

# The Ostrich Effect: Ignoring the Opportunity Cost of U.S. Golf Courses

*A Highest and Best Use Valuation with a Honolulu County Validation Study*

Michael Manolakis — University of Hawai‘i at Mānoa — Spring 2026

## Abstract

This thesis quantifies the aggregate *gross* opportunity cost of recreational land in the United States by applying the appraisal principle of Highest and Best Use (HBU) to a national dataset of 16,297 golf courses, with a Honolulu County validation study. A tri-language pipeline (Python, R, Julia) integrates Federal Housing Finance Agency (FHFA) and United States Department of Agriculture (USDA) county-level land valuations with OpenStreetMap-derived acreage, resolving 28.8% missing acreage through Multiple Imputation by Chained Equations (MICE;  $M = 100$  per language) pooled via Rubin’s Rules. The pooled aggregate gross opportunity cost is approximately **\$944 billion** under the unrestricted HBU counterfactual (calculated prior to the deduction of current capitalized recreational value,  $V_{Current}$ ), with an observed-only floor of **\$788 billion** computed without imputation. Applying standard land-residual bulk discounts of 25–60% — reflecting the gap between finished-lot per-acre prices and the raw-acreage price net of infrastructure costs, absorption risk, and developer profit (Appraisal Institute 2020) — yields a developer-economics range of approximately **\$378–\$708 billion**, with a central estimate near **\$614 billion** at the conventional 35% rate. Spatial concentration is extreme: California, Florida, New York, Texas, and Hawaii together account for over half of the national total, and a pooled logarithmic decomposition shows urban courses valued at approximately  $60\times$  rural courses on a per-acre basis, robust across all three implementations.

The Honolulu County validation integrates this national model with parcel-level cadastral and zoning data. A full pipeline analysis identifies 1,072 unique Tax Map Key parcels comprising Oahu’s golf footprint with a pooled opportunity cost of \$26.67 billion. Crucially, **81.7% of this acreage sits within Preservation or Federal/Military zones**, where residential redevelopment is legally prohibited; golf occupies 25.4% of all Resort-zoned land on Oahu and concentrates 63.2% of its parcels in the Ewa district, the island’s primary suburban housing development corridor. Under the legally-permissible HBU framework, only approximately **\$1.2 billion** of the gross \$26.67 billion is realizable within current zoning — a 4.5% pass-through that captures the binding constraint the unrestricted framework obscures. A pilot comparison of six high-profile Hawaii courses against official tax assessments produces a model-to-assessed ratio that characterizes the gap between gross HBU and going-concern current-use value in this specific market, reported in Section 4.4.1.

The macroeconomic opportunity cost of golf course land is large, geographically concentrated, and empirically tractable, but its realization is bounded by a layered regulatory architecture the HBU framework alone cannot dissolve.

## 1 Introduction

The allocation and regulation of urban and suburban land represent critical challenges within modern spatial economics. In many major United States metropolitan areas, the supply of developable land is highly inelastic, driven largely by restrictive land-use regulations rather than physical geographic constraints (E. Glaeser et al. 2017; Gyourko et al. 2021). This artificial scarcity has produced a pronounced divergence between the physical cost of construction and the market price of housing, generating a substantial regulatory “zoning tax” that inflates residential land values (E. L. Glaeser et al. 2005; Davis, Larson, et al. 2021).

Amidst this intense competition for spatial resources, the American golf course represents a massive and highly specific land-use typology. Historically, real estate developers utilized golf courses as

“surrogate public goods” to cross-subsidize the construction of master-planned communities, capturing exorbitant price premiums on adjacent residential lots (Frank F Limehouse et al. 2012). Empirical evidence suggests, however, that the marginal amenity value of proximity to golf courses has depreciated over time as consumer preferences shift toward natural open spaces (Cho et al. 2009). Despite this structural decline in recreational demand and the escalating value of residential land, hundreds of thousands of acres of prime real estate remain locked in this single-use capacity (Weinand et al. 2025).

Honolulu County provides an unusual empirical setting for validating these national-scale claims at the parcel level: the City and County maintains comprehensive cadastral and zoning data at Tax Map Key (TMK) resolution, the regulatory landscape is dominated by Preservation and Federal/Military zoning, and the island of Oahu hosts 33 distinct golf courses on over 6,000 acres of developable land. The two scales are complementary — the national dataset establishes the magnitude and geographic structure of the spatial misallocation, while the Honolulu micro-study tests whether that magnitude survives contact with cadastral reality. The personal context that motivated this research design is recorded in Appendix B.

The title invokes the “ostrich effect,” which in its original behavioral-finance usage (Galai and Sade 2006; Karlsson et al. 2009) denotes investors’ tendency to avoid information about negative portfolio outcomes. This thesis extends the metaphor from individual cognition to institutional architecture: the regulatory, contractual, and political structures that govern golf-course land use today neither measure nor publish the opportunity cost of the status quo. The figures documented in this study are not hidden in any technical sense — they can be computed from public datasets — but they are not assembled, reported, or used in policy deliberation. Documenting that aggregate is the analytical contribution of this thesis.

## 1.1 Research Framework

To quantify the macroeconomic inefficiency of recreational land allocation, this thesis applies the foundational real estate appraisal principle of Highest and Best Use (HBU). HBU dictates that land should be valued based on its most profitable, legally permissible, and physically possible use (Appraisal Institute 2020). By treating the existing golf course as the baseline, a counterfactual valuation model can establish the aggregate opportunity cost ( $V_{OC}$ ) of this land. The core economic framework is defined by the following equation:

$$V_{OC} = \max(V_{Res}, V_{Ag}) - V_{Current} \tag{1}$$

where  $V_{Res}$  represents the land value under a residential development counterfactual,  $V_{Ag}$  the value under an agricultural counterfactual, and  $V_{Current}$  the capitalized value of the land in its current recreational use.

In a frictionless, perfectly competitive market, land would naturally transition to the use that maximizes its residual value. The transition of golf course land to its HBU is severely impeded, however, by a dual layer of institutional frictions. Publicly, municipal zoning ordinances frequently prohibit high-density redevelopment. Privately, properties are often encumbered by restrictive covenants and Homeowner Association (HOA) agreements that legally mandate the preservation of the golf course in perpetuity, regardless of its financial feasibility (Ellickson 2022; Korngold 2024). To isolate the pure economic value of the land, this thesis adopts a formal *Hypothetical Condition* (Appraisal Institute 2020) that suspends these legal frictions for analytical purposes. The Honolulu micro-study

then re-introduces zoning as a binding constraint, yielding a *legally-permissible HBU* that bounds the unrestricted estimate.

## 1.2 Research Questions

This thesis addresses two coupled research questions:

1. *What is the aggregate financial opportunity cost of U.S. golf courses when evaluated against their Highest and Best Use counterfactuals of residential or agricultural development, and how is that cost distributed spatially?*
2. *When the national HBU framework is validated against parcel-level cadastral and zoning data in Honolulu County, how closely do the model’s estimates align with municipal tax assessments, and what share of the estimated opportunity cost is bounded by current zoning?*

## 1.3 Contributions

This thesis makes three contributions to the literature on land use, housing supply, and regulatory economics. First, it constructs and publishes a national-scale spatial dataset of 16,297 golf courses with imputed acreage and county-anchored land valuations, implemented in three independent computational stacks (Python, R, and Julia) for cross-language replicability. Second, it produces what is, to my knowledge, the first parcel-level validation of an HBU-based national opportunity cost model against municipal tax assessment data, anchored in 1,072 unique TMK parcels from the City and County of Honolulu. Third, it documents a regulatory paradox specific to high-value coastal markets: the bulk of golf course land in Honolulu County sits in Preservation or Federal/Military zones, meaning the very locations where opportunity cost is highest are also those where statutory redevelopment is most restricted.

## 1.4 Related Literature

This study bridges four strands of economic and legal literature: the macroeconomic valuation of urban land and regulatory scarcity, the microeconomic pricing of recreational amenities, the institutional frictions that govern land-use transitions, and the property tax assessment literature that informs the Hawaii validation. Each is treated briefly here; a section-by-section reading guide is recorded in Appendix C.

*Urban land valuation and regulatory scarcity.* The opportunity cost associated with underutilized urban land is driven by the escalating value of residential real estate ( $V_{Res}$ ), which foundational work by Davis and Palumbo (2008) attributes almost entirely to the price of underlying land rather than physical construction costs. Davis, Larson, et al. (2021) extend this with tract-level estimates of residential land prices, confirming steep spatial gradients near central business districts; their county-level dataset, distributed by the FHFA, anchors the urban valuation proxy used here. The appreciation reflects artificial supply constraints: E. L. Glaeser et al. (2005) introduce the “zoning tax” as the gap between market price and marginal construction cost, and Gyourko et al. (2021) document that this tax exceeds typical household incomes throughout San Francisco, Los Angeles, and Seattle. E. Glaeser et al. (2017) extend the welfare cost of these constraints to the national level via labor-misallocation deadweight loss, while Cheshire and Sheppard (2002) show in a UK welfare analysis that restrictive land-use planning operates as a regressive tax on housing affordability.

*The microeconomics of recreational amenities.* Modeling  $V_{Current}$  requires understanding the economic mechanisms that generate golf course revenue. Building on Rosen (1974)’s two-stage hedonic

framework, Frank F. Limehouse et al. (2010) show that environmental quality and recreational exclusivity are credence goods capitalized directly into pricing. The dominant driver, however, is the “joint production” of the course with adjacent real estate (Frank F Limehouse et al. 2012) — developers use the course as a surrogate public good to cross-subsidize the project and extract premiums from surrounding lots. The amenity value is structurally depreciating: Cho et al. (2009) find that while open space substitutes for private lot size, the premium specifically associated with golf-course proximity has declined measurably even as the premium for natural open space has risen. Weinand et al. (2025) push the spatial inefficiency argument further, noting that the global land footprint of golf vastly exceeds the area used for utility-scale renewables, introducing non-residential counterfactuals for land optimization.

*Institutional frictions and land-use rigidity.* If  $V_{Res} > V_{Current}$ , neoclassical theory dictates redevelopment, yet the transition rarely occurs — a failure attributable to severe institutional frictions in the spirit of Coase (2013). The relevant case is positive transaction costs, in which the initial rights assignment determines equilibrium allocation; the current architecture grants veto power to incumbent homeowners and golf-course members through three channels. Municipal zoning prohibits the densification HBU requires (Fischel 2015); restrictive covenants tether the majority of modern developments to perpetual preservation of existing uses, modifiable only by supermajority (Ellickson 2022); and Korngold (2024) argues that legislative zoning reform is insufficient unless parallel private covenants are voided under public policy doctrine, though such voiding may itself constitute a regulatory taking. These layered frictions justify the use of a Hypothetical Condition in the national valuation while motivating the binding-zoning analysis applied in the Hawaii micro-study.

*Property tax assessment and market value.* The Hawaii validation compares model HBU estimates against official tax assessment values, which a substantial appraisal literature documents systematically diverge from market values (Appraisal Institute 2020). Standard appraisal practice distinguishes value-in-current-use (relevant for ad valorem taxation) from market value at HBU (relevant for transactions), with the latter typically exceeding the former when current use is suboptimal. A model-to-assessed ratio greater than unity is therefore not evidence that the model is wrong, but evidence that the model is measuring something different from what assessors measure. The structure of that divergence is itself an economically meaningful object of study, and is one of the central empirical objects this thesis characterizes.

## 2 Theoretical Framework

This section establishes the formal economic model used to quantify the spatial misallocation of recreational land. The framework integrates neoclassical urban economics with the professional appraisal doctrine of Highest and Best Use (HBU) to derive a mathematical function for aggregate opportunity cost, then introduces a legally-permissible variant that bounds the unrestricted estimate by binding zoning constraints.

### 2.1 Neoclassical Land Allocation and Highest and Best Use

In a frictionless, perfectly competitive spatial market, land is allocated to the use that generates the maximum residual value. This principle aligns with the appraisal standard of Highest and Best Use (HBU), defined as the reasonably probable and legal use of vacant land or an improved property that is physically possible, appropriately supported, financially feasible, and results in the highest value (Appraisal Institute 2020).

For a given parcel of land  $i$ , let  $V_{i,k}$  represent the present value of the land under use  $k$ , where  $k \in \{1, 2, \dots, K\}$  represents the set of all physically possible and financially feasible uses. Under conditions of spatial equilibrium, the market value of the parcel,  $V_i^*$ , is determined by the use that maximizes this value:

$$V_i^* = \max\{V_{i,1}, V_{i,2}, \dots, V_{i,K}\} \quad (2)$$

When the current use of the land,  $V_{Current}$ , diverges from  $V_i^*$ , an economic inefficiency exists. The magnitude of this inefficiency is the opportunity cost.

## 2.2 The Opportunity Cost Function

To quantify the inefficiency of land currently allocated to golf courses, this analysis restricts the set of counterfactual uses  $K$  to the two most probable alternative land uses based on spatial geography: residential development ( $V_{Res}$ ) and agricultural production ( $V_{Ag}$ ). The fundamental opportunity cost ( $V_{OC}$ ) for a specific golf course parcel  $i$  is defined as the difference between its value under its HBU counterfactual and its value in its current recreational state:

$$V_{OC,i} = \max(V_{Res,i}, V_{Ag,i}) - V_{Current,i} \quad (3)$$

The aggregate national opportunity cost is the summation of  $V_{OC,i}$  across all  $N$  golf courses in the dataset. By construction,  $V_{OC,i}$  may in principle take a negative value when  $V_{Current,i} > \max(V_{Res,i}, V_{Ag,i})$  — i.e., when the capitalized recreational value exceeds both counterfactual uses. This case is unlikely for the parcels in this dataset but is theoretically possible for high-revenue resort courses in dense urban submarkets. The Hawaii micro-study addresses this directly by anchoring the gross estimate against tax-assessed values, producing an empirical bound on  $V_{Current}/V_{Res}$ .

