded: $G'(\phi_{CL}) = -F(\phi_{CL}) \theta A_c^\theta \phi_{CL}^{-\theta-1}$ carries the factor A_c^θ , which vanishes as $A_c \rightarrow 0$, so a ϕ_{CL} shift moves negligible crop acreage on marginal soil and bites only where mass sits near the cutoff. The model thus predicts no marginal-concentrated commodity-crop expansion – commodity cropland on the marginal class is flat to good-land-concentrated – and on the marginal base the livestock share moves through ϕ^* , expanding when farming is retained on the livestock-dominated marginal land.
- *Displacement and decoupling (Prediction 2, acreage and revenue).* Under development pressure the rising land-rent bar (the development option R^O) prices the marginal, lowest-value-per-acre farm use off the land first: the marginal extensive-grazing band is priced out as the exit cutoff ϕ^* climbs through it, so acreage falls and livestock revenue per acre collapses, while the surviving mix tilts toward input-intensive cropping and input use rises (Appendix B.8). Per-acre revenue falls because the priced-out grazing dominates the affected marginal land, not through any selection

gain in survivor quality. Beyond the fringe the cutoff moves the other way under net reach (ϕ_{none}^* down), so marginal grazing land is retained (Prediction 1) and acreage rises – the two destinations of the same marginal band, sorted by pressure. Prediction 2 is therefore a pressured-stratum statement in acreage and revenue per acre; the per-acre value level is flat as the residual of the revenue decline and the Prediction 4 capitalization wedge, not a compositional gain.

Fringe/non-fringe mixture and the exit destination. A share M_c of developable parcels lie in the fringe (outside option $R^{O,\text{press}} - \kappa$, exit = develop), $1 - M_c$ beyond it (outside option R^N , exit = revert). The crop–livestock cutoff ϕ_{CL} is common to both groups while $\phi_{\text{press}}^* > \phi_{\text{none}}^*$. County shares:

$$S_{ct}^{\text{farm}} = M_c G(\phi_{\text{press}}^*) + (1 - M_c) G(\phi_{\text{none}}^*), \quad S_{ct}^{\text{crop}} = G(\phi_{CL}) \text{ (moderate regime)}. \quad (80)$$

In the moderate regime crop parcels (above ϕ_{CL}) out-bid development, so pressure develops only the livestock band: crop count is M_c -independent and livestock concentrates off-fringe ($G(\phi_{\text{press}}^*) < G(\phi_{\text{none}}^*)$). In the intense regime ($R^{O,\text{press}} - \kappa > \bar{R}$) the marginal crop parcels near ϕ_{CL} develop, contracting cropland under pressure on high-suitability cells. The model *qualitatively signs* this contraction but, under unbounded Fréchet, does not quantify it, so the high-suitability-under-pressure cropland contraction is a confirmatory direction, not a fitted moment.

Table A23: Within-Agriculture Comparative Statics by Development-Pressure Stratum

Outcome	Axis	Predicted direction
Agricultural land share (S_{ct}^A) (survival margin)	Marginal class (FAO N)	Retained
	Suitable classes (S2–S1)	Net farm-gate incidence (A1)
Cropland (survival margin)	Marginal class (FAO N)	Retained / livestock base
	Suitable classes (S2–S1)	Net farm-gate incidence (A1)
Livestock (cattle inventory) (composition margin)	Unpressured $\rightarrow$ pressured	Rises $\rightarrow \approx 0$

Notes. Directional predictions from Prediction 1 and the livestock implication, under Assumptions A1–A3. The survival-margin rows (agricultural land share, cropland) run on a *soil* axis, not a pressure contrast: the marginal farm use is uniformly livestock, but comparative advantage makes low-suitability counties livestock-dominated (the crop share $G(\phi_{CL})/G(\phi^*)$ is monotone increasing in A_c), so the marginal class (FAO N) is retained through grazing against the reversion counterfactual R^N while the suitable classes turn on the ambiguous net farm-gate incidence of A1. Any access-induced shift in the crop–livestock cutoff ϕ_{CL} loads on good land, since the parcel density $G'(\phi_{CL}) \propto A_c^\theta$ vanishes as $A_c \rightarrow 0$; the model accordingly predicts no marginal-concentrated commodity-crop expansion. The survival gradient is present at $M_c = 0$ and is not mediated by development pressure. The composition row (cattle inventory) rises on the retained marginal no-pressure base; in the fringe rising pressure sheds the marginal livestock band to development first, so the model signs the marginal-class level and predicts the development-pressure interaction near zero. Suitability scales the magnitude of each response but not its sign.