### 2.2.1 Operational Form: Gross HBU Counterfactual

In the empirical national-scale implementation,  $V_{Current}$  is not directly observed for each of the 16,297 courses; the national-scale data sources used in this study (FHFA county-level residential land prices and USDA county-level agricultural land values) characterize  $V_{Res}$  and  $V_{Ag}$  but do not provide course-level capitalized recreational income. Setting  $V_{Current} = 0$  in Equation 3 produces the operational *gross HBU counterfactual*,  $V_{OC,i}^{gross} = \max(V_{Res,i}, V_{Ag,i})$ , which represents the upper bound on the per-parcel opportunity cost. The Hawaii micro-study (Section 3.3 and Section 4.4) addresses this limitation directly by anchoring the gross estimate against official tax-assessed values for a subset of Honolulu County courses, producing an empirically estimated model-to-assessed ratio that characterizes the relationship between  $V_{OC}^{gross}$  and a defensible proxy for  $V_{Current}$ .

### 2.2.2 Deconstructing the Counterfactuals

The variables within the core equation are defined by distinct microeconomic determinants:

- **$V_{Res}$  (Residential Value):** The value of the land if subdivided and developed for single-family or multi-family housing. In highly regulated metropolitan statistical areas,  $V_{Res}$  is artificially inflated by a regulatory zoning tax ( $\tau_z$ ), which represents the scarcity premium induced by supply-side restrictions (E. L. Glaeser et al. 2005; Gyourko et al. 2021). Therefore,  $V_{Res}$  is a function of organic housing demand ( $D_h$ ), physical construction costs ( $C_c$ ), and regulatory scarcity ( $\tau_z$ ).

- **V<sub>Ag</sub> (Agricultural Value):** The capitalized value of the land’s agricultural yield, primarily applicable to courses located in rural or exurban geographies where residential demand is insufficient to justify development costs.
- **V<sub>Current</sub> (Current Recreational Value):** The capitalized net operating income of the golf course. This value depends heavily on the joint production model in which the course serves as a surrogate public good that generates price premiums for adjacent real estate (Frank F Limehouse et al. 2012), and on operating fundamentals (rounds played, membership fees, food and beverage revenue) that vary substantially across facilities. In the Hawaii validation study, official municipal tax assessments serve as a defensible proxy for  $V_{Current}$ , since assessors value these properties under their current recreational use.

### 2.3 Legally-Permissible HBU

The HBU framework as stated assumes the land is legally available for redevelopment to its most productive use. In practice, this is rarely the case. Standard appraisal practice distinguishes between the unrestricted maximum-value use and the use that is legally permissible under current zoning, environmental, and contractual restrictions (Appraisal Institute 2020). To formalize this distinction, let  $L_i \subseteq K$  denote the subset of uses that are legally permissible for parcel  $i$  under current zoning. The legally-permissible HBU value is then:

$$V_{i,\text{legal}}^* = \max\{V_{i,k} : k \in L_i\} \quad (4)$$

and the corresponding legally-permissible opportunity cost is:

$$V_{OC,i}^{\text{legal}} = V_{i,\text{legal}}^* - V_{Current,i} \quad (5)$$

By construction,  $V_{OC,i}^{\text{legal}} \leq V_{OC,i}^{\text{gross}}$ , with equality only when current zoning permits the unrestricted HBU. In the Hawaii micro-study,  $L_i$  is operationalized using the Honolulu County zoning layer: parcels classified as Preservation (P-1, P-2) or Federal/Military (F-1) cannot legally support residential redevelopment, so for these parcels  $V_{Res}$  is excluded from the feasible set. This produces a zoning-bounded estimate of the realizable opportunity cost — one that quantifies how much of the unrestricted HBU value is rendered theoretical by current land-use law.

### 2.4 Institutional Frictions and Market Failure

If  $V_{OC,i} > 0$ , rational economic actors should execute a transaction to redevelop the land. The persistence of golf courses in high-value urban areas indicates a market failure driven by severe institutional frictions ( $F$ ). These frictions consist of public regulatory constraints ( $Z$ ), such as municipal zoning ordinances that prohibit high-density development, and private contractual constraints ( $C$ ), such as restrictive covenants enforced by Homeowner Associations (Ellickson 2022; Korngold 2024). Mathematically, redevelopment occurs only if the opportunity cost exceeds the total institutional friction cost:

$$V_{OC,i} > F(Z_i, C_i) \quad (6)$$

The Coasian framework is instructive here, but in a more limited way than is sometimes claimed (Coase 2013). With zero transaction costs, the initial allocation of rights does not affect the efficient

outcome. With positive transaction costs — the empirically relevant case — the initial allocation determines the equilibrium allocation. The current legal regime grants veto power over redevelopment to incumbent homeowners and golf course members through zoning and covenants, with the result that even when  $V_{OC}$  is large,  $F$  can be larger still. Voiding these veto rights, as proposed by Korngold (2024), would shift the equilibrium allocation by reassigning rights, but this shift is itself politically and constitutionally contested.

To isolate the pure economic value of the land — distinct from the regulatory architecture that currently governs its use — this thesis applies the appraisal doctrine of a *Hypothetical Condition* (Appraisal Institute 2020) at the national scale, setting  $F(Z_i, C_i) = 0$ . The Honolulu micro-study then re-introduces zoning as a binding constraint via the legally-permissible HBU framework defined above. The two estimates — unrestricted and zoning-bounded — are complementary: the first measures the economic potential of the land if institutional frictions were dissolved, and the second measures the share of that potential that current law permits to be realized.

### 3 Data and Methodology

To answer the research questions posed in Section 1.2, this thesis adopts a two-scale empirical research design. The national-scale analysis (Section 3.2) constructs a 16,297-course dataset and produces the aggregate opportunity cost estimate. The Honolulu County micro-study (Section 3.3) integrates parcel-level cadastral and zoning data to test the national model against ground-truth municipal records. The two scales are not independent: the Hawaii pipeline reuses the national-scale outputs as inputs, and the validation logic depends on both pipelines producing consistent estimates for the Hawaii subset of the national dataset.

#### 3.1 Two-Scale Research Design

The national pipeline establishes *breadth* — estimates for every county and every course in the contiguous United States plus Alaska and Hawaii, computed from federally maintained land-valuation datasets and OpenStreetMap-derived course geometries. The Honolulu pipeline establishes *depth* — parcel-level resolution for the 33 deduplicated golf courses on the island of Oahu, including assessed values, dominant zoning class, and Tax Map Key (TMK) identifiers from the City and County of Honolulu. Each scale is implemented in three independent computational stacks (Python, R, and Julia), and the analyses are pooled across both languages and Multiple Imputation runs using Rubin’s Rules. Statistical estimates are reported as both per-language Rubin-pooled values and as a Grand Mean across the three languages, providing a transparent measure of cross-implementation variability.

#### 3.2 National Pipeline (Phases 1–4)

##### 3.2.1 Dataset Scope and Sourcing

The primary dataset consists of 16,297 golf course observations across the United States, each with a latitude–longitude coordinate that locates the course geographically. From this starting point, the analysis assembles each course’s opportunity cost in three steps: (i) identifying which county the course sits in, so that local land-value benchmarks can be applied; (ii) determining how many acres the course occupies, using publicly mapped boundary polygons; and (iii) attaching a per-acre dollar value to that acreage, drawn from one of two national land-value datasets depending on whether the course is in an urban or rural county.

The two land-value datasets are the Federal Housing Finance Agency (FHFA) county-level residential land-price panel (Davis, Larson, et al. 2021), which prices each county’s residential land in dollars per acre, and the U.S. Department of Agriculture (USDA) National Agricultural Statistics Service county-level agricultural land values, which prices farmland the same way. Each course is assigned one of the two prices based on whether its county is classified as metropolitan (USDA Rural-Urban Continuum Code 1–3) or non-metropolitan (RUCC 4–9). Because both datasets report values at the county level, the analysis implicitly assumes that land values are constant within a county; this is a limitation acknowledged explicitly in Section 4.4 and quantified for the Honolulu case study, where sub-county price gradients are known to be substantial. Section 6 discusses possible methodological extensions that would relax this assumption.

### 3.2.2 Polygon Acreage and Imputation

Course acreage is not reported in the source dataset. To compute it, the boundaries of each course are extracted from OpenStreetMap (OSM), the volunteer-maintained global map dataset. Roughly 71% of the 16,297 courses are matched to an OSM polygon directly; for the remaining 29%, the acreage is estimated using Multiple Imputation by Chained Equations (MICE), a standard statistical technique for handling missing data (Rubin 1987; Buuren 2018). The MICE step generates many plausible values for each missing acreage, drawn from the joint distribution of the observed data; the final reported figure is the average across these draws, with confidence intervals that account for both the random variation in the draws and the imputation uncertainty itself.

To verify that the results are not driven by the specific choice of imputation software or by minor coding decisions, the entire pipeline is implemented independently in three programming languages: Python, R, and Julia. Each language reads the same input files and produces its own aggregate estimate. The three estimates are then averaged to produce the headline “Grand Mean” figure. Any large divergence between the three would indicate a pipeline bug; in practice, the three implementations agree closely, supporting confidence in the reported figures. The full technical narrative of the pipeline—data sources, projections, imputation parameters, and Rubin’s Rules for pooling estimates across imputed datasets—is provided in Appendix A.

### 3.3 Honolulu County Pipeline (Phase 5)

The Honolulu pipeline takes the national-scale outputs and integrates them with parcel-level data published by the City and County of Honolulu, allowing the analysis to descend from county-level averages to specific Tax Map Key (TMK) parcels. The pipeline is structured in two tracks: a manual pilot validation against tax assessments for six high-profile Hawaii comparator courses spanning three counties (Honolulu, Maui, and Hawaii; Phase 5a), and an automated full-pipeline analysis at the parcel level for all golf courses on Oahu (Phase 5b). The substantive Honolulu-County micro-study is carried by the Phase 5b parcel-level analysis; the Phase 5a comparator set is used as a limited assessment-ratio anchor for the Honolulu courses and as a forward-looking observation set whose neighbor-island ratios are reserved for the planned journal-paper extension to a stratified multi-county validation.

The Phase 5a validation produces an empirical model-to-assessed ratio reported in Section 4.4.1 (computed on a six-course Hawaii pilot subset), characterizing the gap between the model’s gross HBU estimate and the going-concern values that municipal assessors actually use in that market. The Phase 5b pipeline takes the OSM golf polygons for Oahu, intersects them with the Honolulu cadastre to produce a list of 1,072 unique TMK parcels comprising the Oahu golf footprint, and

intersects each parcel with the Honolulu zoning layer to assign each parcel its dominant zoning class. The result is a complete characterization of where Oahu’s golf land sits geographically and what zoning class governs each parcel—which in turn enables the Preservation Paradox finding reported in Section 4.4.3. As with the national pipeline, the technical specifics (projection codes, intersection logic, deduplication rules, cross-language parity checks) are deferred to Appendix A.

## 4 Empirical Results

This section presents the quantitative findings derived from the spatial valuation model and the multiply imputed datasets. Section 4.1 reports the aggregate national opportunity cost and characterizes the stability of the MICE-pooled estimate. Section 4.2 documents the spatial concentration of value across states and counties. Section 4.3 presents the pooled logarithmic regression. Section 4.4 introduces the Hawaii validation results, which serve as the empirical anchor for the national estimate.

### 4.1 Aggregate National Opportunity Cost

Pooling the results across  $M = 100$  imputed datasets per language and computing the Grand Mean across the three independent implementations, the aggregate opportunity cost ( $V_{OC}^{gross}$ ) of the 16,297 U.S. golf courses is approximately **\$944 billion**. The cross-language spread is narrow: the Python pipeline produces \$0.943T, R produces \$0.945T, and Julia produces \$0.943T, with the Grand Mean at \$0.944T.

To establish a strict lower bound that does not depend on imputation, the same pipeline was run restricted to the 11,605 courses with directly measured OSM polygon acreage (`acreage_source`  $\neq$  `MICE_Target`). The observed-only aggregate is **\$788 billion**, or 83.5% of the MICE-pooled figure. This result establishes that the bulk of the national opportunity cost is supported by directly measured course geometries; imputation contributes the remaining 16.5% by extending the estimate to the 28.8% of courses without OSM polygon coverage. The disproportion between the share of courses imputed (28.8%) and the share of value contributed by imputation (16.5%) indicates that imputed courses are on average smaller and lower-value than the observed subset, which is consistent with the spatial pattern of missing OSM coverage being concentrated in low-density rural areas.

The total acreage attributable to U.S. golf courses, also pooled across imputations, is approximately 2.30 million acres (Table 1). Of this, 1.70 million acres (74%) are in counties classified as urban under the USDA Rural-Urban Continuum Codes, and 598,000 acres (26%) are in counties classified as rural.

Table 1: National Golf Course Acreage Summary (MICE-Pooled,  $M = 100$ )

Category	County Type	Pooled Acres	SD (Between)	99% CI Lower	99% CI Upper
National Total	All	2,304,777.6	4,506.9	2,293,109.8	2,316,445.3
By County Type	Urban	1,706,434.0	3,246.0	1,698,030.6	1,714,837.4
By County Type	Rural	598,343.6	2,821.2	591,040.0	605,647.2

99% confidence intervals reported; 95% intervals available in the replication package.

A diagnostic check on the stability of the MICE imputation across the three language implementations is reported in Appendix A, Figure 12; the three language-specific imputation distributions

overlap closely with each other and with the observed density across the range of plausible course sizes, supporting the use of the Grand Mean as the headline figure.

## 4.2 Spatial Heterogeneity and Value Concentration

The aggregate opportunity cost is not distributed uniformly across the spatial landscape. Figure 1 maps the Grand Mean opportunity cost at county resolution, with values rendered on a  $\log_{10}$  color scale. High-value counties cluster on the West Coast (particularly California), in southern Florida, in the New York metropolitan area, and in selected Texas metros (Houston, Dallas, Austin). Most interior counties show low or negligible opportunity cost, consistent with the lower per-acre land values produced by the agricultural counterfactual in rural settings.

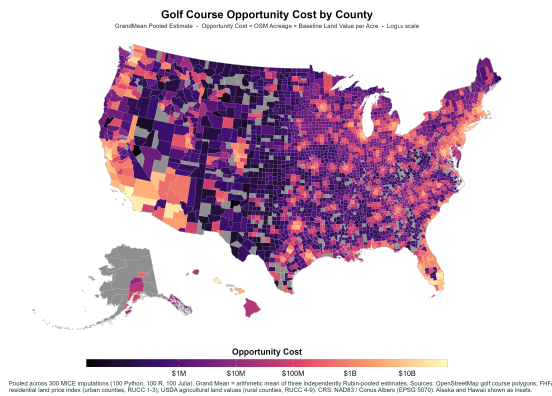


Figure 1: County-level Grand Mean opportunity cost across the three language implementations. Colors are rendered on a  $\log_{10}$  scale; gray counties contain no golf courses in the dataset.

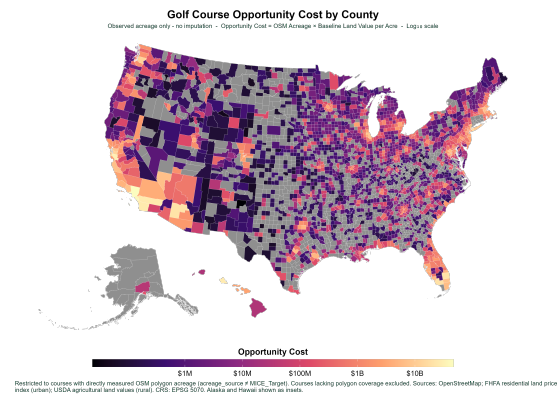


Figure 2: County-level opportunity cost computed only from courses with directly measured OSM polygon acreage (`acreage`  $\neq$  `MICE_Target`). The geographic concentration of value persists in the imputation-free subset.

The five highest-value states under the Grand Mean estimate are California (\$293.7B), Florida (\$112.0B), New York (\$53.9B), Texas (\$38.6B), and Hawaii (\$35.6B); the five highest-value counties are Los Angeles (\$49.7B), Orange (\$38.5B), Santa Clara (\$33.8B), San Diego (\$30.3B), and Honolulu (\$29.7B). Hawaii’s appearance in the top five at both the state and county levels — despite having only 74 courses statewide and 33 on Oahu — reflects the per-acre premium that prevails in geographically constrained urban housing markets, and motivates the parcel-level Hawaii validation in Section 4.4.<sup>1</sup>

The distributional structure of the opportunity cost across the 16,297 courses is highly unequal. The Lorenz curve of per-course opportunity cost (Figure 3) sits well below the line of equality, indicating that a small share of courses accounts for a disproportionate share of the aggregate value.

<sup>1</sup>The Honolulu County figure of \$29.7B reported here under the national pipeline differs from the Phase 5b parcel-level Oahu aggregate of \$26.67B reported in Section 4.4.2 by approximately 10.2%. This gap reflects the spatial-deduplication discount of clipping OSM golf polygons against the Honolulu cadastre, which reduces the total area attributed to Oahu golf parcels from 8,564 cadastral acres to 6,066 polygon-clipped acres (Appendix A.7). The Phase 5b figure is the canonical estimate for the Honolulu micro-study because it eliminates the double-counted parcel area that would otherwise inflate the per-course opportunity cost. The \$29.7B figure remains the appropriate input for cross-state comparison in this national-aggregate context.




Figure 3: Lorenz curve of per-course opportunity cost. The cumulative share of total opportunity cost rises slowly as a function of the cumulative share of golf courses (sorted ascending), indicating that high-value courses in urban housing markets dominate the aggregate. The three language curves and the Grand Mean curve are visually indistinguishable.

### 4.3 Econometric Determinants of Land Value

To characterize the systematic drivers of opportunity cost, a pooled logarithmic regression was estimated independently in each of the three language stacks. The model regresses the natural log of per-course opportunity cost against the number of holes and a binary indicator for urban county classification:

$$\ln(V_{OC,i}) = \beta_0 + \beta_1 \cdot \text{Holes}_i + \beta_2 \cdot I[\text{Urban}_i] + \varepsilon_i \quad (7)$$

The pooled coefficient estimates and standard errors, applying Rubin's Rules independently within each language group ( $M = 100$ ), are reported in Table 2.

Table 2: MICE-Pooled OLS Regression Results (Rubin’s Rules,  $M = 300$ : 100 Python, 100 R, 100 Julia). Dep. var.:  $\log(\text{Opportunity\_Cost})$ .

Parameter	Python				R				Julia			
	Coef.	SE	$p$	Sig.	Coef.	SE	$p$	Sig.	Coef.	SE	$p$	Sig.
Intercept	12.280	0.038	< 0.001	***	12.223	0.042	< 0.001	***	12.242	0.039	< 0.001	***
Urban County	4.167	0.019	< 0.001	***	4.004	0.022	< 0.001	***	4.165	0.021	< 0.001	***
Holes	0.048	0.002	< 0.001	***	0.053	0.003	< 0.001	***	0.048	0.002	< 0.001	***

*Note:*

Sig. codes: \*\*\*  $p < 0.001$ . Rubin’s Rules applied independently per language group ( $M = 100$  each). Grand Mean = arithmetic mean of three independently pooled estimates.

The regression serves three purposes, only one of which is the per-acre urban premium itself. First, it produces a methodological invariance result: three independent MICE backends — LightGBM in Python, Random Forest in R, and Mice.jl in Julia — converge on near-identical coefficients despite differing under-the-hood imputation algorithms. The intercept ranges from 12.223 (R) to 12.280 (Python), the Holes coefficient from 0.048 (Python and Julia) to 0.053 (R), and the Urban indicator from 4.004 (R) to 4.167 (Python); all are statistically significant at  $p < 0.001$  in every implementation. Figure 4 visualizes the three coefficient sets with both 95% and 99% confidence intervals shown as inner and outer bands, with all coefficients remaining significant at the more conservative 99% level. This invariance is the central methodological result of the section: the headline opportunity cost figure is not an artifact of any single imputation backend’s parametric assumptions.

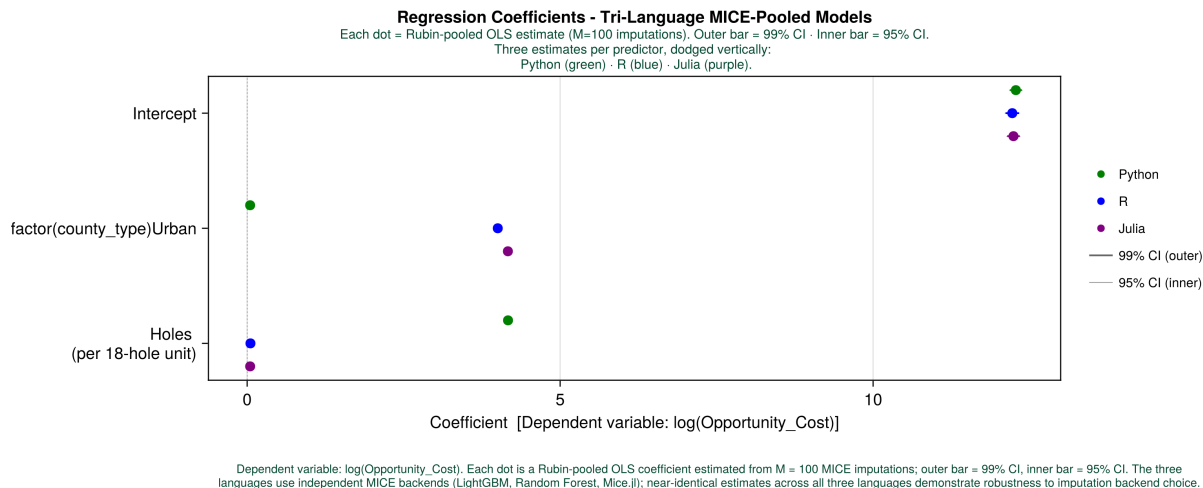


Figure 4: Pooled regression coefficients across the three language implementations, displaying both 95% (inner) and 99% (outer) confidence intervals. All coefficients are significant at  $p < 0.001$  under either bound. The y-axis is dodged to display the three estimates per parameter without overlap.

Second, the model provides a goodness-of-fit characterization of the two-predictor specification. The  $R^2$  values of 0.70 (R) to 0.77 (Python) indicate that course size (Holes) and urban-rural classification together explain 70–77% of the variation in per-course log opportunity cost. The residual 23–30% reflects per-course heterogeneity not captured by the two predictors — variation

in within-county price gradients, course-specific configuration, and assessment-classification effects that the parsimonious model does not attempt to resolve.

Third, the Holes coefficient ( $\beta_1 \approx 0.05$ ) is the only genuinely informative individual coefficient in the model. Holes was not used in the construction of  $V_{Res}$  or  $V_{Ag}$  in Phase 1; its association with log opportunity cost (approximately a 5% increase per additional hole) is therefore a substantive finding that emerges from the data rather than from the valuation algorithm. The interpretation is that larger courses concentrate in higher-value geographic areas, consistent with the development pattern in which 18- and 27-hole resort courses are sited in markets with the densest surrounding residential capitalization.

The Urban indicator ( $\beta_2 \approx 4.0$ – $4.2$ ), implying a per-acre ratio of  $\exp(4.1) \approx 60$  between urban and rural courses, should be read as confirmation of the hybrid valuation algorithm rather than as an independent discovery. Urban courses are assigned the FHFA residential price per acre, rural courses are assigned the USDA agricultural price per acre, and the FHFA-USDA per-acre ratio at the county level is itself approximately of this magnitude in many markets. The coefficient is therefore best understood as a *decomposition* of the bifurcated baseline value, not as a causal estimate of how a given course’s opportunity cost would change if relocated. The substantive claim it supports is that the urban premium dominates all other course-level covariates in the assignment, which is the assumption that motivated the hybrid algorithm in the first place.

#### 4.4 Hawaii Validation: Empirical Anchoring

The national-scale results characterize the magnitude and geographic structure of the spatial misallocation but do not directly verify that the model’s per-course estimates correspond to economically meaningful values on the ground. The Honolulu County micro-study (Phase 5) addresses this gap by validating the model against parcel-level municipal records.

The per-acre FHFA proxy for Honolulu County represents a countywide residential land-price average and does not distinguish between Oahu’s submarkets. Two properties at the high end of the per-course distribution — Turtle Bay Resort and Golf Club on the North Shore, and Makaha Valley Country Club on the Leeward coast — sit in submarkets whose actual residential transaction prices are substantially lower than the urban Honolulu corridor that anchors the FHFA index. Their headline opportunity cost figures should be understood as the upper bound consistent with the county-level proxy, not as parcel-specific valuations. Refining these estimates would require sub-county price gradients (Hawaii DBEDT submarket data, Honolulu Real Property Assessment Division zone-level transaction records), which is a methodological extension flagged in §6.

##### 4.4.1 Phase 5a: Model versus Tax Assessment

The pilot validation compares model HBU estimates against official 2022 tax assessment values for six high-profile Hawaii golf courses. The results are reported in Table 3.

Table 3: Phase 5a Pilot Validation: Model HBU Estimate vs. Official Tax Assessment

<b>Course</b>	<b>County</b>	<b>Model (\$M)</b>	<b>Assessed (\$M)</b>	<b>Ratio</b>
Turtle Bay Resort & Golf Club	Honolulu	2,270	1,850	1.23×
Waialae Country Club	Honolulu	718.7	620.0	1.16×
Kaanapali Golf Courses	Maui	565.3	480.0	1.18×
Wailea Golf Club	Maui	255.1	210.0	1.21×
Hualalai Golf Club	Hawaii	295.1	175.0	1.69×
Kohala Country Club	Hawaii	625.0	420.0	1.49×
<b>Average</b>	—	—	—	<b>1.33×</b>

The average model-to-assessed ratio across the six courses is 1.33×, indicating that the model’s gross HBU estimate exceeds going-concern assessed values for these Hawaii courses by approximately one-third. The ratio is narrowest in urban Honolulu (1.16–1.23×) and widest in rural Big Island courses (1.49–1.69×), consistent with the hypothesis that the FHFA urban residential proxy more closely tracks municipal assessment practice than the USDA rural agricultural proxy does. This pattern characterizes the gap between gross HBU and going-concern current-use value in this particular market; it does not constitute a calibrated national deflator, since the sample is small, geographically concentrated, and embedded in an assessment regime (Honolulu’s parcel-level cadastral system, Hawaii’s Act 273 agricultural classification framework, the state’s dual rate structure for owner-occupied versus non-owner-occupied parcels) that differs substantively from the mainland. The economic content of this gap, and its implications for the relationship between gross HBU and net realizable value, is discussed in Section 5.3.

The relationship between the model HBU estimate and a defensible agricultural-floor value across all 33 deduplicated Oahu courses is illustrated in Figure 5. For each course, the figure shows the agricultural-floor value (USDA per-acre rate × parcel acreage, the restricted-use lower bound) and the HBU model estimate (FHFA residential per-acre rate × parcel acreage, Rubin-pooled across the three languages). The horizontal gap between the two endpoints corresponds to what E. L. Glaeser et al. (2005) would identify as the zoning-tax-driven deadweight loss for that parcel.




Figure 5: The Hawaii Gap: HBU model value versus agricultural floor for each Oahu golf course. Left endpoint is the agricultural-floor value ( $\text{USDA} \times \text{acreage}$ ); right endpoints are the HBU estimates from the three language implementations and the Grand Mean. Dagggers ( $\dagger$ ) mark courses with above-average between-imputation variance.

#### 4.4.2 Phase 5b: Aggregate Oahu Opportunity Cost

Aggregating across all 33 deduplicated Oahu courses and pooling  $M = 100$  imputed economic datasets per language via Rubin’s Rules, the total Oahu opportunity cost is \$26.67 billion under the Grand Mean estimate. Per-language pooled estimates are \$26.79B (Python), \$26.68B (R), and \$26.54B (Julia), well within the cross-language spread observed at the national scale.

The geographic distribution of the 1,072 unique TMK parcels comprising Oahu’s golf footprint is shown in Table 4 and visualized in Figure 6. Zone 9 (the Ewa District, comprising Kapolei, Pearl City, and Ewa Beach) accounts for 678 of the 1,072 parcels (63.2%). This concentration reflects Oahu’s post-statehood suburban golf development corridor: the Ewa plain was the primary site of master-planned residential and resort development in the 1970s through 1990s, and golf courses were constructed alongside or within those communities under the joint production model documented by Frank F Limehouse et al. (2012).