Mapping to Agricultural Census outcomes. Predictions 1–3 determine how much land exits agriculture; Prediction 4 isolates the capitalization of the outside option, and the livestock implication determines the product composition among surviving parcels. The predictions are complementary and are jointly tested in the pressure-stratified results of Section 6.

Appendix B.17 The within-county pressure-share mixture

This subsection derives the area-weighted mixture and the first-order dose-response of Section 3.3, and states the curvature the linear form approximates. Within county c , partition the *undeveloped (convertible)* parcels – the population over which S_{ct}^A is defined – by frontier membership: a share M_c face the development bid $R_{ct}^O - \kappa_{ct}$ (exit = develop) and a share $1 - M_c$ face the reversion $R_{ct}^N - \kappa_N$ (exit = revert), with $R_{ct}^O - \kappa_{ct} > R_{ct}^N - \kappa_N$. The parcel draw ω_{jc} is independent of frontier membership and A_c is common within the county, so each subgroup is a Fréchet population whose marginal farm use (livestock) faces its own threshold through equation (6):

$$\omega_{m,ct}^* = \frac{\rho_{ct}^m + c_L(\tau_{ct})}{b_L(\tau_{ct}) A_c}, \quad m \in \{\text{none}, \text{press}\}, \quad \rho^{\text{press}} = R^O - \kappa, \quad \rho^{\text{none}} = R^N - \kappa_N, \quad (81)$$

which gives $\omega_{\text{press},ct}^* > \omega_{\text{none},ct}^*$. The subgroup land shares follow from equation (7), and the county land share is the area-weighted average of the two subgroups,

$$S_{ct}^A = M_c S^A(\omega_{\text{press},ct}^*) + (1 - M_c) S^A(\omega_{\text{none},ct}^*), \quad (82)$$

which is equation (9). Differencing across the access treatment, with $\Delta_m \equiv \Delta S^A(\omega_{m,ct}^*)$ the treatment effect on subgroup m ,

$$\Delta S_{ct}^A = M_c \Delta_{\text{press}} + (1 - M_c) \Delta_{\text{none}} + (\Delta M_c) [S^A(\omega_{\text{press}}^*) - S^A(\omega_{\text{none}}^*)]. \quad (83)$$

Because M_c is built from 1940 centers and 1940 geography it is predetermined and does not respond to the interstate, so $\Delta M_c = 0$ and the third term drops. Holding Δ_{press} and Δ_{none} fixed within a suitability class then leaves

$$\Delta S_{ct}^A = \Delta_{\text{none}} + M_c (\Delta_{\text{press}} - \Delta_{\text{none}}), \quad (84)$$

a level Δ_{none} and a slope $\Delta_{\text{press}} - \Delta_{\text{none}}$ in the pressure share. This is exact in M_c once the two subgroup responses are treated as constants, but the constancy is the approximation: Δ_m depends on the subgroup threshold through the nonlinear land-share map (7), whose curvature is $\partial^2 S^A / \partial (\omega^*)^2 = -\theta (\omega^*)^{-\theta-2} \exp(-(\omega^*)^{-\theta}) [\theta (\omega^*)^{-\theta} - (\theta + 1)]$. The linear dose-response is therefore a first-order reading valid where the two thresholds bracket a region of roughly constant slope; because M_c is right-skewed with most counties near zero (Section 4.1), the dose-response is evaluated mostly near the unpressured corner, and the binary $M_c^b = 1\{M_c > 0\}$ contrasts counties within reach of a city against those with no urban center in range, while the structural GMM further splits the positive-pressure mass at its median into low and high strata.

Appendix B.18 Construction of the development-pressure index M_c

This appendix gives the full construction summarized in Section 4.1.