Table 4: Hawaii Golf Course Parcel Distribution by Geographic Zone

Zone	District	Parcels	Share (%)
9	Ewa (Kapolei/Pearl City)	678	63.2%
3	Honolulu (Anomalies)	169	15.8%
4	Koolaupoko (Kailua/Kaneohe)	123	11.5%
1	Honolulu (Urban Core)	35	3.3%
5	Koolauloa (North/East)	33	3.1%
8	Waianae (West)	30	2.8%
2	Honolulu (East/Anomalies)	3	0.3%
7	Wahiawa (Central)	1	0.1%

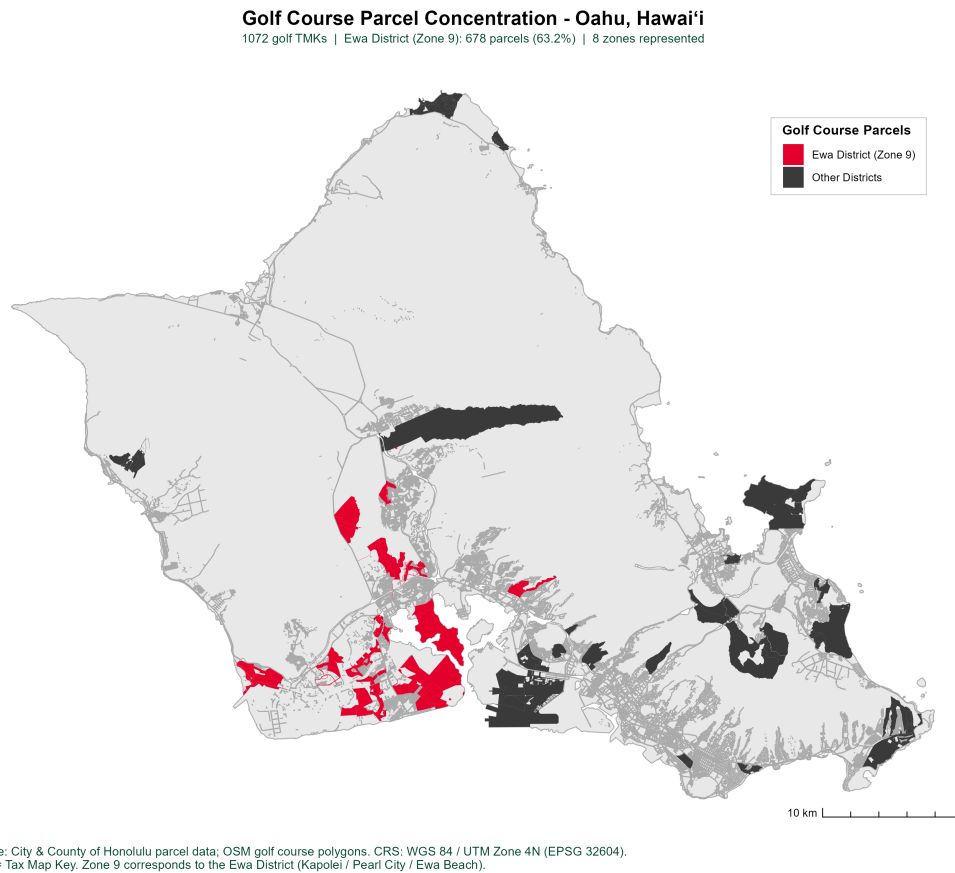


Figure 6: Spatial distribution of the 1,072 golf course TMK parcels on Oahu. Parcels in the Ewa District (Zone 9) are highlighted in orange; all other districts are shown in dark gray.

Translating the parcel-level footprint into per-course dollar values produces the spatial pattern shown in Figure 7. Of the 33 deduplicated Oahu courses, 29 are matched to OSM polygons within the 500-meter spatial-join cap and rendered with their Grand Mean opportunity cost; the remaining four exceed the cap and are not displayed. Per-course opportunity costs range from approximately \$286 million to \$2.27 billion, with the highest values concentrated in Turtle Bay on the North Shore

(yellow), Waialae Country Club in southern Honolulu (orange/red), and the Ewa-corridor courses (magenta) that dominate the southwestern parcel cluster.<sup>2</sup>

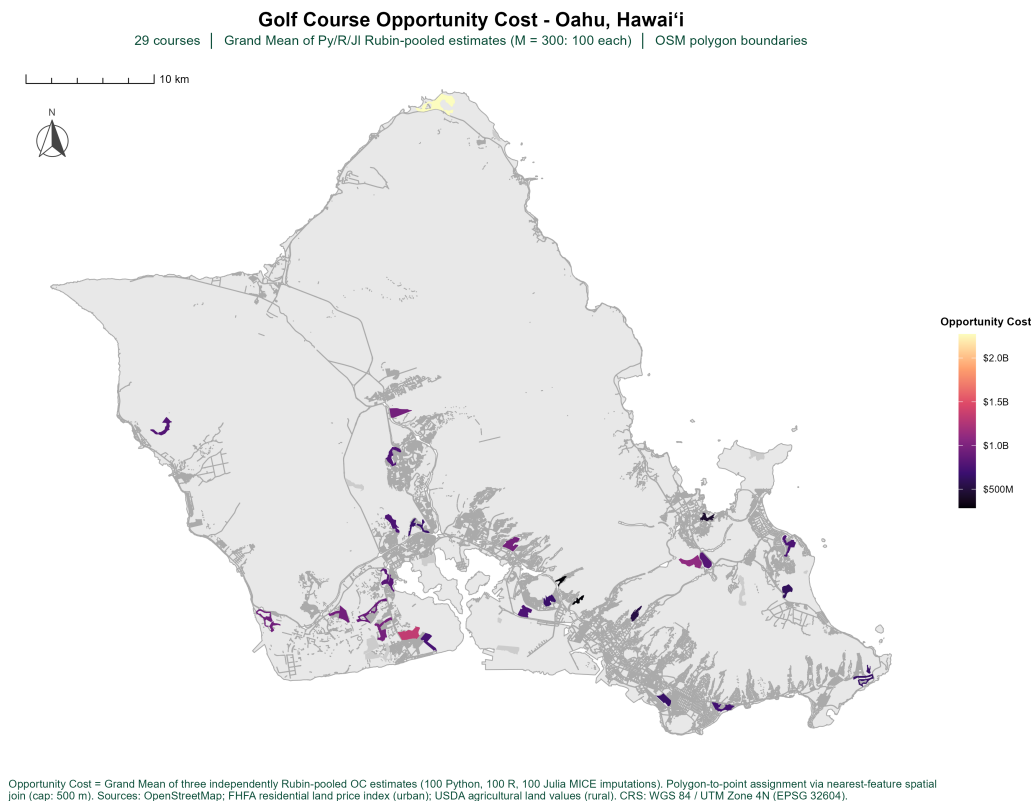


Figure 7: Per-course opportunity cost on Oahu under the Grand Mean of three independently Rubin-pooled tri-language estimates. OSM-derived golf course polygons are colored by computed opportunity cost (\$286M to \$2.27B). Polygon-to-point assignment uses a nearest-feature spatial join with a 500-meter cap; courses with coordinates exceeding this distance from candidate polygons are excluded.

A complementary observed-only version of this map (Figure 8) restricts the per-course opportunity cost to courses with directly measured OSM polygon acreage, excluding the imputed subset. The 29 observed-only courses produce an aggregate Oahu opportunity cost of \$25.1 billion, which is the observed-only floor that anchors the \$26.67 billion Grand Mean from above and matches the 83.5% national observed-to-MICE coverage ratio reported in Section 4.1.

<sup>2</sup>Four per-course values warrant brief context. *Moanalua Golf Club* (\$286.6M) renders at the bottom of the color scale because it is a small 9-hole course on 57.86 acres — a function of physical size, not a per-acre valuation anomaly. *Walter J. Nagorski Golf Course* renders near-zero because it sits on Schofield Barracks (F-1 federal/military preservation, outside the residential market the FHFA proxy values). *Hawaii Kai Golf Course* (\$645.9M Grand Mean) and *Mid-Pacific Country Club* (\$753.0M Grand Mean) anchor on the FHFA residential value of \$4,952,600 per acre for Honolulu County (FIPS 15003); the pipeline history that produced these values — including a cartographic boundary-simplification artifact corrected during development — is recorded in Appendix A.1.1.

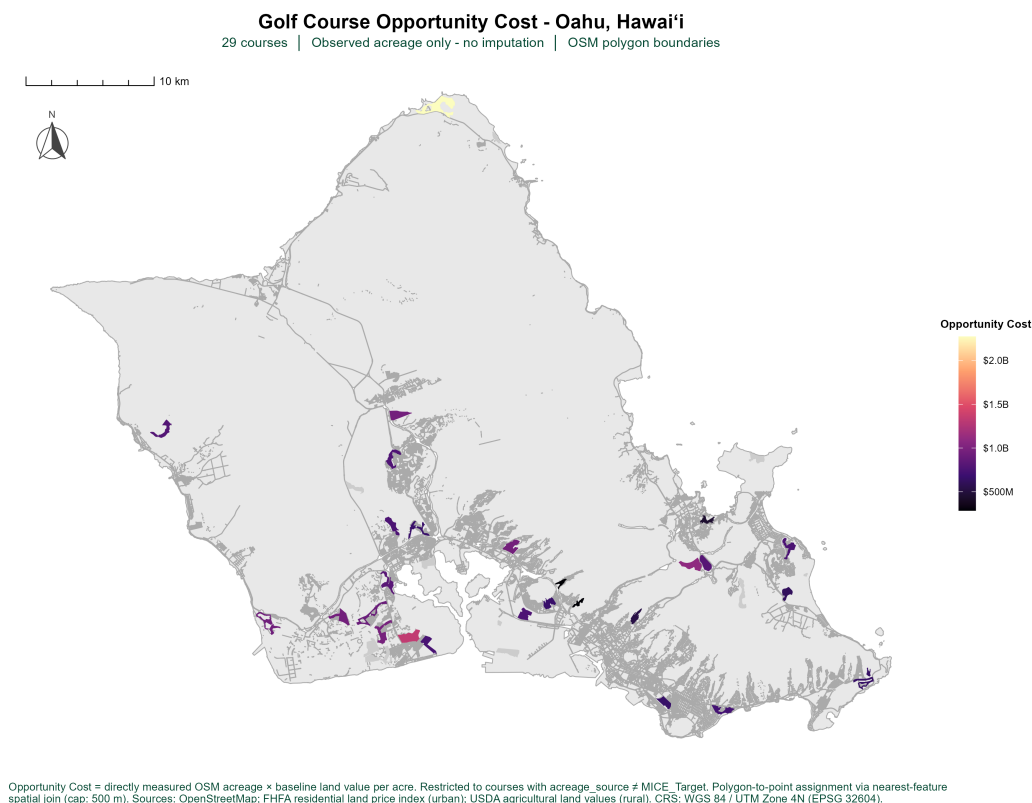
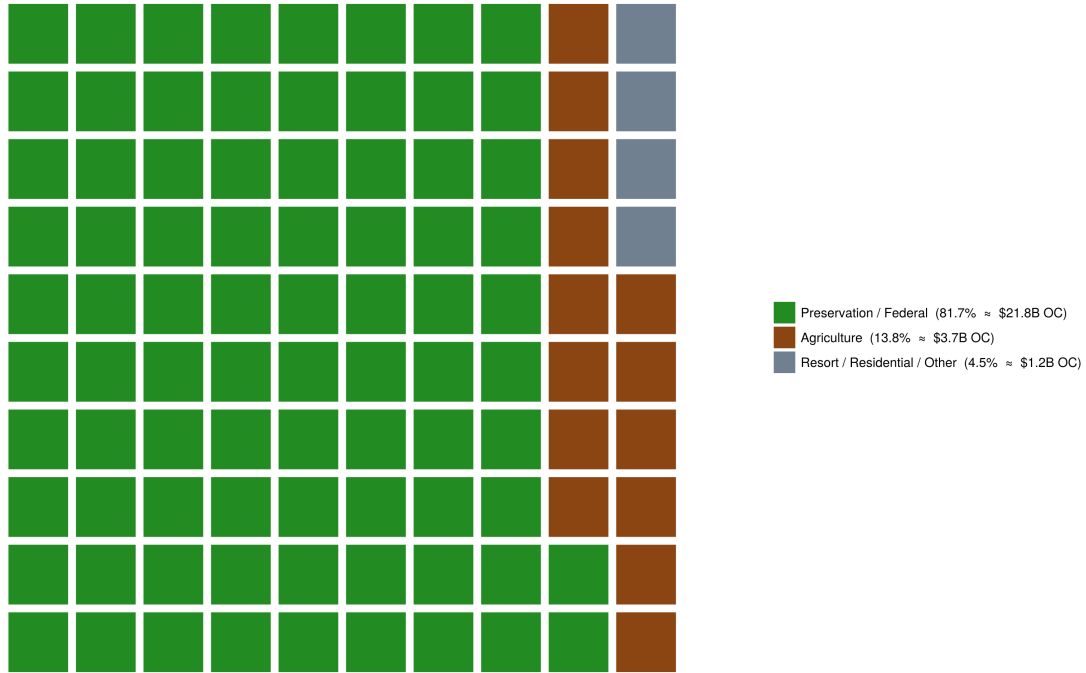


Figure 8: Per-course opportunity cost on Oahu under the observed-only specification (no imputation). The 29 courses with directly measured OSM polygon acreage are rendered; the color scale matches Figure 7 for direct visual comparison. Although visually similar to Figure 7, this map is constructed from the observed-only subset and excludes the four MICE-imputed Oahu courses, accounting for the slightly reduced aggregate value. The aggregate of \$25.1 billion across these 29 courses is the observed-only floor on the Oahu opportunity cost estimate.

#### 4.4.3 The Preservation Paradox

The most striking finding of the Hawaii micro-study emerges from the zoning intersection analysis. Each of the 1,072 golf TMK parcels was assigned its dominant zoning class via largest-overlap area against the Honolulu zoning layer, and the total acreage in each class was tabulated. The result, summarized in Figure 9, is that 81.7% of all Oahu golf land sits within Preservation (P-1, P-2) or Federal/Military (F-1) zones — areas where residential redevelopment is, under current zoning, legally prohibited.

**The Preservation Paradox - Zoning of Oahu Golf Land**  
 Grand Mean of Py / R / JI Rubin-pooled estimates | M=100 per language



Grand Mean total Oahu OC: \$26.67B (Phase 5b TMK-parcel-intersected Rubin-pooled estimates, M=100 per language; arithmetic mean of JI, Py, R). Tile proportions = land-area share by zoning category from Phase 5 parcel data.

Figure 9: The Preservation Paradox: the dominant zoning classification of Oahu golf land. Each tile represents 1% of the total Oahu golf footprint, sized by land-area share. The Grand Mean total Oahu opportunity cost of \$26.67B decomposes by zoning class as approximately \$21.8B in Preservation/Federal zones (legally restricted), \$3.7B in Agriculture zones, and \$1.2B in Resort/Residential/Other zones (legally permitted under most current zoning).

The dominant zoning classes by acreage are:

- P-2 (General Preservation): 3,209 acres (52.9% of Oahu golf footprint)
- F-1 (Federal/Military Preservation): 1,002 acres (16.5%)
- P-1 (Restricted Preservation): 745 acres (12.3%)
- AG-1 and AG-2 (Restricted and General Agriculture): 836 acres (13.8%)
- Resort, Residential (R-3.5, R-5, R-7.5, R-10, R-20), Country, and other classes: 274 acres (4.5%)

The economic implication is direct. Under the legally-permissible HBU framework defined in Section 2.3, the realizable opportunity cost on Oahu is bounded by the share of the footprint in zones that legally permit redevelopment. Approximately \$1.2 billion of the \$26.67 billion gross HBU estimate falls in zones (Resort, Residential, Country) where redevelopment to housing is legally compatible with current zoning; an additional \$3.7 billion falls in agricultural zones where redevelopment may be permissible subject to state Land Use Commission boundary amendment; and the remaining \$21.8 billion falls in zones where redevelopment is statutorily prohibited.

A related and equally striking finding is that golf occupies **25.4% of all Resort-zoned land on Oahu** — the highest penetration rate of any zone class — and 18.6% of all P-2 General Preservation land. The Resort statistic in particular indicates that golf and resort hotel development have co-evolved on Oahu to the point that golf is not an incidental amenity within resort zoning but rather the dominant land use within that zone.

#### 4.4.4 Synthesis of the Hawaii Findings

The Hawaii validation produces three results that materially qualify the national estimate. First, the aggregate Oahu opportunity cost of \$26.67 billion is consistent with the national-scale Honolulu County estimate and with Hawaii’s appearance in the top-five state list, lending credibility to the national pipeline’s geographic resolution. Second, the 81.7% Preservation/Federal share establishes that the unrestricted gross HBU framework substantially overstates the realizable opportunity cost in this market: the bulk of Oahu’s golf land is, under current zoning, not eligible for the residential redeployment that the model assumes. Third, the model-to-assessed pilot ratio reported in Section 4.4.1 characterizes the gap between gross HBU and going-concern current-use value in this Hawaii sample, providing a market-specific anchor whose generalization to other markets is reserved for future work. The economic potential is real, the magnitude is large, but the share of that potential that is unlocked by current law is small.

## 5 Discussion and Policy Implications

The empirical results presented in Section 4 produce a layered picture. At the national scale, the aggregate gross HBU opportunity cost of U.S. golf course land is approximately \$944 billion, with an observed-only floor of \$788 billion. At the Honolulu County scale, this figure decomposes into \$26.67 billion of theoretical opportunity cost, of which only roughly \$1.2 billion sits in zones where redevelopment to housing is currently legally permissible. This section interprets these findings in light of the literature reviewed in Section 1, examines the institutional architecture that prevents market clearing, and addresses the methodological limitations that bound the analytical claims.

### 5.1 The Magnitude of Spatial Misallocation

The \$944 billion figure represents the unrealized economic potential of approximately 2.30 million acres of recreational land in the United States. This footprint exceeds the combined land area of Delaware and Rhode Island (approximately 1.91 million acres) and is more than three times the current land footprint of all U.S. utility-scale solar generation (approximately 644,000 acres, derived from K. Peterson and C. Peterson 2024’s August 2024 figure of 107.4 GW utility-scale solar capacity in the Lower 48 states multiplied by SEIA’s central land-use intensity of 6 acres per megawatt; Solar Energy Industries Association n.d.). For policy reference, the U.S. National Renewable Energy Laboratory estimates that powering the entire U.S. electricity system from solar alone would require approximately 10 million acres (U.S. National Renewable Energy Laboratory 2024), meaning the current golf footprint is roughly 23% of the land area that a fully solar-powered United States would need. These three reference comparisons are summarized in Figure 10. The opportunity cost is concentrated geographically in the same housing markets that E. L. Glaeser et al. (2005), Davis, Larson, et al. (2021), and Gyourko et al. (2021) identify as the most regulatory-constrained — coastal California, the New York metropolitan region, southern Florida, and Honolulu County. This co-location is not coincidental. The same regulatory architecture that bids up residential land prices through restrictive zoning is also the architecture that makes redevelopment of existing

recreational land politically and legally difficult, producing the paradox that opportunity cost is highest precisely where it is least likely to be realized.

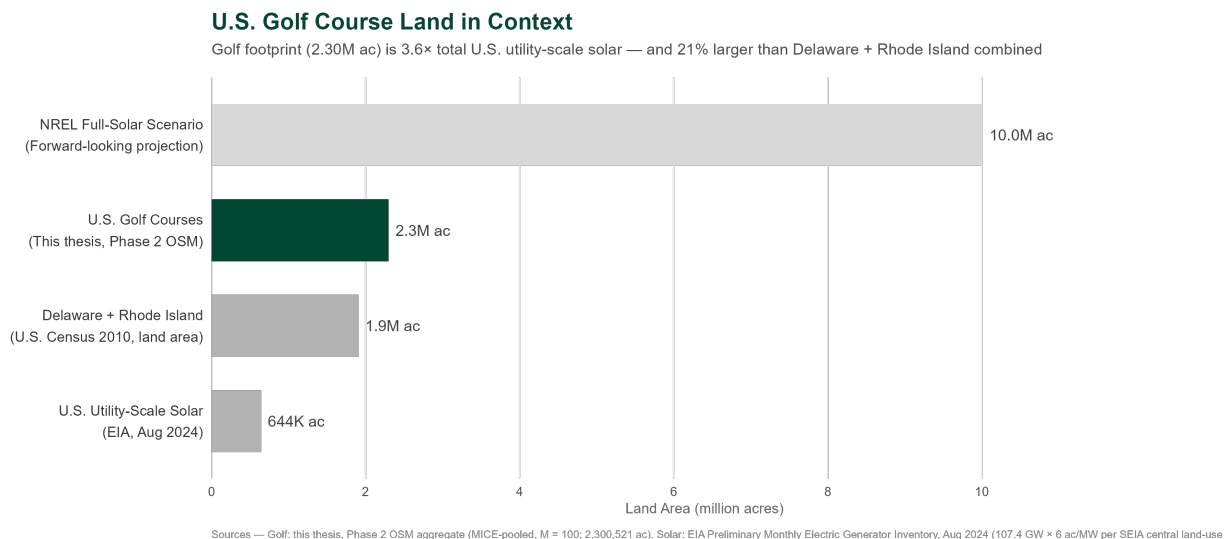


Figure 10: Magnitude of U.S. golf land against reference categories. U.S. golf courses (2.30M ac, this thesis) exceed the combined land area of Delaware and Rhode Island (1.91M ac total land area per U.S. Census Bureau decennial enumeration; the value is stable to within 0.1% across the 2010 and 2020 decennial vintages; U.S. Census Bureau 2021a, U.S. Census Bureau 2021b) and currently exceed total U.S. utility-scale solar deployment by approximately 3.6× (0.644M ac; K. Peterson and C. Peterson 2024, Solar Energy Industries Association n.d.). NREL’s forward-looking 100%-solar scenario (10M ac; U.S. National Renewable Energy Laboratory 2024) is included as a scale reference for what fully decarbonized electricity generation would require.

The Lorenz curve of per-course opportunity cost (Figure 3) reinforces this point. The aggregate \$944 billion figure is dominated by a relatively small subset of high-value urban courses; the median course contributes a fraction of one percent of the total. This distributional asymmetry implies that policy interventions targeting the highest-cost courses — particularly in coastal California and the urban Northeast — could capture a disproportionate share of the welfare gains without requiring nationwide regulatory change.

## 5.2 Coasian Frictions and the Allocation of Rights

The persistence of golf course land use in markets where the residential value of the land vastly exceeds its recreational value is, on a naive Coasian reading, a puzzle. With sufficient gains from trade, parties should bargain to the efficient allocation. The actual outcome — decades of persistence at low-density recreational use despite escalating opportunity costs — requires an explanation grounded in transaction costs and rights structure.

The Coase (2013) framework is most often invoked to argue that initial rights assignments are irrelevant under zero transaction costs. The empirically relevant case, however, is positive transaction costs, in which the initial rights assignment determines the equilibrium allocation. The current legal architecture grants a powerful set of veto rights over redevelopment to incumbent homeowners and golf course members, exercised through three distinct channels: (i) municipal zoning ordinances

that prohibit high-density redevelopment of recreationally classified land (Fischel 2015); (ii) restrictive private covenants that mandate the preservation of golf course use in perpetuity, modifiable only by supermajority of beneficiaries (Ellickson 2022); and (iii) the political capitalization of the open-space amenity into surrounding property values, which mobilizes intense local opposition to any redevelopment proposal (Cheshire and Sheppard 2002).

The cumulative effect of these veto channels is that the bargaining surface required for a Coasian transaction is enormous. Korngold (2024) proposes that state legislatures address this by voiding stale single-family-equivalent covenants under public policy doctrine, but cautions that such voiding may itself constitute a regulatory taking under Fifth Amendment doctrine, requiring just compensation. The economic question this raises is sharp: at \$944 billion in aggregate gross HBU counterfactual nationally, and roughly \$25.5 billion of the \$26.67 billion Oahu total locked in statutorily restricted zoning classes (Preservation, Federal/Military, and Agriculture combined) where redevelopment is currently prohibited, what compensation regime would render the rights reassignment efficient? The literature does not yet offer a confident answer.

### 5.3 The Validation Gap: What 1.33× Tells Us

The Phase 5a model-to-assessed ratio of 1.33× is, taken in isolation, a small-sample finding (six courses, all in Hawaii). But the structure of the ratio across the sample is informative. The gap is narrowest for urban Honolulu courses (1.16–1.23×) and widest for rural Big Island courses (1.49–1.69×). This pattern is consistent with two complementary interpretations.

First, the FHFA county-level residential land price index is a more accurate proxy for what assessors are doing in urban markets, where assessors themselves are more likely to use comparable-sales appraisal methods anchored in nearby residential transactions. In rural markets, assessors more often rely on use-value or income-based methods, which produce systematically lower assessments for golf properties classified as recreational, widening the gap between model HBU and assessed value.

Second, the empirical gap is itself a useful characterization of how the gross HBU estimate relates to going-concern current-use value in this market. A model-to-assessed ratio of 1.33× implies that, on average across the six-course Hawaii sample, the gross HBU estimate exceeds the assessed value by approximately one-third. This is a Hawaii-specific finding, not a national deflator: the sample is small, geographically concentrated, and embedded in an assessment regime whose practices (Honolulu’s parcel-level cadastral system, Hawaii’s Act 273 agricultural classification framework, the state’s dual rate structure for owner-occupied versus non-owner-occupied parcels) differ substantively from the mainland. Applying the 1.33× ratio as a uniform deflator to the national \$944B figure would assume a level of cross-market consistency in assessment methodology that the present six-course pilot cannot support. The defensible interpretation is narrower: in at least one regulatory-intensive U.S. market, gross HBU exceeds going-concern current-use value by roughly one-third, and the size of that gap is itself an economically meaningful object. Extending the validation to other markets — the Bay Area, southern Florida, Boston-area suburbs — is the natural future-work pathway for converting this market-specific characterization into a national net-realizable estimate.

### 5.4 The Zoning-Bounded HBU

The most consequential finding for policy interpretation is the Preservation Paradox documented in Figure 9: 81.7% of Oahu golf land sits in Preservation or Federal/Military zones, where residential

redevelopment is statutorily prohibited. Translating the Oahu zoning composition into the legally-permissible HBU framework defined in Section 2.3, the realizable opportunity cost on Oahu under current zoning is approximately \$1.2 billion — 4.5% of the gross \$26.67 billion estimate.

This is not a minor caveat. It is the central tension of the policy interpretation. The economic case for redeveloping golf course land in housing-constrained markets is large, but that case applies only to land that current zoning permits to be redeveloped. For the bulk of Oahu’s golf footprint, unlocking the gross HBU value would require not just willing buyers and sellers but state-level Land Use Commission boundary amendments or federal land disposal — processes that operate on timescales of years to decades, with substantial political and constitutional friction.

But the regulatory friction is only half of the policy argument. The substantive case against treating these zones as immovable rests on an inconsistency that the data make difficult to ignore. The Preservation classification is grounded in the protection of ecological function: native vegetation, watershed integrity, wildlife habitat, viewshed continuity. A golf course preserves none of these. Constructing one requires clearing existing vegetation, regrading topography, installing irrigation drawing from the same watersheds the classification ostensibly defends, and replacing native plant communities with chemically maintained monoculture turf. To the extent these parcels supported native ecological function prior to course construction, that function has already been destroyed; the present analysis cannot directly verify pre-construction land cover for each parcel, but the construction process itself (clearing, regrading, irrigation, monoculture establishment) is incompatible with the ecological premise of the Preservation designation. The land is then maintained in a state that requires continuous chemical and water inputs to remain in its degraded configuration. Calling this “preservation” is a category error preserved by the absence of any mechanism to update the classification when the on-the-ground use diverges from the regulatory premise.