Towns. The baseline measure counts the 1940 incorporated census places of at least 25,000 residents, taken from the complete full-count census-place layer; each enters as a point at its census coordinates carrying its 1940 population. This cutoff isolates the settlements large enough to project a genuine development frontier, and leaves a well-defined control of counties with no urban center in reach. The measure uses these census places directly, rather than urban blocks delineated from the built-up raster, so it inherits no raster measurement error on the town side. The built-up record enters only through the convertible denominator below.

Convertible land. The convertible denominator is built at the native 250 m resolution of the HISDAC-US built-up record and aggregated to a 1 km working lattice. A 250 m cell is developed when its 1940 built-up intensity is at least 1,000 m² of gross indoor floor area; the threshold de-dusts the record, which counts every structure type including agricultural buildings (Leyk and Uhl, 2018; McShane et al., 2022), and is set non-circularly as the level that best separates incorporated census places from rural background. Open water (USGS National Atlas water bodies) is excluded as well. Convertible land is the remainder — undeveloped, non-water 1940 land. Thresholding at the native resolution and aggregating the resulting indicator to the 1 km lattice, rather than averaging intensity and thresholding the average, keeps the threshold at the scale at which it was calibrated.

Coverage-gap footprint imputation. The built-up record under-records towns in some regions: of the 7,306 towns, 83 percent sit on a developed core in the 1950 raster, 14 percent are absent but flagged by HISDAC’s uncertainty layers, and 3 percent are absent and unflagged. A coverage-gap town’s own built core would otherwise be counted as convertible land, inflating its measured pressure. We therefore impute that core: log developed-footprint area is regressed on log population and Census region across the towns the record does cover, and each coverage-gap town is assigned the predicted footprint for its population and region, netted out of its convertible denominator. The imputation uses population and region only — never the under-recorded rural built-up record, and no market-access information — and matters most for the larger coverage-gap towns, whose cores a flat default would badly under-size.

The reach. Each town j is assigned a circular reach of radius $r_j = \min\{k P_j^{1/3}, \bar{r}\}$. The cube-root exponent follows from urban allometry: a settlement’s built-up area scales with population to roughly the two-thirds power, so its spatial extent scales to the one-third. The constant k is pinned by a single anchor — a city of one million residents reaches sixty kilometers — which implies reaches of about eight, eighteen, twenty-eight, and sixty kilometers at populations of 2,500, 25,000, 100,000, and one million. The cap $\bar{r} = 75$ km guards the metropolitan tail and binds for almost no 1940 city. Because the reach is a function of the census population alone, it requires no calibration on realized post-1940

growth and carries no free per-town parameter: a larger town's development frontier reaches farther, sub-linearly. The coastline is handled geometrically, by clipping the reach to land.

The measure and cutoffs. For each town the reach disc is drawn over the lattice; the discs are unioned, so a county within range of several towns is counted once. M_c is the convertible-land-weighted share of the county that falls inside the union, equation (11), with developed and water cells carrying no weight in either the numerator or the denominator. The convertible denominator is identical across the four town-population cutoffs — only the set of towns drawing a reach changes with the cutoff — so the four surfaces share a denominator and differ only in which towns project pressure.

Distinctness and cutoff sensitivity. The distribution of M_c and its distinctness from market access vary across the 2,500, 10,000, 25,000, and 50,000 town cutoffs. Raising the cutoff drops smaller towns from the city set, so pressure concentrates on larger urban centers, coverage falls, and the surface pulls away from the broad market-access gradient. The baseline $\geq 25,000$ cutoff sits in this range: it leaves roughly two-thirds of counties with no urban center in reach, a clean no-pressure control, while retaining enough coverage for the pressured stratum to identify the dose-response. Development pressure and market access nonetheless co-move through shared proximity to cities, so the moderator is not orthogonal to the market-access gradient; identification does not require orthogonality, because market potential enters the DID first stage year-interacted, so the pressure coefficient is read holding the market-access gradient fixed. Figure A14 shows the four national surfaces, which are genuinely distinct rather than rescalings of one another; the $\geq 25,000$ surface is the baseline.




Figure A14: County development pressure M_c under four town-population cutoffs ($\geq 2,500$, $\geq 10,000$, $\geq 25,000$, $\geq 50,000$). Raising the cutoff drops smaller towns from the city set, concentrating pressure on larger urban centers and shrinking coverage. The $\geq 25,000$ surface is the baseline.