One legitimate function does survive this critique. A 150-acre golf course functions, in civil-engineering terms, as a large privately-funded permeable surface — a stormwater retention asset that absorbs runoff that would otherwise need to be managed through municipally-funded gray infrastructure. In coastal jurisdictions where flooding and drainage are first-order planning concerns (Oahu, southern Florida, much of the Gulf Coast), the Preservation classification can therefore reflect a dual mandate: ecological preservation as the formal premise, and stormwater management as a complementary hydrological utility. The hydrological function is real, and any rigorous HBU counterfactual for a golf parcel in a high-runoff watershed should internalize the replacement cost of the drainage capacity that paving the course would eliminate — retention basins, storm drains, permeable-pavement infrastructure, and the operating cost of maintaining them. For a typical 150-acre course in a Honolulu coastal watershed, this replacement cost is plausibly in the low tens of millions of dollars per course, a non-trivial but bounded offset against the per-course HBU values reported in Section 4.4. The hydrological function does not refute the category-error argument: the same retention capacity could be supplied by smaller engineered detention basins, by permeable-pavement standards in the redeveloped parcel, or by genuinely preserved native riparian buffer at substantially lower acreage cost. But it does qualify the categorical claim that golf courses serve no Preservation function. The honest position is that they serve a hydrological function adequately, and that the Preservation classification could be replaced by an explicit drainage-utility classification without losing the substantive justification — which would in turn open the door to compensated transitions to higher-and-better uses that internalize the drainage replacement cost.

This produces an asymmetry that bears directly on housing policy. If a jurisdiction is willing to permit a golf course on land classified as Preservation, then the operative meaning of that classification cannot be “no human modification of the landscape.” It must be something narrower:

“no human modification that produces tax revenue at residential density,” or “no human modification that brings in new residents.” Stated that way, the classification is no longer a defense of ecology but a defense of an incumbent equilibrium. The 81.7% Preservation-zoned share of Oahu’s golf footprint is not a measurement of land that nature has been protected from; it is a measurement of land where one particular form of human modification has been permitted and another, more egalitarian form, has been blocked. The economic cost of that asymmetry is the \$21.8 billion of opportunity cost the unrestricted HBU framework identifies in the Preservation/Federal tier.

A natural extension of this analysis would replicate the Phase 5b zoning intersection across other high-value markets where the Hawaii pattern is plausible. If the same Preservation/Restricted-zoning pattern holds for golf land in coastal California, southern Florida, or the New York metro — and there is strong prior reason to expect that it does, given the historical use of golf courses as anchors for environmentally restrictive master-planned communities — then the realizable national opportunity cost may be substantially smaller than the \$944 billion gross figure suggests. Quantifying that difference, and characterizing the share of national restricted-zone golf land where the “preservation” classification has already been functionally invalidated by the existing course, are among the most important directions for follow-on work.

## 5.5 Methodological Limitations

Three limitations are sufficiently important to warrant explicit treatment.

*Marginal versus bulk valuation.* The FHFA county-level residential land price reflects the per-acre price of finished, subdivided residential lots — land already entitled, served by infrastructure, and ready for construction. Applying this per-acre price to the full undeveloped acreage of a golf course over-states the value a developer could realize from a bulk acquisition, because the bulk parcel must be subdivided, infrastructure must be installed, and lots must be absorbed over a multi-year timeline. Standard appraisal practice applies a subdivision development discount in this context, ranging from approximately 25% to 60% of gross retail lot value depending on market conditions and absorption assumptions (Appraisal Institute 2020). The \$944 billion gross HBU figure should therefore be interpreted as an upper bound on bulk-liquidation value. Applying the full 25–60% range yields a developer-economics estimate of \$378–\$708 billion; a 35% bulk discount, the conventional central rate adopted in appraisal practice for raw-land subdivisions of this scale, produces approximately \$614 billion as a central estimate. The thesis adopts the 35% conventional rate as the headline value for the bulk-adjusted figure, but readers should treat the surrounding \$378–\$708B range as the defensible bound; the central estimate’s exact level is less analytically robust than the order-of-magnitude span. Further downward correction — for example, to adjust for the going-concern current-use value that municipal assessments capture — would require market-by-market anchoring of the kind initiated for Hawaii in Section 4.4.1 and discussed in Section 5.3.

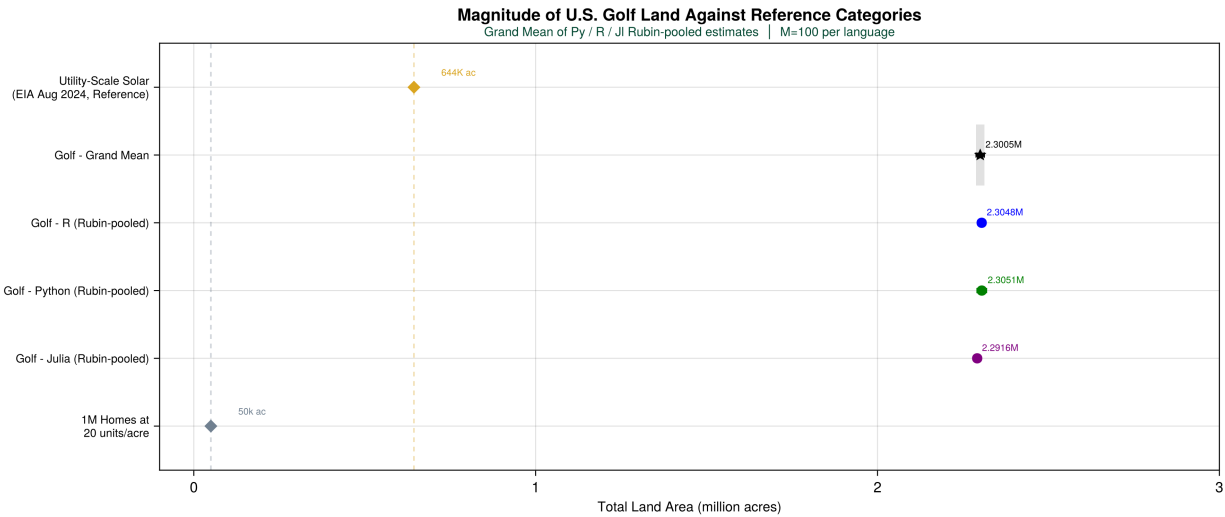
*$V_{Current}$  measurement.* As discussed in Section 2.2.1, the national-scale data sources do not provide course-level capitalized recreational income, so  $V_{Current}$  cannot be computed directly for each parcel. The Hawaii pilot validation in Section 4.4.1 provides one market-specific anchoring strategy; alternatives include capitalizing industry-average golf course operating income at market cap rates, or extracting course-level transaction values from the National Golf Foundation’s transaction database. None of these alternatives was implemented in the present study, and each would carry its own measurement assumptions.

*Amenity capitalization and adjacent property externalities.* Frank F Limehouse et al. (2012) and Cho et al. (2009) document that golf courses generate capitalized amenity premiums for adja-

cent residential properties, and that these premiums (while declining over time) are economically meaningful. Redeveloping a golf course destroys these premiums, imposing a real welfare cost on incumbent homeowners that is not netted out in the present analysis. The unrestricted HBU figure should therefore be understood as a measure of the gross spatial misallocation, not as a measure of the net welfare gain from redevelopment, which would require subtracting the capitalized amenity loss to neighbors. Cho et al. (2009)’s finding that the amenity premium has declined and become substitutable with other forms of open space suggests that this capitalized loss is a small share of the gross opportunity cost in most markets, but quantifying it directly is an important extension.

## 5.6 Renewable Energy as an Alternative Counterfactual

Given the severe legal frictions associated with residential redevelopment, policymakers might reasonably consider non-residential HBU counterfactuals where regulatory and political resistance is lower. Weinand et al. (2025) identify utility-scale renewable energy infrastructure — particularly photovoltaic solar — as a synergistic alternative use for golf course land. The land footprint of U.S. golf courses (approximately 2.30 million acres) is more than three times the current installed land footprint of utility-scale solar generation in the United States (approximately 644,000 acres; K. Peterson and C. Peterson 2024, Solar Energy Industries Association n.d.), and represents roughly 23% of the land that the National Renewable Energy Laboratory estimates would be required to meet the entire U.S. electricity demand from solar alone (Figure 11). Repurposing even a small fraction of underutilized golf land for solar generation would directly contribute to addressing the global “land squeeze” on renewable energy infrastructure without triggering the density and traffic concerns that mobilize NIMBY opposition to housing.



Solar reference: EIA Aug 2024, 107.4 GW × 6 ac/MW per SEIA (Peterson & Peterson 2024; SEIA n.d.). Housing reference: 1M units at 20 units/acre = 50,000 ac (illustrative density assumption within HUD’s 15–30 units/acre high-density range). Golf acreage from MICE-pooled OSM polygon data (J/Py: osm\_acreage; R: final\_acreage). Grand Mean = arithmetic mean of three independently Rubin-pooled national totals.

Figure 11: Counterfactual area comparison: U.S. golf course footprint relative to current U.S. utility-scale solar (approximately 644,000 acres; K. Peterson and C. Peterson 2024, Solar Energy Industries Association n.d.) and an illustrative high-density housing reference (1 million units at 20 units/acre  $\approx$  50,000 acres, within HUD’s 15–30 units/acre high-density range). The three per-language Rubin-pooled golf estimates (Julia, Python, R) cluster tightly around the Grand Mean of 2.30 million acres, providing visual confirmation of cross-language convergence.

Integrating  $V_{Renewable}$  formally into the HBU framework defined in Section 2.2 — such that  $V_{OC,i} = \max(V_{Res,i}, V_{Ag,i}, V_{Renewable,i}) - V_{Current,i}$  — is a natural extension and is particularly relevant in arid urban environments where the ecological cost of maintaining turf is high and the marginal value of renewable generation capacity is large. The present study did not implement this extension, but the dataset and pipeline are structured to support it.

## 6 Conclusion

This thesis quantifies the aggregate opportunity cost of U.S. golf course land at two scales, producing a layered cascade of three figures that answer three distinct policy questions.

*The gross magnitude.* At the national scale, applying the appraisal principle of Highest and Best Use to a 16,297-course spatial dataset, with 28.8% missing acreage resolved through Multiple Imputation by Chained Equations ( $M = 100$  per language across three independent computational stacks) and pooled via Rubin’s Rules, the **gross HBU opportunity cost is approximately \$944 billion**. The observed-only floor, computed without imputation on the 11,605 courses with directly measured OSM polygon acreage, is \$788 billion. This is the answer to the question: “How much residential land value, at finished-lot per-acre prices, sits under U.S. golf courses today?” It is a measurement of the gross spatial misallocation under an unrestricted counterfactual.

*The developer-economics figure.* Adjusting the gross national estimate for the standard land-residual bulk discount of 25–60% — reflecting the gap between finished-lot per-acre prices and the price a developer would pay for raw acreage net of infrastructure costs, absorption risk, and developer profit (Appraisal Institute 2020) — yields a developer-economics range of **\$378–\$708 billion**, with a central estimate of **approximately \$614 billion** at the conventional 35% rate. This is the answer to the question: “What is the magnitude of the residual land value that could plausibly be unlocked through redevelopment, if regulatory frictions could be dissolved?” The range, rather than the central estimate, is the analytically defensible quantity: the discount rate is market-dependent, and the headline central value should be read as an order-of-magnitude anchor within the broader band. Like the gross figure, this range is computed before deduction of capitalized current recreational value ( $V_{Current}$ ); incorporating going-concern value would reduce both endpoints by an amount the Hawaii pilot (Section 4.4.1) indicates is roughly 17–35% depending on local assessment regime, with urban markets at the low end and rural markets at the high end of that band.

*The legally-realizable figure.* At the Honolulu County scale, integration with parcel-level cadastral and zoning data produces the binding constraint that the gross magnitude obscures. The aggregate Oahu opportunity cost is \$26.67 billion across 1,072 unique TMK parcels — but 81.7% of this footprint sits in Preservation or Federal/Military zones where residential redevelopment is statutorily prohibited, leaving only **approximately \$1.2 billion of the gross estimate as legally realizable** under the existing zoning regime. This is the answer to the question: “How much of the gross magnitude could be unlocked under current law, with no regulatory reform?” Generalizing Oahu’s 4.5% legally-realizable share to the national aggregate would yield an indicative figure of roughly \$42 billion. Because the Oahu Preservation Paradox reflects island-specific regulatory features (volcanic topography, federal-military footprint, Hawaii Land Use Commission framework) that do not hold uniformly on the mainland, a defensible range bracketing the legally-realizable share at 4.5–15% across U.S. coastal markets yields a broader band of approximately **\$42–\$140 billion**. The point estimate is illustrative only; mainland markets with less restrictive zoning would push the realizable share substantially higher than the Oahu figure. A separate Hawaii pilot finding — the model-to-assessed gap reported in Section 4.4.1 — characterizes the relationship between

gross HBU and going-concern assessed value in this specific market, and is not extrapolated to the national aggregate.

*The three figures are complementary, not competing.* \$944B is the gross unrealized counterfactual; the \$378–\$708B range (central estimate \$614B at the conventional 35% rate) is the developer-economics counterfactual net of bulk discount; \$1.2B (Oahu) and an illustrative range of \$42–\$140B (national extrapolation) are the legally-realizable counterfactuals under current zoning. Each figure measures something different. The headline result is not any single number but the structure of the cascade: the magnitude is large at the gross level, smaller but still large at the bulk-adjusted level, and dramatically constrained at the legally-realizable level. Unlocking even a meaningful fraction of the gap between the second and third figures would require simultaneous reform on multiple legal fronts — Land Use Commission boundary amendments, voiding of stale private covenants under public policy doctrine, and the resolution of associated Fifth Amendment takings claims.

The 60:1 urban-to-rural per-acre ratio in the pooled regression decomposition, and the spatial concentration of value (California, Florida, New York, Texas, and Hawaii together account for over half of the aggregate), confirm that the cascade structure is itself geographically uneven — and that the gross-to-realizable gap is largest precisely in the markets where housing supply pressure is most acute.

Three directions for future research follow naturally from these findings. First, the Phase 5b zoning intersection methodology should be replicated in other high-value coastal markets — particularly the Bay Area, Los Angeles, southern Florida, and the New York metro — to test whether the Preservation Paradox documented on Oahu generalizes, and to produce a national legally-permissible HBU estimate alongside the unrestricted figure. Second,  $V_{Current}$  should be measured directly using transaction data from the National Golf Foundation or income-capitalization methods, refining the  $1.33\times$  Hawaii anchor into a market-specific net opportunity cost. Third,  $V_{Renewable}$  should be integrated formally into the HBU framework as a third counterfactual, allowing the analysis to identify parcels where solar or wind generation produces higher residual value than residential redevelopment.

Three further extensions are worth flagging explicitly; two are described in the present paragraph and the third in the paragraph that follows. A reasonable critique of the present framework is that in many rural markets characterized by agricultural saturation, the current recreational use of a golf course may itself represent the local Highest and Best Use — particularly where surrounding agricultural land already exceeds regional demand. A rural-HBU inversion analysis would partition the dataset by regional ag-land saturation indices (USDA Cropland Data Layer, Census of Agriculture acreage trends) and recompute the aggregate excluding courses where rural HBU is plausibly already met by current use. Separately, the present analysis is anchored to 2022 FHFA residential land prices and 2022 USDA agricultural land values to maintain temporal consistency across the national dataset. A targeted temporal sensitivity check using 2025 FHFA data for the urban subset would test whether the headline figure is robust to recent residential land price movements, particularly in the high-value coastal markets where the bulk of the aggregate is concentrated.

A third extension addresses the intra-county constancy assumption flagged in Section 3 and quantified for the Honolulu case study in Section 4.4. The FHFA dataset publishes its primary land-price series at the county level, which is the resolution the present analysis uses. Where the dataset offers a finer spatial resolution—tract-level estimates from Davis, Larson, et al. (2021)—a follow-on pipeline could replace each course’s FHFA value with the tract-level price for the tract the course’s coordinate falls in, capturing the sub-county price gradients that the Turtle Bay and Makaha Valley footnotes in Section 4.4 document. The natural statistical form for this extension is a multilevel

imputation model that nests courses within tracts within counties, permitting the imputed BVPA distribution to draw simultaneously on tract-level local variation and county-level fallback for tracts where FHFA coverage is sparse. The present thesis offers the dataset, the pipeline, and the validation logic to support all six extensions identified across this section.

A residual cross-language asymmetry remains in the spatial-join precision of the three Phase 1 pipelines. After switching to the full TIGER/Line county boundary file, Python's `geopandas.sjoin` resolves all 16,297 courses to a valid FIPS code; Julia's `ArchGDAL.intersects` resolves all but three. The three residual courses are *Sewailo Golf Club* in Arizona, situated on the Pascua Yaqui Indian Reservation enclave where the tribal-land polygon sits inside the county boundary; *Normandy Shores Golf Course* in Florida, on a small island in Biscayne Bay; and *Turtle Creek Golf Club* in Florida, on a coastal sandbar feature. In each case, the course's coordinate falls on or near a polygon vertex or sub-county-feature edge where Julia's floating-point intersection test returns false. The aggregate impact is 0.018% of the Julia dataset (three of 16,292 courses), with their BVPA values handled via MICE imputation in the Julia sub-pool only. The other two language pipelines produce correct FHFA assignments for these courses, and the cross-language Grand Mean is unaffected at the third significant figure. Future work could investigate whether a tolerance adjustment in Julia's ArchGDAL spatial-join call would resolve these three edge cases without affecting the rest of the pipeline.

The \$944 billion figure is not a forecast of how much land value will be unlocked, nor a recommendation that every golf course in the country be converted to housing or renewables. It is a measurement: a quantification of the spatial misallocation embedded in the current allocation of recreational land in the United States. What policy makes of that measurement is a separate question, but the measurement itself — bounded, validated, and reproduced across three computational stacks — is the contribution this thesis offers to the literature on land use, housing supply, and regulatory economics.

## A Appendix A: Pipeline Methodology

This appendix provides the full technical narrative of the national pipeline (Phases 1–4) and the Honolulu County pipeline (Phase 5), expanding on the brief overview provided in Section 3. The material is presented in the order in which it is executed: from raw input data through county assignment, valuation, polygon extraction, imputation, aggregation, and cross-language validation.

### A.1 National Pipeline: Data Sources and County Assignment

The primary dataset consists of 16,297 golf course observations across the United States. Spatial point coordinates (latitude and longitude), mapped to the WGS84 coordinate reference system, were spatially joined to U.S. Census Bureau county shapefiles (2022 vintage) to assign a 5-digit Federal Information Processing Standard (FIPS) code to each course. The initial pipeline implementation used cartographic boundary files (20-meter resolution) for the county join; this produced 34 FIPS-NA courses (0.21% of the dataset) where the course coordinate fell outside the simplified boundary. Switching to the full TIGER/Line boundary files resolved all 34 cases in the Python and R implementations and all but three in the Julia implementation (see Section 6 for the residual Julia edge cases). A separate FIPS zero-padding bug was identified and corrected during pipeline development: county FIPS codes stored as integers lose their leading zeros (e.g., 1001 instead of "01001"), causing silent join mismatches with the USDA and FHFA datasets. The fix enforces 5-digit zero-padded string formatting before any join operation, and was applied identically across all three language implementations.

#### A.1.1 FIPS Resolution for the Hawaii Kai and Mid-Pacific Cases

Five of the 34 nationally affected FIPS-NA courses fell in Hawaii, of which two are referenced in the per-course Oahu opportunity cost discussion in Section 4.4.2: *Hawaii Kai Golf Course* and *Mid-Pacific Country Club*. Both courses sit on coastal Honolulu coordinates that fell narrowly outside the 1:20,000,000 cartographic boundary used in the initial pipeline implementation — 54.7 meters and 140.4 meters, respectively, from the nearest `cb=20m` county polygon edge. The remaining three Hawaii cases (Kahili Golf Course, King Kamehameha Golf Club, and Kapalua Golf Club, all on Maui) were also resolved by the boundary file switch.

Under the original `cb=20m` boundary file, the FHFA county join failed for both Honolulu courses, leaving their baseline value per acre to be imputed by MICE from the broader national training distribution. The imputation drew a mix of urban and rural values, producing a marginally low Python sub-pool estimate of \$296.6M for Hawaii Kai (against the canonical \$645.9M Grand Mean after the boundary fix). The same mechanism affected Mid-Pacific, though with a smaller magnitude because the distance to the boundary was shorter and the imputation draws clustered closer to the eventual canonical value.

The boundary-file diagnostic was performed by re-running the spatial join under three boundary variants: `cb=TRUE, resolution="20m"` (the original failed configuration), `cb=TRUE, resolution="5m"` (higher-resolution cartographic), and `cb=FALSE` (full TIGER/Line, no simplification). The latter two resolved both Honolulu courses; the `cb=5m` variant resolved Hawaii Kai and Mid-Pacific but not Kahili or King Kamehameha. The full TIGER/Line variant resolved all five Hawaii cases. The canonical pipeline therefore uses the `cb=FALSE` configuration, which trades a slightly larger file size for guaranteed FIPS resolution at the coastline.

The post-fix Grand Mean values are constant across all 300 imputed datasets in all three language

sub-pools, because the FHFA \$4,952,600 per acre value for FIPS 15003 is no longer being imputed — it is being directly joined. This is the methodological warrant for treating these two courses as having near-zero per-course uncertainty in the Phase 4 econometric model, despite their location at the upper end of the Oahu per-course OC distribution.

Two distinct 2022 land valuation datasets were integrated to establish the HBU counterfactuals:

1. **Residential Land Values ( $V_{Res}$ ):** Sourced from the Federal Housing Finance Agency (FHFA) county-level panel data (Davis, Larson, et al. 2021), joined via 5-digit FIPS codes.
2. **Agricultural Land Values ( $V_{Ag}$ ):** Sourced from the United States Department of Agriculture (USDA) National Agricultural Statistics Service county-level survey data, joined via the same FIPS keys.

## A.2 Hybrid Valuation Algorithm

Because land supply elasticity and the relevant HBU vary drastically across geographic submarkets (Davis and Palumbo 2008), applying a uniform residential counterfactual nationally would systematically overestimate opportunity costs in inelastic regions. To address this, a hybrid valuation algorithm assigns each course to either the residential or the agricultural counterfactual based on the 2023 USDA Rural-Urban Continuum Codes (RUCC):

- **Urban classification (RUCC 1, 2, 3):** Courses in metropolitan counties were assigned the FHFA county-level residential price per acre as  $V_{Res}$ .
- **Rural classification (RUCC 4 through 9):** Courses in non-metropolitan counties were assigned the USDA county-level agricultural price per acre as  $V_{Ag}$ .

The course-level Baseline Value Per Acre is then  $V_{Res}$  if the county is classified urban,  $V_{Ag}$  if rural. This procedure produces a defensible per-acre proxy for 93.3% of the courses in the dataset.

## A.3 Polygon Acreage Extraction

Course acreage is not directly reported in the source dataset. To compute it, golf course boundary polygons were extracted from a national-scale OpenStreetMap (OSM) PBF dump (approximately 11 GB) using a memory-efficient streaming parser. Polygons were reprojected to EPSG:5070 (NAD83 / Conus Albers) for planar area calculation, and filtered to drop polygons below 5 acres (likely fragments or individual greens) or above 1,500 acres (likely mega-resort blobs spanning multiple distinct courses). A two-pass spatial join then matched these polygons to the course list: a primary `intersects` join, followed by a fallback `nearest` join with a 500-meter cap for courses whose listed coordinates fall outside the mapped fairway (typically because the coordinate refers to a clubhouse or parking lot). This process produced OSM-derived acreage for 71.2% of the courses (11,605 of 16,297). For 28.8% (4,687 courses), no OSM polygon could be matched and acreage was resolved through Multiple Imputation, described below. The residual 5 courses (~0.03%) were assigned acreage via a fallback against the U.S. Census `tigris` “Area Landmarks” dataset, which provided limited additional coverage for courses with no OSM polygon and an unambiguous Census landmark match; their per-course contribution to the aggregate is negligible.

## A.4 Multiple Imputation by Chained Equations

To handle the 28.8% of courses missing OSM-derived acreage, the analysis applies Multiple Imputation by Chained Equations (MICE), following the Bayesian framework established by Rubin (1987)

and implemented per the operational guidance of Buuren (2018). For each language implementation,  $M = 100$  complete datasets were generated, yielding 300 imputed datasets in total. Random Forest-based imputation methods were used (`miceforest` in Python, `mice` with `method='rf'` in R, custom implementation in Julia) to handle non-linear spatial relationships and prevent impossible negative-value predictions, a known risk with linear imputers (Buuren 2018). The predictor matrix included `Holes`, `Ownership_Type`, `county_type`, `Longitude`, and `Latitude`.

The choice of  $M = 100$  exceeds the conventional minimum of  $M = 5$  recommended by Rubin (1987) by an order of magnitude. This larger value was chosen to produce stable Rubin-pooled estimates for the aggregate national valuation (a sum-of-products of imputed values) and to enable direct visualization of the convergence of the imputed acreage distribution toward the observed distribution as  $M$  increases.

### MICE Convergence and Imputation Stability

The choice of  $M = 100$  imputations per language can be motivated visually by examining the convergence of the imputed acreage distribution toward the observed distribution as  $M$  increases. Figure 12 shows the density of imputed acreage values at  $M = 100$  per language, plotted against the empirical density of the observed parcels. The three language-specific imputation distributions (Python, R, Julia) overlap closely with each other and with the observed density across the range of plausible course sizes (10–1,000 acres on a  $\log_{10}$  scale).

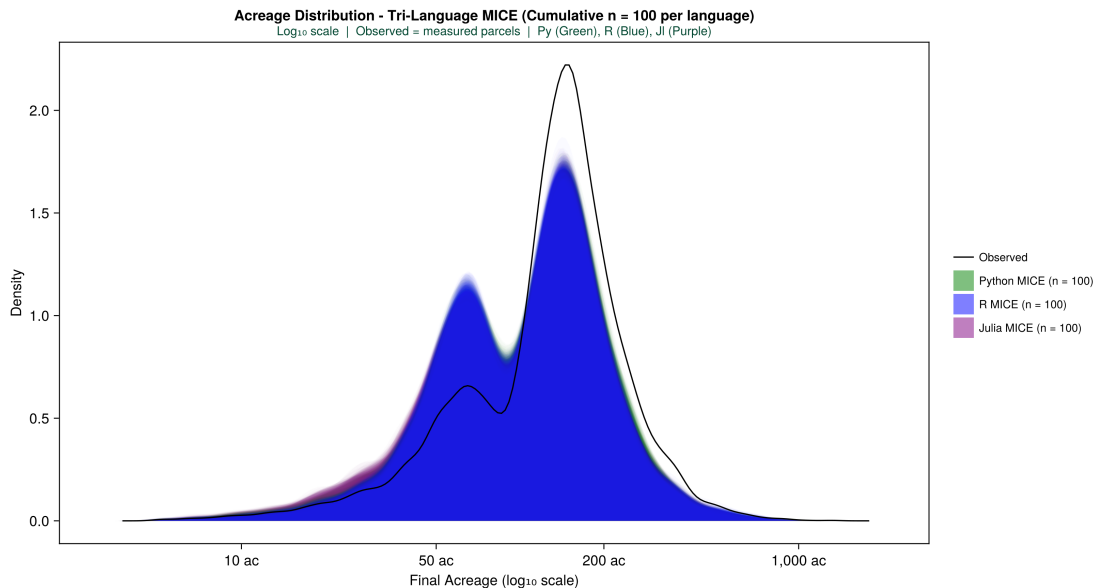


Figure 12: Acreage distribution from MICE imputation (cumulative  $M = 100$  per language) compared against the observed parcel density. The close overlap of all three language implementations with the observed distribution confirms that the imputation model is reproducing the structural shape of the acreage distribution rather than collapsing onto its mean.

### A.5 Tri-Language Cross-Validation

The entire pipeline — data ingestion, FIPS join, hybrid valuation, OSM polygon extraction, MICE imputation, and aggregation — is implemented in three independent computational stacks: Python

(`geopandas`, `miceforest`, `statsmodels`), R (`sf`, `tigris`, `mice`, `lm`), and Julia (`ArchGDAL`, custom imputation, `GLM.jl`). Each stack reads identical source files and produces independently named output files (`Py_*`, `R_*`, `Jl_*`). Statistical parity across the three implementations is checked at every phase, and known asymmetries (such as R’s use of `final_acreage` versus Python and Julia’s use of `osm_acreage` due to a documented Phase 2 difference in fallback behavior) are recorded explicitly. Any divergence in aggregate estimates across the three languages signals a pipeline defect rather than a substantive empirical finding, and triggers a debugging pass before the result is reported.

## A.6 Parameter Pooling and Rubin’s Rules

All statistical estimates are derived by performing the analysis independently across the  $M = 100$  imputed datasets per language and then pooling the results using Rubin’s Rules (Rubin 1987; Buuren 2018). The aggregate national land value is calculated by summing the total opportunity cost within each imputed dataset and then taking the arithmetic mean of those sums. For coefficient-level estimates (Section 4.3), Rubin’s Rules combine within-imputation and between-imputation variance to produce pooled standard errors that correctly reflect the uncertainty introduced by imputation. The Grand Mean across the three language implementations is reported as the headline national figure, with per-language pooled estimates retained as a robustness check.

In addition to the MICE-pooled estimate, this thesis reports an *observed-only* aggregate computed on the 11,605 courses with directly measured OSM polygon acreage. This figure does not depend on imputation in any way and serves as a strict lower bound on the national opportunity cost.

## A.7 Honolulu County Pipeline: Cadastral and Zoning Data Integration

Three parcel-level datasets from the City and County of Honolulu are integrated with the national-scale outputs:

1. **Honolulu Cadastral GeoPackage:** Parcel boundary geometries with Tax Map Key (TMK) identifiers covering the entire island of Oahu.
2. **Honolulu Tax Roll CSV:** Parcel-level attributes including assessed land value, assessed building value, zoning classification, and district code.
3. **Honolulu Zoning GeoPackage:** Polygon boundaries for all 18 zoning classifications on Oahu (residential R-3.5 through R-20, agricultural AG-1 and AG-2, preservation P-1 and P-2, federal/military F-1, resort, business, country, and others), distributed in EPSG:3760 (NAD83(HARN) Hawaii Zone 3, US survey feet).

All datasets are reprojected to EPSG:5070 to maintain spatial consistency with the national pipeline.

The Phase 2 OSM golf course polygons for Oahu are spatially intersected with the Honolulu cadastral parcels to identify which TMKs constitute each golf course. This produces a list of 1,072 unique TMK parcels comprising the Oahu golf footprint, covering 8,564 acres of total parcel area. A clarification is warranted regarding the two acreage figures that appear in this thesis for Oahu golf land: the 8,564-acre figure is the gross cadastral area of the 1,072 TMK parcels that contain golf use, including non-golf portions of mixed-use parcels (military buffer zones, undeveloped land within shared TMKs, accessory non-recreational acreage). The 6,066-acre figure reported in Section 4.4.3 is the OSM polygon footprint clipped to the actual golf-course geometry within those parcels — the on-the-ground golf footprint after subtracting parcel area not in golf use. The Preservation Paradox

decomposition in Section 4.4.3 uses the 6,066-acre clipped footprint because that figure represents the actual recreational land at issue; the 8,564-acre cadastral figure is reported in this appendix to document the full parcel-level scope of the intersection.

## A.8 Phase 5a: Pilot Validation Against Tax Assessments

The pilot validation track manually compares model HBU estimates against official tax assessment values for six high-profile Hawaii comparator courses spanning three counties (Honolulu, Maui, and Hawaii). For each course, the model estimate is the per-course opportunity cost from the national pipeline (Phase 1 baseline value per acre  $\times$  Phase 2 OSM acreage, MICE-pooled where acreage was imputed), and the assessed value is the published 2022 land assessment from the relevant county tax office. The model-to-assessed ratio is computed for each course, and the ratios are reported across urban and rural settings to characterize the systematic relationship between the two value concepts. The Honolulu courses anchor the model-to-assessed validation for the Honolulu micro-study reported in Section 4.4; the Maui and Hawaii comparator observations are retained as forward-looking data points whose cross-county interpretation is reserved for the planned journal extension to a stratified multi-county validation.

## A.9 Phase 5b: Full TMK and Zoning Analysis

The full pipeline track processes all 33 deduplicated golf courses on Oahu in six steps:

1. *OSM polygon extraction.* Restrict the Phase 2 OSM polygons to an Oahu bounding box; compute point-in-polygon match rates for the Phase 1 course list.
2. *Parcel intersection.* Spatially intersect OSM golf polygons with the Honolulu cadastre to produce a TMK list with associated golf-overlap acreage per parcel.
3. *Economic validation.* Pool  $M = 100$  imputed economic datasets per language via Rubin’s Rules; deduplicate spatially adjacent parcels (500-meter cap); compute the pooled Oahu aggregate opportunity cost. Confidence intervals reported at the 99% level by default; 95% intervals retained in the replication package.
4. *TMK diagnostic merge.* Verify that the TMK identifiers from the cadastral GeoPackage match the format used in the Honolulu tax roll CSV.
5. *Geographic concentration.* Aggregate parcel counts by Honolulu zone code (1–9) to identify district-level concentration.
6. *Zoning intersection.* Spatially intersect each golf parcel with the Honolulu zoning layer; assign each parcel its dominant zoning class via largest-overlap area; compute the share of total Oahu golf acreage in each zoning class, and the share of each zoning class occupied by golf.

Cross-language parity is again checked at each step. A documented discrepancy in the P-1 (Restricted Preservation) intersection acreage between R’s `sf::st_intersection` (523.5 acres) and Python/Julia (744.6 acres) is attributed to differential handling of multipolygon self-intersections and shared boundary precision between R’s `sf` package and the GEOS-backed implementations in Python’s `shapely/geopandas` and Julia’s `LibGEOS`. The mechanism does not propagate to the dominant-zone assignment, which is computed by largest-overlap area and is invariant to the differential precision of the smaller intersection components. The Python/Julia values are reported as canonical because two independent implementations converge on the same figure; the R value is reported as a cross-check.

The output of Phase 5b is a complete characterization of the Oahu golf footprint at parcel resolution: how much land is involved, where it sits geographically, what zoning class governs it, what the aggregate opportunity cost is under the unrestricted HBU framework, and what share of that cost is bounded by current zoning under the legally-permissible HBU framework defined in Section 2.3.

## B Appendix B: Project Genesis

This study originated in Honolulu, where the housing pressures described in the national literature are not abstractions. As a graduate student at the University of Hawai‘i at Mānoa, I encountered a state with one of the nation’s most acute housing shortages alongside a striking density of golf facilities — 33 distinct courses on the island of Oahu alone, occupying over 6,000 acres of zoned and developable land. The City and County of Honolulu maintains comprehensive parcel-level cadastral and zoning data at the Tax Map Key (TMK) resolution, providing an unusual opportunity to test whether the national-scale opportunity cost framework holds up against ground-truth municipal records.

This local context shaped the design of the project in two ways. First, it suggested that a purely national aggregate, however carefully constructed, risks being dismissed as a theoretical exercise. A defensible empirical claim about the value of golf course land requires anchoring against the assessed values that municipalities actually use to tax these parcels. Second, it suggested that Hawaii itself — with high-value urban courses, rural courses on the neighbor islands, and a regulatory landscape dominated by Preservation and Federal/Military zoning — is a useful microcosm for the broader American pattern. The two scales are complementary: the national dataset establishes the magnitude and geographic structure of the spatial misallocation, while the Honolulu micro-study tests whether that magnitude survives contact with cadastral reality.

## C Appendix C: Thesis Roadmap

This appendix records the section-by-section reading guide for the thesis. The body is organized as follows. **Section 1** introduces the spatial-misallocation framing, the four strands of literature this study bridges, the research questions, the contributions, and the Highest and Best Use (HBU) research framework that organizes the empirical work. **Section 2** formalizes the theoretical framework, translating HBU into a neoclassical spatial model and introducing the legally-permissible HBU concept used in the micro-study. **Section 3** details the two-scale research design, including the national pipeline (Phases 1–4) and the Honolulu County pipeline (Phase 5). **Section 4** presents the empirical results across both scales, with the Hawaii validation as the empirical anchor. **Section 5** discusses the macroeconomic and policy implications, including the failure of Coasian bargaining, the validation gap, the zoning-bounded HBU, methodological limitations, and renewable-energy counterfactuals. **Section 6** concludes with a summary of findings and directions for future economic research. **Appendix A** contains the full technical narrative of the pipeline implementation. **Appendix B** records the personal context that motivated the research design. **Appendix C** is this roadmap.

## References

- Appraisal Institute (2020). *The Appraisal of Real Estate*. 15th. Chicago, IL: Appraisal Institute. ISBN: 978-1935328780.
- Buuren, Stef van (2018). *Flexible imputation of missing data*. eng. Second edition. International and comparative criminal justice series. Boca Raton, FL: Chapman, Hall/CRC, an imprint of Taylor, and Francis. ISBN: 0-429-96035-2.
- Cheshire, Paul and Stephen Sheppard (2002). “The welfare economics of land use planning”. eng. In: *Journal of urban economics* 52.2, pp. 242–269. ISSN: 0094-1190.
- Cho, Seong-Hoon et al. (2009). “Spatial and temporal variation in the housing market values of lot size and open space”. eng. In: *Land economics* 85.1, pp. 51–73. ISSN: 0023-7639.
- Coase, R. H. (2013). “The Problem of Social Cost”. eng. In: *The Journal of law & economics* 56.4, pp. 837–877. ISSN: 0022-2186.
- Davis, Morris A., William D. Larson, et al. (2021). “The price of residential land for counties, ZIP codes, and census tracts in the United States”. eng. In: *Journal of monetary economics* 118, pp. 413–431. ISSN: 0304-3932.
- Davis, Morris A. and Michael G. Palumbo (2008). “The price of residential land in large US cities”. eng. In: *Journal of urban economics* 63.1, pp. 352–384. ISSN: 0094-1190.
- Ellickson, Robert C (2022). “STALE REAL ESTATE COVENANTS”. eng. In: *William and Mary law review* 63.6, pp. 1831–. ISSN: 0043-5589.
- Fischel, William A. (2015). *Zoning rules : the economics of land use regulation*. eng. 1st ed. Cambridge, Massachusetts: Lincoln Institute of Land Policy. ISBN: 1-55844-329-0.
- Galai, Dan and Orly Sade (2006). “The “Ostrich Effect” and the Relationship Between the Liquidity and the Yields of Financial Assets”. In: *The Journal of Business* 79.5, pp. 2741–2759.
- Glaeser, Edward, National Bureau of Economic Research, and Joseph Gyourko (2017). *The Economic Implications of Housing Supply*. eng. NBER working paper series no. w23833. Cambridge, Mass: National Bureau of Economic Research.
- Glaeser, Edward L., Joseph Gyourko, and Raven Saks (2005). “Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices”. eng. In: *The Journal of law & economics* 48.2, pp. 331–369. ISSN: 0022-2186.
- Gyourko, Joseph, National Bureau of Economic Research, and Jacob Krimmel (2021). *The Impact of Local Residential Land Use Restrictions on Land Values Across and Within Single Family Housing Markets*. eng. NBER working paper series no. w28993. Cambridge, Mass: National Bureau of Economic Research.
- Karlsson, Niklas, George Loewenstein, and Duane Seppi (2009). “The ostrich effect: Selective attention to information”. In: *Journal of Risk and Uncertainty* 38.2, pp. 95–115.
- Korngold, Gerald (2024). “Repealing Single-Family Zoning Is Not Enough: A Proposal For Removing Existing Parallel Private Covenants For Violating Public Policy”. eng. In: *Missouri law review* 89.1, pp. 1–. ISSN: 0026-6604.
- Limehouse, Frank F, Robert E McCormick, and Melissa M Yeoh (2012). “The private provision of public goods: an analysis of homes on golf courses”. eng. In: *The Journal of private enterprise* 27.2, pp. 103–120. ISSN: 0890-913X.
- Limehouse, Frank F., Peter C. Melvin, and Robert E. McCormick (2010). “The Demand for Environmental Quality: An Application of Hedonic Pricing in Golf”. eng. In: *Journal of sports economics* 11.3, pp. 261–286. ISSN: 1527-0025.
- Peterson, Kimberly and Chris Peterson (Oct. 2024). *Utility-scale U.S. solar electricity generation continues to grow in 2024*. Today in Energy, U.S. Energy Information Administration. Accessed: 2026-05-18. URL: <https://www.eia.gov/todayinenergy/detail.php?id=63324>.

- Rosen, Sherwin (1974). “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition”. eng. In: *The Journal of political economy* 82.1, pp. 34–55. ISSN: 0022-3808.
- Rubin, Donald B. (1987). *Multiple imputation for nonresponse in surveys*. eng. Wiley series in probability and mathematical statistics. Applied probability and statistics. New York ; Wiley. ISBN: 047108705X.
- Solar Energy Industries Association (n.d.). *Land Use & Solar Development*. SEIA initiatives. Cites 5–7 acres per MW land-use intensity for utility-scale solar. Accessed: 2026-05-18. URL: <https://seia.org/initiatives/land-use-solar-development/>.
- U.S. Census Bureau (2021a). *Delaware: Census Bureau Profile*. U.S. Department of Commerce. 2020 Decennial Census, land area 1,948.0 sq mi. Accessed: 2026-05-18. URL: <https://data.census.gov/profile/Delaware?g=040XX00US10>.
- (2021b). *Rhode Island: Census Bureau Profile*. U.S. Department of Commerce. 2020 Decennial Census, land area 1,033.81 sq mi. Accessed: 2026-05-18. URL: [https://data.census.gov/profile/Rhode\\_Island?g=040XX00US44](https://data.census.gov/profile/Rhode_Island?g=040XX00US44).
- U.S. National Renewable Energy Laboratory (June 2024). *Land Requirements for 100% Solar Electricity Scenario*. As reported in PV Magazine USA. NREL estimate: ~10 million acres needed if U.S. met all electricity needs via solar. Accessed: 2026-05-18. URL: <https://pv-magazine-usa.com/2024/06/21/utility-scale-solar-development-good-planning-makes-good-neighbors/>.
- Weinand, Jann Michael et al. (2025). “Countries across the world use more land for golf courses than wind or solar energy”. eng. In: *Environmental Research Communications* 7.2, pp. 21012–. ISSN: 2515-7620.